

Reality, Causality, and Quantum Theory



John-Mark A. Allen
St. John's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Hilary Term 2017

“I daresay you haven’t had much practice. When I was your age, I always did it for half-an-hour a day. Why, sometimes I’ve believed as many as six impossible things before breakfast.”

– *The White Queen*¹

“When you are a Bear of Very Little Brain, and you Think of Things, you find sometimes that a Thing which seemed very Thingish inside you is quite different when it gets out into the open and has other people looking at it.”

– *Pooh Bear*²



¹*Through the Looking Glass*, Lewis Carroll

²*The House at Pooh Corner*, A. A. Milne

To

Mary

Peter

John

Caroline

Tim

Phil

Acknowledgements

Down the long wall of the main stairwell in the Scottish National Gallery of Modern Art One there are many names, neatly arranged in unassuming columns of clean Helvetica on the white gallery wall. They appear as a very standard (if impressively long) list of benefactors to the gallery. Look closer and you can recognise some famous names of modern art in the crowd. Regain some awareness of your surroundings and you notice the label identifying it as “List of Names (Random)” by Douglas Gordon, an ongoing piece listing all of the people the artist can remember meeting.

Thinking about writing these acknowledgements, my mind keeps coming back to that wall. Acknowledgements tend to end up looking like a long list of loosely related and largely unfamiliar names, perhaps sprinkled with some you recognise. But mostly I end up thinking of all the people Gordon must have forgotten to include. It seems horrible and cruel to forget people, yet failing to mention everyone is inevitable. So I guess I’m starting with an apology—a long-winded and somewhat pretentious apology to anybody who has helped me at all in the making of this thesis but I have failed to mention here. I’m sorry and I am certainly grateful.

Foremost thanks most definitely go to Jon Barrett. As a supervisor, collaborator, source of encouragement, fantastic explanations, and good humour he has been wonderful throughout. In particular should be noted his clarity and patience when explaining an open problem to me for the fourth time (when my memory refuses to play along) and his encouragement in finding further research opportunities and collaborations in Oxford, Waterloo, and elsewhere.

Dom Horsman also deserves more than a special mention here. He has selflessly kept me sane and on-track in many ways over many conversations. He has been both amazingly practical and delightfully irreverent, each exactly when they were needed. I have done nothing to deserve such a wonderful and effective pair to guide me through my time at Oxford and I am forever grateful.

They are far from the only people from the Quantum Group who need thanking, however. In particular Ciarán Lee (ah, yeah, it’s been grand), Matty Hoban, John Selby, and Stefano Gogioso have been exceptionally accommodating to being distracted by sometimes-interesting sometimes-banal questions. Samson Abramsky, Bob Coecke, Carlo Maria Scandolo, Nathan Walk, Niel de Beaudrap, Subhayan Roy Moulik, Robin Lorenz, Jamie Vicary, Shane Mansfield, Sean Tull, Miriam Backens, Will Zeng, Brendan Fong, Aleks Kissinger, Chris Heunen, Dan Marsden, Sina Salek, and Ognyan Oreshkov all need to be thanked for interesting and productive conversations, advice, and teaching.

I have also benefited from working closely with several collaborators, from whom I have learned much. These include Owen Maroney, Stefano Gogioso, Ciarán Lee, Jon Barrett, and Dom Horsman in Oxford. I’m indebted to Rob Spekkens, Matt Pusey, Matt Leifer, and Elie Wolfe at Perimeter Institute and especially Lucien Hardy for hosting my

visit there. Special thanks also to Lída del Rio in Zürich for hosting a productive and enjoyable visit.

Back in the classical world, there are many more people who have been a joy to spend time with. They have helped me to keep my head and humour throughout my DPhil process in Oxford. Alma Brodersen has been unyieldingly delightful and encouraging. Solja Höft is absolutely brilliant and a magnificent friend. Jan Cosgrave has been determined to keep me topped up with sleep and dessert wine (dude, with Ross Haines, also deserves utmost respect for slogging all the way through this thesis in search of typos). Sally Le Page has been both inspiring and comforting. The list of wonderful people who've kept me sane and happy goes on and on with Sarah Penington, Marina Lambrakis, Jon Templeman, Katherine Fender, Laura Grima, Kate Radford, Ali Ward, Sam Forbes, Tim Bourns, Max Emmerich, Ellie Milnes-Smith, Martin Bittner, Chris Arran, and Alan Percy. In fact, many people from the St. John's MCR could happily fit on this list. One group that deserves particular note is Oriel College Choir including David Maw, Claire Lowe, Alasdair Cameron, Johanna Hockmann, and Lizzie Searle to name but a few, because there's nothing like depressed Tudor Catholic to brighten up your week.

The list of wonderful friends who have helped in many ways extends well beyond Oxford. In particular, Ed Whittle, Harriet Cook, Chris Hallam, Kenny Masters, Kieran Franklin, Katie Dooley, Emma Wilson, Emily Baker, Becca Lee, Josh Rhodes, Bonnie Barker, Darryl Hoving, Lizzie Briggs, Giulia Postal, Katharine Elliot, Sophie Reed, Elena Teh, Heather Parsons, Olivia Cleary, and Jazzy Ramsay Gray should all be thanked. Tort, Richard, Emma, and Madeleine Olver have all been hugely generous, fun, and kind.

Of course there are several institutions which have practically enabled this thesis. The Department of Computer Science, especially Julie Sheppard and Destiny Chen, has been very supportive in many ways. St. John's College has acted as a home and source of great support. The Engineering and Physical Science Research Council (and, by extension, the UK taxpayer) funded this research, for which I am forever grateful. I also received generous funding support from ETH Zürich and the Perimeter Institute for Theoretical Physics for visits.

My family, to whom this thesis is dedicated, quite literally cannot be thanked enough. Brothers, uncles, aunts, cousins, I only become increasingly aware of how absurdly lucky I am to have you. Most of all, my parents and grandparents have always supported me with a compassion and grace that is difficult to describe. Thank you all, for everything that has led up to this thesis and everything that will follow.

Finally, Catherine. There is certainly nobody who understands the personal content of this work better than you. Thank you for your patience and your impatience. Your understanding and your refusal to understand. Your support and your drive. Your excitement and your confusion. Your celebrations and your commiserations. Your presence. Your letters. Your thoughts. Your belief. Thank you.

Abstract

Quantum theory describes our universe incredibly successfully. To our classically-inclined brains, however, it is a bizarre description that requires a re-imagining of what fundamental reality, or “ontology”, could look like. This thesis examines different ontological features in light of the success of quantum theory, what it requires, and what it rules out. While these investigations are primarily foundational, they also have relevance to quantum information, quantum communication, and experiments on quantum systems.

The way that quantum theory describes the state of a system is one of its most unintuitive features. It is natural, therefore, to ask whether a similarly strange description of states is required on an ontological level. This thesis proves that almost all quantum superposition states for $d > 3$ dimensions must be real—that is, present in the ontology in a well-defined sense. This is a strong requirement which prevents intuitive explanations of the many quantum phenomena which are based on superpositions. A new theorem is also presented showing that quantum theory is incompatible with macro-realist ontologies, where certain physical quantities must always have definite values. This improves on the Leggett-Garg argument, which also aims to prove incompatibility with macro-realism but contains loopholes. Variations on both of these results that are error-tolerant (and therefore amenable to experimentation) are presented, as well as numerous related theorems showing that the ontology of quantum states must be somewhat similar to the quantum states themselves in various specific ways. Extending these same methods to quantum communication, a simple proof is found showing that an exponential number of classical bits are required to communicate a linear number of qubits. That is, classical systems are exponentially bad at storing quantum data.

Causal influences are another part of ontology where quantum theory demands a revision of our classical notions. This follows from the outcomes of Bell experiments, as rigorously shown in recent analyses. Here, the task of constructing a native quantum framework for reasoning about causal influences

is tackled. This is done by first analysing the simple example of a common cause, from which a quantum version of Reichenbach's principle is identified. This quantum principle relies on an identification of quantum conditional independence which can be defined in four ways, each naturally generalising a corresponding definition for classical conditional independence. Not only does this allow one to reason about common causes in a quantum experiments, but it can also be generalised to a full framework of quantum causal models (mirroring how classical causal models generalise Reichenbach's principle). This new definition of quantum causal models is illustrated by examples and strengthened by it's foundation on a robust quantum Reichenbach's principle. An unusual, but surprisingly fruitful, setting for considering quantum ontology is found by considering time travel to the past. This provides a testbed for different ontological concepts in quantum theory and new ways to compare classical and quantum frameworks. It is especially useful for comparing computational properties. In particular, time travel introduces non-linearity to quantum theory, which brings (sometimes implicit) ontological assumptions to the fore while introducing strange new abilities. Here, a model for quantum time travel is presented which arguably has fewer objectionable features than previous attempts, while remaining similarly well-motivated. This model is discussed and compared with previous quantum models, as well as with the classical case.

Together, these threads of investigation develop a better understanding of how quantum theory affects possible ontologies and how ontological prejudices influence quantum theory.

Contents

Related Publications	v
Notation	vii
1 Introduction	1
1.1 “Ontology”, “Epistemic”, and “Causality”	4
1.2 Thesis Overview	5
2 Ontology of Quantum States and Superpositions	9
2.1 Theorising About State Ontology	9
2.1.1 The Ontological Models Framework	10
2.1.2 The Desire for Simpler Ontologies	14
2.1.3 Classifying Ontologies	17
2.1.4 Previous Ontology Theorems	20
2.1.5 Limitations and Loopholes	21
2.1.6 Chapter Overview	23
2.2 The Asymmetric Overlap	23
2.2.1 Properties of the Asymmetric Overlap	25
2.2.2 Anti-distinguishability	29
2.3 Must Superpositions be Real?	32
2.3.1 Distinguishing Ontic Superpositions	32
2.3.2 Almost All Superpositions are Real	32
2.4 Must any Specific States be Real?	35
2.4.1 State-Specific Ontology	35
2.4.2 No States Can Be ψ -Epistemic	37
2.4.3 Towards Error-Tolerance and Experiments	39
2.4.4 Justifying Preparation Contextuality for Stabiliser Unitaries	41
2.5 Communication Bounds from Ontology Results	42
2.5.1 Relating Ontology and Communication	43

2.5.2	A Simple Exponential Bound	45
2.5.3	Generating Further Bounds	46
2.6	Summary	47
3	Ontology and Macro-realism	51
3.1	The Meaning of Macro-realism	51
3.1.1	Introducing Macro-realism	52
3.1.2	Defining Macro-realism	53
3.2	The Leggett-Garg Inequalities	54
3.2.1	Outline of the Argument	54
3.2.2	Sub-classes of Macro-realism	55
3.2.3	Macro-realism for Quantum Systems	58
3.2.4	Loopholes in the Leggett-Garg Argument	58
3.3	A Stronger Theorem Against Macro-realism	60
3.3.1	Theorem	61
3.3.2	Relation to Other Ontology Results	61
3.4	An Error-Tolerant Argument	62
3.4.1	Error-Tolerance and the ϵ -Asymmetric Overlap	63
3.4.2	Properties of the ϵ -Asymmetric Overlap	64
3.4.3	Error-Tolerant Results	68
3.5	Summary and Discussion	73
4	Quantum Conditional Independence and Quantum Common Causes	77
4.1	The Need for Quantum Conditional Independence	77
4.1.1	The Reality of Causal Influences	77
4.1.2	Causation in Bell Experiments	78
4.1.3	Metaphysically-Neutral Quantum Conditional Independence	80
4.2	Classical Conditional Independence and Reichenbach's Principle	80
4.2.1	Reichenbach's Principle in Two Parts	80
4.2.2	Justifying the Quantitative Part	81
4.3	Quantum Conditional Independence and Quantum Reichenbach	84
4.3.1	The Choi-Jamiołkowski Isomorphism	84
4.3.2	Justifying Quantum Reichenbach	85
4.3.3	Alternative Expressions for Quantum Conditional Independence	88
4.3.4	Proving Propositions 4.1 and 4.2	89
4.3.5	Circuits for Common Cause Channels	94
4.3.6	Examples	97

4.3.6.1	Generic Unitary Transformations	98
4.3.6.2	Coherent and Incoherent Copies	98
4.3.6.3	Bell Experiments	101
4.3.7	Generalisation to k Outputs	103
4.4	Summary	105
5	Quantum Causal Models	107
5.1	The Need for Quantum Causal Models	107
5.1.1	The Reality of Causal Influences	107
5.1.2	Causation in General Quantum Experiments	108
5.1.3	Previous Approaches to Quantum Causal Models	109
5.2	Classical Causal Models Generalise Reichenbach's Principle	111
5.2.1	Classical Causal Models in Two Parts	111
5.2.2	Justifying the Markov Condition	113
5.3	Quantum Causal Models Generalise Quantum Reichenbach	114
5.3.1	Defining Quantum Causal Models	114
5.3.2	Bayesian vs. do-conditionals	116
5.3.3	Making Predictions	117
5.3.4	Classical Limits	118
5.3.5	Examples	120
5.3.5.1	Confounding Common Cause	120
5.3.5.2	Simple Bayesian Updating with a Complete Common Cause	122
5.4	Summary and Discussion	124
6	Causal Loops: Time Travel to the Past	129
6.1	Closed Timelike Curves and the Ontology of Time Travel	129
6.1.1	Why Study Time Travel?	129
6.1.2	Closed Timelike Curves and General Relativity	131
6.2	Modelling Time Travel with Quantum Circuits	132
6.2.1	The Standard Form of Time Travel Circuits	132
6.2.2	Time Travel Paradoxes and the Classical Model	133
6.3	Quantum Time Travel Models, Non-linearity, and Ontology	135
6.3.1	Overview of D-CTCs	135
6.3.2	Overview of P-CTCs	137
6.3.3	The Problems of Non-linearity	139
6.3.4	The Role of Ontology	141

6.4	Alternative Time Travel Models	144
6.4.1	Integrating over Quantum States	144
6.4.2	Desiderata	145
6.4.3	Some Alternative Models	147
6.4.4	The Uniqueness Ambiguity and Epistemic Reasoning	149
6.4.5	T-CTCs	150
6.4.5.1	Simplification of the T-CTC Equation of Motion	151
6.4.5.2	The P Operator	153
6.4.5.3	Paradoxes	153
6.4.5.4	Computation	154
6.4.5.5	Mixed States and Non-linearity	155
6.4.5.6	Relation to Other Alternatives	157
6.5	Summary and Discussion	157
7	Conclusions and Further Work	161
7.1	State Ontology and Macro-realism	161
7.2	Quantum Causality	165
7.3	Time Travel and CTCs	168
7.4	To Conclude	170
	Appendix A Proofs of Theorems 2.3 and 2.4	171
	Appendix B Proof of Theorem 2.5	175
	Bibliography	182

Related Publications

- [All14] **JMAA**, *Treating Time Travel Quantum Mechanically*. *Physical Review A* **90**(4), 042107 (2014). arXiv:1401.4933 [quant-ph]
- [All16] **JMAA**, *Quantum Superpositions Cannot be Epistemic*. *Quantum Studies: Mathematics and Foundations* **3**(2), 161–177 (2016). arXiv:1501.05969 [quant-ph]
- [ABH⁺17] **JMAA**, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, *Quantum Common Causes and Quantum Causal Models*. *Physical Review X* **In preparation** (2017). arXiv:1609.09487 [quant-ph]
- [AMG16] **JMAA**, O. J. E. Maroney, and S. Gogioso, *A Stronger Theorem Against Macro-realism* (2016). arXiv:1610.00022 [quant-ph]

Notation

Most of the notation used in this thesis is based on standard usage in quantum foundations or physics in general. In order to streamline the text, some of the most common notation used is summarised here. If the reader comes across an unfamiliar symbol in the text, this should be their first port of call for a definition.

The symbol $\stackrel{\text{def}}{=}$ is used to indicate that the expression is a definition of the item on the left. Read “is defined as being equal to”.

For logarithms, \log is used for the base-two logarithm (as in information theory), while \ln is used for base- e natural logarithms.

Quantum Systems and States

Quantum systems will normally be given capital Latin letters, such as A .

- \mathcal{H}_A is the *Hilbert space* (bounded inner-product complex vector space) corresponding to system A .
- $d_A \stackrel{\text{def}}{=} \dim \mathcal{H}_A$ is the dimension of the system A .
- $\mathbb{1}_A$ is the identity operator on \mathcal{H}_A .
- $\mathcal{H}_A^* = (\mathcal{H}_A)^*$ is the dual Hilbert space to \mathcal{H}_A , with dimension $d_{A^*} = d_A$ and identity $\mathbb{1}_{A^*}$.
- $\mathcal{P}(\mathcal{H}_A) \stackrel{\text{def}}{=} \{|\psi\rangle \in \mathcal{H}_A : \|\psi\| = 1, |\psi\rangle \cong e^{i\theta}|\psi\rangle, \forall \theta \in \mathbb{R}\}$ is the set of *pure physical states* in A , viz. unit vectors in \mathcal{H}_A where vectors equal up to a global phase are considered equivalent.
- $\mathcal{D}(\mathcal{H}_A)$ is the set of normalised *density operators* (i.e. mixed states) of A , viz. positive Hermitian trace-one operators ρ on \mathcal{H}_A .
- A *superposition* state with respect to some given orthonormal basis \mathcal{B} of \mathcal{H} is any $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ such that $|\psi\rangle \notin \mathcal{B}$.

- AB , where A and B are each distinct quantum systems, is the bipartite contemporary system comprised of the two with Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$.
- If a state for a multipartite system AB is given by, e.g, ρ_{AB} , then the same symbol with omitted indices indicates a partial trace of that state, such as $\rho_A = \text{Tr}_B \rho_{AB}$.
- Products of states that do not share all of their subsystems are taken to implicitly include identities on those systems, e.g. $\rho_A \rho_B \stackrel{\text{def}}{=} (\rho_A \otimes \mathbb{1}_B)(\mathbb{1}_A \otimes \rho_B)$ and $\rho_{AB} \rho_{BC} \stackrel{\text{def}}{=} (\rho_{AB} \otimes \mathbb{1}_C)(\mathbb{1}_A \otimes \rho_{BC})$.
- A *measurement* M of a quantum system consists of a set of outcomes $E \in M$ which can be obtained when the measurement is performed.
- A *basis measurement* $M = \mathcal{B} = \{|i\rangle\}_i$ is a measurement where the outcomes are pure states $|i\rangle$ which form an orthonormal basis \mathcal{B} for the Hilbert space of system. The probability of obtaining outcome $|i\rangle$ when the system is in state ρ is $\langle i|\rho|i\rangle$.
- A *POVM measurement* $M = \{E_i\}_i$ is a measurement where the outcomes are Hermitian positive definite operators on the Hilbert space that form a partition of unity $\sum_i E_i = \mathbb{1}$. The probability to obtain outcome E_i when the system is in state ρ is $\text{Tr}(\rho E_i)$.
- A *von Neumann measurement* $M = \{E_i\}_i$ is a POVM measurement where all outcomes are projectors $E_i E_j = \delta_{ij} E_i$. If all operators are one-dimensional projectors then this is equivalent to a basis measurement.
- A *quantum channel* from A to B , $\mathcal{E}_{B|A}$, is a completely-positive trace-preserving (CPTP) map from states in $\mathcal{D}(\mathcal{H}_A)$ to states in $\mathcal{D}(\mathcal{H}_B)$.
- A *quantum instrument* from A to B , $\{\mathcal{E}_{B|A}^k\}_k$, is a set of completely-positive trace-non-increasing maps from states in $\mathcal{D}(\mathcal{H}_A)$ to sub-normalised density operators on \mathcal{H}_B . Each map is labelled by a classical outcome k and they sum to a quantum channel. These represent general operations on a quantum system that include transformations and measurements. The interpretation is that one such map is applied and the corresponding outcome observed, with probabilities given by the trace of the corresponding output state. The channel obtained by summing the maps outputs the proper mixed state obtained by ignoring this classical outcome.
- $\mathcal{S}_{|\psi\rangle}$ is the set of *stabiliser unitaries* for the pure state $|\psi\rangle \in \mathcal{P}(\mathcal{H})$. That is, $\mathcal{S}_{|\psi\rangle} \stackrel{\text{def}}{=} \{U \text{ on } \mathcal{H} : U|\psi\rangle = |\psi\rangle\}$. In other words, $\mathcal{S}_{|\psi\rangle}$ is the stabiliser subgroup of the unitary group over \mathcal{H} with respect to $|\psi\rangle$.

Quantum Circuit Diagrams

A standard box-and-wire notation for quantum circuits will be used, specifically with the following conventions.

- Circuits proceed from bottom to top.
- Wires represent systems and may or may not be labelled with the appropriate symbol or Hilbert space.
- Boxes represent unitary or CPTP evolutions from (the tensor product of) input systems (below) to output systems (above) and are labelled with the appropriate operation.
- Semi-circles represent state preparations of the outgoing systems (above) and are labelled with the appropriate state.
- The unitary CNOT gate is illustrated as $\dot{\text{---}}\text{---}\oplus$ where the solid dot is on the control system and the \oplus is on the target system.
- The unitary SWAP gate may be illustrated as $\ast\text{---}\ast$.

Ontological Models and Overlap Measures

Ontological models are defined in Sec. 2.1.1 with the notation summarised here for convenience.

- Λ is the set of *ontic states* $\lambda \in \Lambda$.
- Σ is the sigma-algebra such that (Λ, Σ) is a measurable space—the *ontic state space*.
- Δ_P is the set of *preparation measures* $\mu \in \Delta_P$ for a preparation P .
- $\Delta_{|\psi\rangle}$ is therefore the set of preparation measures for all preparation methods that produce quantum state $|\psi\rangle$ if the system is a quantum system.
- $\mu \in \Delta_P$ gives the probability $\mu(\Omega)$ for the resulting ontic state being in $\Omega \in \Sigma$ for some preparation method of P .
- Γ_T is the set of *stochastic maps* $\gamma \in \Gamma_T$ for a transformation T of the system.
- Γ_U is therefore the set of stochastic maps for all transformation methods corresponding to a unitary U if the system is a quantum system.

- $\gamma \in \Gamma_T$ gives the probability $\gamma(\Omega|\lambda)$ for the resulting ontic state being in $\Omega \in \Sigma$ given that some method for T was applied to a system in ontic state $\lambda \in \Lambda$.
- $\mu \overset{\gamma}{\rightsquigarrow} \nu$ indicates that if a system is prepared according to μ and then a transformation is applied according to γ , then ν is the preparation measure that describes that composite process, Eq. (2.2).
- Ξ_M is the set of *response functions* $\mathbb{P}_M \in \Xi_M$ for a measurement M .
- If $M = \{|i\rangle\}_i$ is a basis measurement then Ξ_M is therefore the set of response functions for all methods of performing the measurement M on a quantum system.
- $\mathbb{P}_M \in \Xi_M$ gives the probability $\mathbb{P}_M(E|\lambda)$ that the outcome $E \in M$ is obtained given that some method for measuring M was performed on a system in ontic state $\lambda \in \Lambda$.
- Similarly, $\mathbb{P}_M(E|\mu) \stackrel{\text{def}}{=} \int_{\Lambda} d\mu(\lambda) \mathbb{P}_M(E|\lambda)$ is the probability of obtaining outcome $E \in M$ given that some method for measuring M was performed on a system prepared according to μ .
- $\varpi(\cdot|\cdot)$ is the *asymmetric overlap* as defined and explained in Sec. 2.2 [AMG16, All16, Bal14, LM13, Mar12].
- $\omega(\cdot, \cdot)$ is the *symmetric overlap* as defined and explained in Sec. 2.4.3 [AMG16, All16, BCLM14, Bra14, Lei14b, Lei14a, Mar12].
- $\varpi_{\epsilon}(\cdot|\cdot)$ for some $\epsilon \in [0, 1)$ is the ϵ -*asymmetric overlap* as defined and explained in Sec. 3.4.1. 0-asymmetric overlap is the normal asymmetric overlap.
- \bar{k} is a mapping from measurable functions $g : \Lambda \rightarrow [0, 1]$ to measurable sets over Λ . $\bar{k}(g) \stackrel{\text{def}}{=} \ker(1 - g) \in \Sigma$ as defined in Def. 2.1.
- \bar{k}_{ϵ} , for $\epsilon \in [0, 1)$ generalises \bar{k} as defined in Def. 3.1.

Classical Random Variables

Classical random variables are normally given capital Latin letters, such as X . Values that the variable can take are normally given by the corresponding lower-case Latin letter, such as x . Random variables are assumed to be discrete unless otherwise indicated.

- $\mathbb{P}(X = x)$ is the probability that random variable X has value x . $\mathbb{P}(X)$ is the probability distribution for X . Alternatively, $\mathbb{P}(x) \stackrel{\text{def}}{=} \mathbb{P}(X = x)$ is used as shorthand where there is no ambiguity in doing so.
- $\mathbb{P}(X, Y)$ is the joint probability distribution for random variables X and Y . $\mathbb{P}(X, Y = y)$ is the distribution of probabilities that $X = x \wedge Y = y$ for all possible x .
- $\mathbb{P}(X|Y)$ is the conditional probability distribution for X given that Y takes any given value.

Entropies and Information

Classical *Shannon entropies* are usually denoted with H . Quantum *von Neumann entropies* are usually denoted with S . Recall that \log is used for base-two logarithms. The definitions here assume that all random variables are discrete and all quantum systems are finite-dimensional.

Note that the definitions of quantum entropies can all be applied to both quantum systems and quantum states, depending on context.

- $H(X)$ is the classical *entropy* of a classical random variable, defined $H(X) \stackrel{\text{def}}{=} -\sum_x \mathbb{P}(x) \log \mathbb{P}(x)$.
- $H(X, Y)$ is the classical *joint entropy*, defined $H(X, Y) \stackrel{\text{def}}{=} -\sum_{x,y} \mathbb{P}(x, y) \log \mathbb{P}(x, y)$.
- $H(X|Y) \stackrel{\text{def}}{=} H(X, Y) - H(Y)$ is the classical *conditional entropy*.
- $I(X : Y) \stackrel{\text{def}}{=} H(X) + H(Y) - H(X, Y)$ is the classical *mutual information*.
- $I(X : Y|Z) \stackrel{\text{def}}{=} H(X, Z) + H(Y, Z) - H(Z) - H(X, Y, Z) = H(X|Z) - H(X|Y, Z)$ is the classical *conditional mutual information*.
- $S(\rho)$ is the quantum *entropy* of a state ρ (equivalently, of a quantum system in state ρ), defined $S(\rho) = -\text{Tr}(\rho \log \rho)$.
- $S(A)$ is the *entropy* of a quantum system A , defined $S(\rho_A)$ where ρ_A is the state of A .
- $S(AB)$ is the quantum *joint entropy* of bipartite system AB , defined $S(\rho_{AB})$ where ρ_{AB} is the state of AB .
- $S(A|B) \stackrel{\text{def}}{=} S(AB) - S(A)$ is the quantum *conditional entropy*.

- $I(A : B) \stackrel{\text{def}}{=} S(A) + S(B) - S(AB)$ is the quantum *mutual information*.
- $I(A : B|C) \stackrel{\text{def}}{=} S(AC) + S(BC) - S(C) - S(ABC) = S(A|C) - S(A|BC)$ is the quantum *conditional mutual information*.
- Quantum and classical mutual informations, despite using the same symbol, can be told apart by the types of system on which they act.

1

Introduction

“... So it will be difficult. But the difficulty, really, is psychological and exists in the perpetual torment that results from your saying to yourself “But how can it be like that?” Which really is a reflection of an uncontrolled, but I say utterly vain, desire to see it in terms of some analogy with something familiar. I will not describe it in terms of an analogy with something familiar. I’ll simply describe it.

...

So that’s the way to look at the lecture—is not to try to understand—well, you have to understand the English, of course. But in any sense in terms of something else. And don’t keep saying to yourself, if you can possibly avoid it, “But how can it be like that?” Because you’ll get down the drain. You’ll get down into a blind alley in which nobody has yet escaped. Nobody knows how it can be like that.” [Fey64]

These words of Richard Feynman, delivered to the audience of a public lecture on quantum theory in 1964, contain both a warning and an invitation. You must not get lost down a blind alley trying to understand quantum theory. But nobody knows how it can be like this; which, to a theorist, is an invitation if ever there was one.

More than fifty years later, it is still true that nobody knows how or why nature is successfully described by quantum theory. Some would certainly claim that they do, but unfortunately such people generally disagree with one another [SBKW10, BH95, FMS14,

DE10]¹. However, the study of quantum theory, especially in responding to Feynman's challenge, has changed significantly. Quantum foundations has bloomed into its own sub-field on the borders of physics, philosophy, mathematics, information science, and computer science. The most productive insights from this field have not come from stumbling into the alley, but by carefully peering in and trying and make out its main contours.

This thesis contributes to that effort while concentrating on *quantum ontology*. In foundations of quantum theory, *ontology* refers to any potential "actual state of affairs" or fundamental description, essentially the *reality* of the physical system. What does quantum theory require or suggest about the behaviour of physics on the most fundamental level? Does the success of quantum theory require certain ways of understanding the actual state of affairs of a system, or does it rule out others? Rather than attempting to give a comprehensive account of an exact quantum ontology, such questions can be carefully drawn out in specific ways, revealing the important features of any plausible or actual ontology.

It is helpful to note that quantum theory can fruitfully be considered as a *framework* rather than a *theory* as such², a framework being understood as a low-level set of rules that theories built from that framework must respect. The classical framework, for example, might be expressed as the facts that: systems occupy exactly one state from a space at any given point, joint states of multiple systems are found by taking the Cartesian product, transformations can move systems between states in such a way that probability is conserved, *etc.* Out of this framework, specific theories such as Maxwellian electromagnetism, statistical mechanics, and Newtonian orbital mechanics can be built. Similarly, the quantum framework can be understood as the Hilbert space structure of system states, the unitary/CPTP (completely-positive trace-preserving) nature of transformations, the Born rule for measurement outcomes *etc.* and out of this, theories such as quantum optics and the standard model can be built. This distinction between framework and theory goes a long way to explaining why quantum theories are uniquely difficult for us to understand compared to any other physical theory: they are all predicated on an unfamiliar non-classical framework. It also allows for easier comparisons between quantum and classical predictions in general by concentrating on what is allowed by these general abstract frameworks, rather than getting lost in the exact physics of specific systems. However,

¹Of course, this is not to say that none of these people are correct, but it certainly shows that most are likely to be incorrect.

²I have heard this point being made by various people, but have been unable to find a citation for it. This is therefore presented with apologies to anybody who may lay claim to its conception.

“quantum theory” will continue to be the term used throughout this thesis since “quantum framework” is clunky and “quantum mechanics” is no more specific.

The mother of the modern approach to quantum ontology is Bell’s theorem [Bel66], developed in response to controversy over hidden variables arguments in quantum theory in general and the EPR argument [EPR35] in particular. This approach starts by precisely defining an ontological feature of interest in a way that is independent of either the quantum or classical frameworks. In the case of Bell’s theorem, this generic property has come to be known as “Bell locality” [Shi13]. The content of the theorem is to show, using as few assumptions as possible, that this property can be violated in quantum theories and, ultimately, in quantum experiments. The conclusion: quantum theory does not admit ontologies with this property. The key to the extraordinary success of Bell’s theorem is in keeping all assumptions to a bare minimum.

Since Bell, many others have followed this broad outline when investigating quantum ontology and other foundational issues. Highlights include the Kochen-Specker theorem [KS67, Hel13], the Leggett-Garg inequality [LG85], and the PBR theorem [PBR12]. The most immediately impressive results are often *no-go theorems*, proving that quantum theory precludes certain ontological features (as in Bell’s theorem). However, there has also been good progress in finding formal examples of quantum ontologies which are compatible with quantum theory and do have other interesting ontological features, for example in Refs. [LJBR12, ABCL13, KS67, HR07]. In these ways, certain major features and limitations required of a plausible quantum ontology can be carefully discerned.

In these discussions, a dichotomy is often drawn between ontic and *epistemic* features or explanations. An epistemic feature is an artefact of a particular agent’s description of a physical system, subject to their (often incomplete) knowledge of it. This is in contrast to ontic features, which are considered objective states of affairs. Most often, epistemic explanations make use of probability distributions (or, more generally, probability measures) representing the agent’s (imperfect) knowledge. Note, however, that while this is consistent with their use in quantum foundations, philosophers may use the terms “ontic” and “epistemic” somewhat differently.

Quantum theory is unrivalled in its experimental success, having never faced an incompatible experimental result. This is certainly impressive, but much more understandable when quantum theory is understood as a framework rather than a theory—comparisons to Maxwellian electromagnetism, for instance, are unfair. This has led to theorists often taking quantum predictions (in the broad sense of the quantum framework) to be generically correct and any incompatibility with them to be damning. However, in science experimental results should still reign supreme. Any theoretical result should be seen as

preliminary until it can be subject to experimental verification. That said, it is common for a strict theoretical incompatibility between quantum theory and some ontology to be found before that result can be extended to something amenable to experiment.

This thesis further characterises necessary or plausible quantum ontologies in three ways. In the first, the subject is ontology of quantum states. This is the most common setting for results in quantum ontology, as the states of quantum systems are what most obviously differs compared to classical systems. In the second, the focus is on quantum causality. This is a much more recent subject in the study of quantum theory, but certainly of no less importance than the ontology of quantum states for understanding the reality of quantum systems. Third, a rather different approach is taken by considering quantum theory supplemented with time travel to the past, as this gives interesting alternative perspectives on many of the problems of quantum ontology.

1.1 “Ontology”, “Epistemic”, and “Causality”

Before proceeding, some comments are necessary on the use of terms like “ontology”, “epistemic”, and “causality” in this thesis. These terms have all been inherited from philosophy, in particular metaphysics. However, this thesis is in the tradition of quantum foundations which, as a field, tends to use these terms somewhat differently (and rather more loosely) than its more philosophical cousins.

For the avoidance of doubt, therefore, it will be useful to quickly outline what the intended meanings of such words are in this thesis. These are not philosophically robust definitions, nor are they meant to be, as the intended audience is primarily physicists.

By “ontology”, what is meant is all or part of some actual or conceivable final and fundamental description of reality. That is, supposing that there is an objective reality, this is described by the ontology. Not everybody will agree that such a thing needs to exist, of course. As noted above, “epistemic” describes a feature that is part of an agent’s subjective knowledge about a system and may often take the form of probability distributions (or measures) over ontic features. By “causality”, what is meant is the study of causal relations where one event or physical value has some actual influence on the occurrence of another, as distinct from mere correlations.

There are vast philosophical literatures on each of these topics, none of which can be effectively engaged with here. It is, however, hoped that the contents of this thesis may be able to inform some work in these philosophical fields. So the intention here is not to ignore the valuable philosophical work of these areas at all, but simply to note that they must unfortunately lie well beyond the scope of this thesis.

An inevitability of working in these areas of quantum foundations is some philosophically-charged language. Despite this, all of the results presented here should be fairly philosophically-neutral. That is, they should be valuable to people of most common philosophical leanings. On the simplest level, this is because the results are mathematical in nature. For example, the results in Chaps. 2, 3, based on the ontological models framework [Sec. 2.1.1], are very general regardless of philosophy, since they apply to any potential underlying physical theory that can be cast as an ontological model—a very large class of theories indeed.

1.2 Thesis Overview

This thesis is about ontology in quantum theory. More specifically, about obtaining a better understanding of the types of ontology that are plausible given the apparent correctness of quantum theory. This is done by carefully obtaining precise results about the nature of quantum ontology under well-defined assumptions. The study of the ontology of quantum states is relatively mature and so this thesis can build on that work to prove some very specific statements. In particular: that almost all superposition states must be real (in a well-defined sense); that quantum theory is incompatible with most types of macro-realism; and that under reasonable assumptions any conceivable overlap between many quantum states on an ontological level must be small if not zero. States are not, however, the only feature of ontology. A newer field of study is that of causes in quantum theory. This thesis constructs an expressive framework for this study by building on the specific, but important, example of a complete common cause. Since causes are typically considered to be ontological features, this framework provides a new setting for discussing ontology in quantum causality. While it is much less typical to consider time travel to the past to be real, it remains a possibility and, even beyond that possibility, time travel to the past provides an interesting playground in which to test and compare ideas in quantum ontology. This thesis therefore does so by comparing different models of quantum time travel and constructing a novel one that addresses their shortcomings. Combining these approaches to quantum ontology, a fuller and more nuanced understanding can be achieved.

In Chap. 2, the ontology of quantum states is studied. After discussing how different ontologies are classified, the current state of theorems both ruling out and demonstrating the possibilities of various classes of ontology is discussed. This discussion identifies two main shortcomings common to many such theorems, to be addressed by new results presented later in the chapter. First, it is proved that almost all superpositions must be real (in a well-defined sense) for any ontology of a quantum system of dimension $d > 3$. The same

techniques used are then applied to proving that no quantum states can be ψ -epistemic [Sec. 2.4.1] and that in large-dimensional quantum systems, many quantum states must be close to ontologically distinct. The effect of these results is to show that potential epistemic uncertainty over the exact ontological state of a quantum system is unlikely to have much power for explaining any features of quantum theory. However, all of these new results assume that quantum predictions are exactly correct and are therefore not immediately relevant to experiments, so a proof-of-concept result for error-tolerant extensions is also given. Finally, it is shown how these foundational techniques can be applied to problems in quantum information theory. In particular, it is shown that exponentially many classical resources are required to simulate a quantum channel. Similar results have appeared in the literature before, but the method used here has some key advantages. For instance, it is a significantly simpler proof than previous results and it can also be easily extended to generate potentially better bounds from new classical error-correction codes in the future.

In Chap. 3, the focus remains within the ontology of quantum states, but shifts slightly to study *macro-realism*. Macro-realism is a particular ontological property that is most commonly associated with the Leggett-Garg inequalities and corresponding no-go theorem. After introducing macro-realism in detail, the Leggett-Garg argument is briefly reviewed and some loopholes in it are identified. That is, the Leggett-Garg argument can only show that quantum theory is incompatible with one of three types of macro-realism. This is followed by a new result showing that quantum systems of $d > 3$ dimensions are incompatible with two of those three types of macro-realism, improving on the Leggett-Garg argument by using the methods of Chap. 2. Since both this result and the result showing that superpositions are real are intolerant to experimental error, these results are then brought together and given error-tolerant variations. The main conclusions of Chaps. 2, 3 are then discussed in terms of their impact and potential for further work.

Large parts of chapters 2, 3 are based on Refs. [All16, AMG16]. In particular, results of Secs. 3.1–3.3 are the result of collaboration with Owen Maroney and Stefano Gogioso. The communication work in Sec. 2.5 is joint work with Jonathan Barrett.

In Chap. 4 quantum causality is studied. After introducing the topic and its primary motivations the first major problem for quantum causality is identified. That is, the problem of finding a satisfactory version of Reichenbach’s principle for a quantum universe [Sec. 4.1.2]. An appropriate quantum Reichenbach’s principle is then carefully justified by developing a notion of quantum conditional independence. This quantum conditional independence has four definitions which are proved to be equivalent. Each definition naturally generalises a corresponding definition for classical conditional independence, lending strength to the proposed quantum Reichenbach’s principle that results. These

new definitions are then illustrated with some examples and generalised from conditional independence of two systems to conditional independence of $k \geq 2$ systems.

Chapter 5 extends this approach to quantum causality to the general case of causal models. Classically, causal models form a general framework based on Reichenbach's principle that allows systematic study of causal relationships. These classical causal models are briefly introduced before being used to motivate an analogous generalisation to quantum causal models. A proposed definition of these quantum causal models is then given that generalises quantum Reichenbach's principle from Chap. 4 in a sensible way. Some examples then illustrate this new framework, including that of Bell's theorem. The constructions of Chaps. 4, 5 are then discussed and compared to alternative approaches to formalising the study of quantum causality. In particular, it is noted that the solid foundation provided by the quantum Reichenbach's principle places confidence in the robustness of the approach taken here.

The results of Chaps. 4, 5 first appeared in Ref. [ABH⁺17] and are the result of joint work with Jonathan Barrett, Dominic Horsman, Ciarán Lee, and Robert Spekkens.

A somewhat different approach is taken in Chap. 6. There, quantum ontology is studied through the lens of possible time travel to the past. The most common way to motivate this is through noting that *closed timelike curves* (CTCs) permitted by general relativity would allow for such time travel. After motivating the approach, previous attempts to model quantum time travel using quantum circuits are briefly covered. There are two such models, called D-CTCs and P-CTCs [Sec. 6.3]. Between them, they highlight the interactions between ontology and possible non-linear extensions of quantum theory. In particular, that non-linearity added to quantum theory forces one to be more specific about ontology to consistently describe a system. Both the successes and shortcomings of D- and P-CTCs are then used to construct two classes of new models for time travel in quantum theory. From these classes, one model—dubbed T-CTCs—is fully fleshed out and compared to D- and P-CTCs at length. These findings are then discussed, with particular attention paid to the roles of non-linearity and ontological understanding in the models.

Finally, Chap. 7 summarises the results of the thesis and lays them out in the context of the further work that they suggest. In particular, the possibilities for combining the approaches to quantum ontology taken in this thesis will be discussed.

Some notational conventions used throughout this thesis have been summarised starting on page vii. As is common in modern literature on quantum foundations, the focus will be on finite-dimensional quantum systems so finite-dimensional Hilbert spaces may often be assumed.

2

Ontology of Quantum States and Superpositions

2.1 Theorising About State Ontology

When considering the ontology of quantum systems the usual questions centre on the ontology of quantum *states*. Questions such as: What properties must the ontological states of the system have, or not have? How can the structure of the space of these ontological states relate to preparations of quantum states? and How similar must this ontological state space be to the quantum state spaces, $\mathcal{P}(\mathcal{H})$ and $\mathcal{D}(\mathcal{H})$?

Many physicists would, for example, be much more comfortable if quantum phenomena could be explained with a concise and elegant realist ontology for states. Some certainly believe that existing explicit realist interpretations—such as Bohmian mechanics [BH95, Boh52a, Boh52b, Bro09] or Everettian interpretations [DG73, SBKW10]—achieve this, but such opinions are hardly uncontroversial [SBKW10, Gol16]. But even regardless of one’s view on interpretations, it is interesting to consider whether certain ontological features can be ruled out *a priori* as being incompatible with the predictions of quantum theory. Moreover, there has also been considerable cross-pollination between foundational ontology results and information-theoretic results [Mon12, BHK05, PJO15, LPZ⁺16, Mon13, Mon15, Mon16b, Lei14a], as discussed in Sec. 2.5.

Some of these questions will be addressed and answered in this chapter. In doing so, the results will fall into a tradition of “ontology theorems” in quantum foundations [Lei14a].

In particular, this chapter will concentrate on the ontology of quantum superposition states, proving that almost all quantum superpositions must be “real” in a well-defined sense [Sec. 2.3]. The techniques used to prove this main result will then be adapted for two purposes. First, to address some common shortcomings of many of the current ontology theorems [Sec. 2.4]. Second, to provide a simple way to exponentially bound the classical resources required to simulate a quantum channel [Sec. 2.5].

This chapter will also lay much of the groundwork for Chap. 3, where the general question of quantum state ontology is applied specifically to the case of macro-realism.

The methods of this chapter and the next are all based on the mathematical framework of ontological models. Because of this, before introducing the relevant background it will be prudent to first introduce this framework in Sec. 2.1.1. This framework itself is independent of quantum theory, however since the focus here is exclusively on quantum systems it will normally be assumed that ontological models exactly reproduce quantum predictions. As well as this, it will be assumed that all quantum systems are finite-dimensional, with the infinite-dimensional case being discussed briefly in Sec. 3.5.

Most of the material in this chapter and the next overlaps with that published in Refs. [All16, AMG16], with the notable exceptions of Secs. 2.5, 3.4.

2.1.1 The Ontological Models Framework

Debates about the ontology of quantum states are at least as old as quantum theory itself. In more recent years, this subject has benefited from a standardisation of definitions and concepts into the *ontological models* framework [HR07, HS10, Lei14a]. The framework of ontological models has been expressly developed to make discussions about ontology in physics precise and is the natural arena for such discussions. It will therefore be useful to lay down the mathematics of this framework before discussing the motivations and history behind the work in this chapter as it will enable a much easier and more precise discussion.

An ontological model is exactly that: a bare-bones model for the underlying ontology of some physical system. Since the ontological models framework is so sparse, very many conceivable ontologies can be understood as ontological models. The system may also be correctly described by some other, higher, theory (or framework of theories)—such as Newtonian mechanics or quantum theory—in which case the possible ontological models considered should be constrained to reproduce the predictions of that theory. By combining these constraints with the very general framework of ontological models, interesting and general conclusions can be drawn about the nature of the ontology. It is important to

note that, while ontological models are normally used to discuss quantum ontology, the framework itself is entirely independent from quantum theory.

In this section the framework of ontological models will be defined and introduced. First on its own and then as applied to quantum theory in both the absence and presence of possible experimental error. This account of ontological models is based on the one given in Ref. [Lei14a].

The framework of ontological models relies on just two core assumptions: (1) that the system of interest has some ontic state λ representing the entirety of the actual state of affairs of the system and (2) that standard probability theory may be applied to these states. That is, at any given time the entire ontology of a system is given by its ontic state. Together, these bring us to consider the ontology of some physical system as represented by some measurable space (Λ, Σ) of ontic states $\lambda \in \Lambda$ which the system might occupy (Σ being a sigma algebra of measurable subsets of Λ). The requirement that ontic states occupy a measurable space simply guarantees that sensible probabilities can be defined over them.

In the lab, a system can be prepared, transformed, and measured in certain ways. Each of these operational processes needs to be describable in the ontological model for it to be capable of describing the system.

Preparation must result in the system ending up in some ontic state λ , though the exact state need not be known. Thus, each use of an operational preparation P gives rise to some *preparation measure* μ over Λ which is a probability measure ($\mu(\emptyset) = 0$, $\mu(\Lambda) = 1$). For every measurable subset $\Omega \in \Sigma$, $\mu(\Omega)$ gives the probability that the resulting λ is in Ω . The set of all such preparation measures for some P is Δ_P . Note how the measure can vary between uses of the same preparation.

Similarly, an operational transformation T of the system will generally change the ontic state from $\lambda' \in \Lambda$ to a new $\lambda \in \Lambda$. Recalling that the ontic state λ' represents the entirety of the actual state of affairs before the transformation, then the final state can only depend on λ' (and not the preparation method or any previous ontic states, except as mediated through λ'). The transformations must therefore be described as *stochastic maps* γ on Λ . A stochastic map consists of a probability measure $\gamma(\cdot | \lambda')$ for each initial ontic state, such that for any measurable $\Omega \in \Sigma$, $\gamma(\Omega | \lambda')$ is the probability that the final λ lies in Ω given that the initial state was λ' ¹. The set of all stochastic maps corresponding to some T is Γ_T .

¹These stochastic maps, viewed as a set of functions $\gamma(\Omega | \cdot) : \Lambda \rightarrow [0, 1]$ (one for each $\Omega \in \Sigma$), must be *measurable functions*. That is, for any measurable set $\mathcal{S} \subseteq [0, 1]$ and any $\gamma(\Omega | \lambda')$, then $\{\lambda' \in \Lambda : \gamma(\Omega | \lambda') \in \mathcal{S}\} \subseteq \Lambda$ is a measurable set.

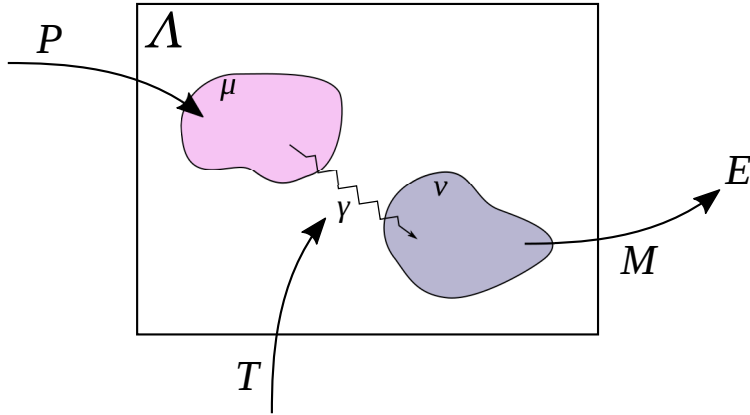


Figure 2.1: Illustration of the basic prepare-transform-measure procedure in an ontological model. The box represents the entire ontic state space Λ , while shaded regions represent the ontic states that can be prepared by the indicated preparation measures. Some preparation P is performed, resulting in an ontic state according to $\mu \in \Delta_P$. Some transformation T is performed, resulting in some stochastic map $\gamma \in \Gamma_T$, effectively transforming μ to ν as in Eq. (2.2). Finally, some measurement M is made, resulting in some outcome E according to probabilities given by Eq. (2.1).

Finally, a measurement M may give rise to some outcome E . Again, which outcome is obtained can only depend on the current ontic state λ' . Therefore a measurement M gives rise to a conditional probability distribution² $\mathbb{P}_M(E | \lambda')$, sometimes called a *response function*. The set of all such response functions corresponding to M is Ξ_M . For this thesis it is only necessary to consider measurements that have countable sets of possible outcomes E .

Putting these parts together: given a system where a preparation P is performed followed by some transformation T and some measurement M then the ontological model for that system must have some preparation measure $\mu \in \Delta_P$, stochastic map $\gamma \in \Gamma_T$, and conditional probability distribution $\mathbb{P}_M \in \Xi_M$ such that the probability of obtaining outcome E is

$$\mathbb{P}_M(E | \nu) = \int_{\Lambda} d\nu(\lambda) \mathbb{P}_M(E | \lambda) \quad (2.1)$$

where

$$\nu(\Omega) \stackrel{\text{def}}{=} \int_{\Lambda} d\mu(\lambda) \gamma(\Omega | \lambda) \quad (2.2)$$

is the effective preparation measure obtained by preparation P followed by transformation T . This is schematically illustrated in Fig. 2.1.

Note that ontological models must be closed under transformations. That is, for any preparation μ and transformation γ in the model then the preparation ν defined by Eq. (2.2)

²These probability distributions, viewed as functions $\Lambda \rightarrow [0, 1]$, must also be measurable functions.

must also exist in the model (since a preparation followed by a transformation is itself a type of preparation).

It is useful to think of stochastic maps acting on preparation measures as an emergent property of them acting on the underlying ontic state space. If an ontic state is sampled from μ and then transformed via γ the effect is the same as sampling from the measure ν defined by Eq. (2.2). Rather than having to say this in such a cumbersome way, the notation $\mu \xrightarrow{\gamma} \nu$ will be used to express the same thing and may be read “ μ transforms to ν under γ ”. The full definition of this notation is, however, via Eq. (2.2) as above.

Summarising, an ontological model for some physical system:

1. defines a measurable space (Λ, Σ) of ontic states for the system;
2. for each possible transformation T defines a set of stochastic maps $\gamma \in \Gamma_T$ from Λ to itself;
3. for each possible preparation P defines a set of preparation measures $\mu \in \Delta_P$ over Λ , ensuring closure under the actions of the stochastic maps as in Eq. (2.2);
4. for each possible measurement M defines a set of response functions $\mathbb{P}_M \in \Xi_M$ over the outcomes given $\lambda \in \Lambda$;

and then produces probabilities for measurement outcomes via Eqs. (2.1, 2.2).

It should be noted that ontological models are usually presented using probability *distributions* rather than the more mathematically involved measures used here. However, as noted in Ref. [Lei14a], this simplification precludes many reasonable ontological models, including the archetypal Beltrametti-Bugajski model [BB95] (see also Sec. 2.1.2). The more accurate approach is taken here both for the accuracy itself and to serve as a resource of how to construct proofs in measure-theoretic ontological models as such proofs are rarely seen in the literature.

The above definition for ontological models *per se* does not have much useful structure. The power is found when the probabilities given by Eqs. (2.1, 2.2) are constrained to match those given by other theories known to accurately describe the system (such as quantum theory) or by experiments. To this end, consider applying ontological models to systems accurately described by quantum theory.

An ontological model is defined by some ontic state space (Λ, Σ) as well as the relevant preparation measures, stochastic maps, and conditional probability distributions. For a quantum system these must include at least the following. For each state $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ there must be a set $\Delta_{|\psi\rangle}$ of preparation measures $\mu_{|\psi\rangle}$ —potentially at least one for each distinct

experimental procedure for preparing $|\psi\rangle$. Similarly, for each unitary operator U on \mathcal{H} there is a set Γ_U of stochastic maps γ_U and for each basis measurement $M = \{|i\rangle\}_{i=0}^d$ there is a set Ξ_M of conditional probability distributions \mathbb{P}_M —again, potentially at least one stochastic map/probability distribution for each experimental procedure for transforming/measuring.

Since the quantum system is accurately described by quantum theory, the ontological model must also reproduce the predictions of quantum theory. That is, for any $|\psi\rangle \in \mathcal{P}(\mathcal{H})$, $\mu \in \Delta_{|\psi\rangle}$, $U, \gamma \in \Gamma_U$, basis M , and $\mathbb{P}_M \in \Xi_M$ it is required that

$$|\langle i|U|\psi\rangle|^2 = \int_{\Lambda} d\nu(\lambda) \mathbb{P}_M(|i\rangle|\lambda), \quad \forall |i\rangle \in M \quad (2.3)$$

where $\mu \rightsquigarrow \nu$ is defined as in Eq. (2.2). Note also that $\nu \in \Delta_{U|\psi}$ since preparing the quantum state $|\psi\rangle$ (via any ontological preparation $\mu \in \Delta_{|\psi\rangle}$) followed by performing the quantum transformation U (via any $\gamma \in \Gamma_U$) is simply a way to prepare the quantum state $U|\psi\rangle$.

So Eq. (2.3) must hold when quantum theory is known to accurately describe the system. What of the case where quantum theory only approximately describes the system? Suppose, for instance, that the probabilities predicted by quantum theory are accurate to within $\pm\epsilon$ for some given $\epsilon \in (0, 1]$. It follows that the ontological model need only reproduce these approximate predictions, so Eq. (2.3) is replaced by

$$|\langle i|U|\psi\rangle|^2 + \epsilon \geq \int_{\Lambda} d\nu(\lambda) \mathbb{P}_M(|i\rangle|\lambda) \geq |\langle i|U|\psi\rangle|^2 - \epsilon, \quad \forall |i\rangle \in M \quad (2.4)$$

where, again, $\nu \in \Delta_{U|\psi}$ is defined as in Eq. (2.2).

This mathematical framework forms the foundation of many modern results in the ontology of quantum states. It may be seen as a successor to, and extension of, the *hidden variables* models used historically [HR07, and references therein]. Recall that, unless otherwise stated, quantum predictions are assumed to be exactly accurate throughout this chapter and Chap. 3. In other words, Eq. (2.3) will be assumed on top of the bare ontological models framework unless otherwise stated.

2.1.2 The Desire for Simpler Ontologies

Given the definitions and mathematical background of the ontological models framework, it is now time to consider the motivation behind wanting a relatively simple ontology for quantum states. Whether or not you believe it to be possible, this motivation is quite easy to understand.

The simplest way to capture textbook quantum theory in an ontological model is to: have the pure quantum states *be* the ontic states $\Lambda \cong \mathcal{P}(\mathcal{H})$, have the pure state preparation

measures give unit probability to the corresponding ontic state, and have transformations and measurements act on Λ exactly as they do on $\mathcal{P}(\mathcal{H})$ in textbook quantum theory. This transliteration of quantum theory to an ontological model is called the *Beltrametti-Bugajski model* [BB95, HR07, Lei14a].

The Beltrametti-Bugajski model is often seen as an unattractive ontology for several reasons. It requires an uncountable ontic state space for any non-trivial system, even when there are only two distinguishable preparations. It also forces one to talk about ontology globally, due to entanglement. Finally, it contains a lot of redundancy: the properties of the quantum state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ are entirely inherited from properties of $|0\rangle$ and $|1\rangle$, it therefore seems ontologically extravagant to describe $|\psi\rangle$ entirely separately from $|0\rangle$ and $|1\rangle$ on an ontological level.

So how does one respond to this easy, but rather ugly, understanding of state ontology? One way is to simply accept that the ontology of quantum states is just like this or to argue that it is not as ugly as it might seem. That view might lead one towards Everettian interpretations of quantum theory, for example. Another is to deny the need for a realist ontology of quantum states altogether—a position variously called “anti-realist”, “neo-Copenhagen”, and “instrumentalist” [Lei14a]—a route with its own conceptual and philosophical hurdles. A third way (and the one of most interest in this thesis) is to seek a more elegant realist ontology that might underlie quantum theory.

This third approach is sometimes called that of the “epistemic realist” and naturally leads to using ontological models because of their extreme generality for realist ontologies. By using the fact that preparing some quantum state may leave some uncertainty over the ontic state (the preparation measure can give non-zero probability to a large number of ontic states), the epistemic realist can hope to explain many features of quantum systems as arising from this “ontological uncertainty”. These features include the indistinguishability of non-orthogonal states, no-cloning, stochasticity of measurement outcomes, and the exponential increase in state complexity with increasing system size [Spe07].

In particular, the epistemic realist can hope to use the fact that both quantum states and ontological model preparations can “overlap”. Two quantum states in $\mathcal{P}(\mathcal{H})$ overlap by an amount quantified by the Born rule (equivalently, their inner product). The corresponding ontic overlap occurs when preparations for different quantum states can prepare some of the same ontic states.

A typical motivation given for the epistemic realist view is that such ontic overlaps might naturally explain the indistinguishability of non-orthogonal quantum states [Spe07, BCLM14, Lei14b]. If two non-orthogonal quantum states have finite ontic overlap, then sometimes they will prepare the same ontic states and, since the ontic state describes




Figure 2.2: Illustration of the epistemic realist explanation for indistinguishability of non-orthogonal quantum states. Suppose that $\mu \in \Delta_{|\psi\rangle}$ and $\nu \in \Delta_{|\phi\rangle}$ are preparation measures over ontic state space Λ for non-orthogonal quantum states $|\psi\rangle$ and $|\phi\rangle$. If they overlap, as shown, there are ontic states such as the $\lambda \in \Lambda$ illustrated that may be prepared by both. In this case, there is no way to distinguish which preparation, μ or ν , was performed, naturally explaining why there is no way to perfectly distinguish between $|\psi\rangle$ and $|\phi\rangle$.

the entire ontology of the system, there is no way to tell them apart. By looking at the probability of this occurring, the ability of ontic overlaps to account for indistinguishability can be quantified [Sec. 2.2.2]. This is illustrated in Fig. 2.2. More thorough discussions of the way that epistemic realist explanations can explain puzzles in quantum foundations can be found in Refs. [Spe07, Spe14].

The epistemic realist perspective on the foundations of quantum theory is not only philosophically attractive but also appears to be tenable. There are theories that explain the quantum state in an epistemically realist manner that reproduce interesting and large subsets of quantum theory including many characteristically quantum features [Spe14, Spe07, BRS12, JL15]. There are also several explicit ontological models exactly describing whole isolated quantum systems that make very good use of ontological uncertainty, though only in $d = 2$ dimensions as discussed below.

A particularly powerful result for the epistemic realist would be to find an explanation of certain quantum superposition states as statistical effects due to ontic overlap. Superpositions are behind quantum interference, the uncertainty principle, wave-particle duality, entanglement, Bell non-locality [Bel87], and the probable increased computational power of quantum theory [JL03]. Perhaps most alarmingly, superpositions give rise to the measurement problem, so captivatingly illustrated by the ‘‘Schrödinger’s cat’’ thought experiment [Gib87]. Explaining superpositions with ontological uncertainty would therefore go a long way to explaining these features of quantum theory.

In quantum theory, if $|\psi\rangle$ is a superposition over, say, $|0\rangle$ and $|1\rangle$ then all of the properties of $|\psi\rangle$ are inherited directly from $|0\rangle$ and $|1\rangle$, mediated by the amplitudes of the

superposition. This is in a very similar way to how a probability distribution over classical states inherits all of its properties from the underlying classical states, mediated by the probabilities. This situation would naturally be explained if the superposition $|\psi\rangle$ was just a statistical effect over the ontic states corresponding to $|0\rangle$ and $|1\rangle$, exactly as ontic overlap could explain indistinguishability. A full mathematical treatment of this intuition is deferred until Sec. 2.3.1.

If the epistemic realist programme is successful in finding an elegant underlying ontology that naturally explains any of these quantum features, it will become a very attractive proposition indeed. The primary problem for the epistemic realist is that no such ontological model has yet been found. Rather, the epistemic realist sits between the success of models that partially reconstruct quantum theory [Spe14, Spe07, BRS12, JL15] and the ontology theorems that constrain their ability to reconstruct the rest, discussed in Sec. 2.1.4.

2.1.3 Classifying Ontologies

In order to effectively discuss the types of ontological models are or are not compatible with quantum theory, it is necessary to identify some classes of ontological models worth discussing. By finding examples of ontological models in some classes compatible with quantum theory and proving the impossibility of such models in other classes, a clearer prognosis for the epistemic realist perspective emerges. The main classifications from the literature used in this chapter are: ψ -ontic, ψ -epistemic, maximally ψ -epistemic, and various contextualities.

An ontological model for a quantum system is ψ -ontic if and only if each ontic state can only be prepared by a single quantum state. That is, if one were able to see the ontology of the system directly, there would be no ambiguity as to which quantum state was prepared. The Beltrametti-Bugajski model is trivially ψ -ontic, as the ontic states *are* the quantum states, but it is possible to consider others (including Bohmian mechanics) which are ψ -ontic but quantum state preparations can result in more than one ontic state.

The opposite of ψ -ontic is ψ -epistemic. That is, an ontological model is ψ -epistemic if and only if it is not ψ -ontic. A ψ -epistemic model is something of a minimum requirement for the epistemic realist since ψ -ontic models leave no room for explaining anything other than indeterminism by ontological uncertainty. If there were a theorem proving that all ontological models that reproduce quantum statistics must be ψ -ontic then the epistemic realist programme would certainly be dead. However, no such theorem can exist as there are ontological models for every finite dimension which are both ψ -epistemic and compatible with quantum theory [LJBR12, ABCL13].

So ψ -ontic delineates one extreme of the spectrum of conceivable ontological models. The other end is marked by *maximally ψ -epistemic* models. Recall that both quantum states and ontological model preparations can overlap. An ontological model is maximally ψ -epistemic if and only if the ontic overlap entirely accounts for the Born rule overlap [Mar12, LM13, Lei14b, Bal14]. As will become clear in Sec. 2.4.1, it is impossible for ontic overlaps to be any larger than this, hence such models are “maximally” ψ -epistemic.

A maximally ψ -epistemic ontology would be ideal for the epistemic realist, as the large amount of ontological uncertainty between quantum states would allow for powerful explanations [Lei14a, Mar12, LM13, Lei14b, Bal14]. Such ontological models do exist in $d = 2$ dimensions, such as the Kochen-Specker model [KS67, HR07], but no such model can exist for $d > 2$, as discussed in Sec. 2.1.4.

These three notions— ψ -ontic, ψ -epistemic, and maximally ψ -epistemic—form a coarse first-order classification of the ontological models of interest to the epistemic realist. Precise definitions of each will be deferred until Sec. 2.4.1 once the appropriate mathematical background has been covered. While this classification is far from nuanced, it serves as an excellent starting point to discuss finer distinctions within the class of ψ -epistemic models.

For example, the concepts behind the ψ -epistemic/ontic dichotomy can also be used to discuss the reality of quantum superpositions. Consider the example of Schrödinger’s cat. Schrödinger’s cat is set up to be in a superposition of $|\text{dead}\rangle$ and $|\text{alive}\rangle$ quantum states. The epistemic realist (and probably the cat) would ideally prefer the ontic state of the cat to only ever be one of “dead” or “alive” (*viz.*, only in ontic states accessible when preparing either the $|\text{dead}\rangle$ or $|\text{alive}\rangle$ quantum states). In that case, the cat’s apparent quantum superposition would be *epistemic*—there would be nothing “real” about the superposition state not already captured by $|\text{dead}\rangle$ and $|\text{alive}\rangle$. Conversely, if there are ontic states which can only obtain when the cat is in a quantum superposition (and never when the cat is in either quantum $|\text{dead}\rangle$ or $|\text{alive}\rangle$ states), then the superposition is unambiguously *ontic*: there are ontological features which correspond to that superposition but not to non-superpositions, so that superposition is real. Precise mathematical definitions will be deferred until Sec. 2.3.1.

It is not yet known to what degree epistemic superpositions are compatible with quantum theory. Some of the epistemic realist theories reproducing subsets of quantum theory noted above do include epistemic superpositions. The question of the reality of superpositions in quantum theory is, therefore, very much open.

Obviously quantum superpositions are different from proper mixtures of basis states. The question here is rather whether quantum superpositions can be understood as distri-

butions over some subset of underlying ontic states, where each such ontic state is also accessible by preparing some basis state.

A neat toy example is found in Spekkens’ toy theory [Spe07], where the “toy-bit” reproduces a subset of qubit behaviour. A toy bit consists of four ontic states, a, b, c, d , and four possible preparations, $|0\rangle, |1\rangle, |+\rangle, |-\rangle$, which are analogous to the correspondingly named qubit states. Each preparation corresponds to a uniform probabilistic distribution over exactly two ontic states: $|0\rangle$ is a distribution over a and b ; $|1\rangle$ a distribution over c and d ; $|+\rangle$ over a, c ; and $|-\rangle$ over b, d . Full details of how these states behave and how they reproduce qubit phenomena are described in Ref. [Spe07]. For the purposes here, it suffices to note that all ontic states corresponding to the superposition states $|+\rangle$ and $|-\rangle$ are also ontic states corresponding to either $|0\rangle$ or $|1\rangle$ —this toy-bit has nothing on the ontological level which can be identified as a superposition. The toy theory superpositions are epistemic. Toy models such as this therefore lend credibility to the idea that quantum superpositions themselves might, in a similar way, fail to have an ontological basis.

Finally, it is common to identify ontological models with various features collectively referred to as “contextuality” [Spe05, HR07, LM13]. Contextuality was first introduced with what is now known as “Kochen-Specker contextuality” [KS67, Hel13]. However, Ref. [Spe05] showed how contextuality can be more broadly be thought of in terms of operational equivalence. Loosely, contextuality refers to when operationally indistinguishable situations are described differently in the ontological model. There are various types of contextuality of interest in various situations (including Kochen-Specker contextuality) but the only one necessary for this thesis is *preparation contextuality*. Roughly, a model is preparation contextual if it contains separate descriptions for preparations that are operationally equivalent. That is, the ontology depends on the “context” of the preparation in a way that is not operationally discernible. A precise mathematical definition will be given in Sec. 2.4.1.

Contextuality of any kind is traditionally thought of as undesirable for the epistemic realist but it is certainly not fatal. It has even been argued that contextuality should be expected in an appropriate epistemic realist theory [Bal14].

While the classifications introduced here are relatively coarse, they are good starting points for identifying more subtle classifications, some of which will be discussed in Sec. 2.4. Before getting there though, it is appropriate first to review the main results in the literature concerning the classifications already introduced.

2.1.4 Previous Ontology Theorems

As noted above, there is a tradition of “ontology theorems” in quantum foundations. These are typically no-go theorems, proving that certain classes of ontological models can never be compatible with the predictions of quantum theory. The most famous ontology theorem is Bell’s theorem [Bel87], which primarily concentrates on locality. However, the ontological models framework as defined here—and in much of the recent literature—avoids talking about composition of local systems, preferring to concentrate on single systems where issues of locality and composition do not arise. The primary reason for this is the PBR theorem.

The PBR theorem [PBR12] proves that all ontological models for $d > 2$ must be ψ -ontic if they satisfy the *preparation independence postulate* (PIP). The PIP is a reasonable extra restriction on the structure of ontological models for multipartite systems that is not present in the bare ontological models framework. The PBR theorem is therefore a very powerful blow to the epistemic realist. However, it does problematically depend on the PIP, which has been challenged [Man16, ESSV13, Hal11, SF14, Wal13] (especially as being similar to Bell locality, which is already ruled out by Bell’s theorem) and without it the PBR theorem is impotent. It is for this reason that recent work has avoided the PIP and related issues by looking only at single systems in order to seek more conclusive ontology results. This will be the perspective taken here.

While ψ -epistemic models are not possible in multipartite systems with the PIP, they are possible for single systems without the PIP. As noted above, explicit ψ -epistemic models have been presented for every finite dimension in Refs. [LJBR12, ABCL13].

The obvious next question is whether maximally ψ -epistemic models are possible without the PIP? In $d = 2$ dimensions they are, as shown explicitly by the Kochen-Specker model [KS67, HR07]. However, they are not possible for $d > 2$. There are several disparate ways of proving this [HR07, Bal14, LM13, Mar12, BCLM14, Bra14, Lei14b]. One of these proofs even existed before the idea of “maximally ψ -epistemic” [HR07, Bal14]. Reference [LM13] proved that “maximally ψ -epistemic” implies Kochen-Specker contextuality and is therefore ruled out by the Kochen-Specker theorem [KS67, Hel13]. The majority of these theorems, however, prove their results by bounding ontic overlaps to be less than maximal for $d > 2$ [Mar12, BCLM14, Bra14, Lei14b, Bal14]. As a result these theorems are able to rule out more than simply maximally ψ -epistemic models; each also rules out some subset of non-maximally ψ -epistemic models, depending on the exact bound found. In particular, the stated aim is often to show that certain ontic overlaps must be small and therefore, whilst ψ -epistemic ontologies are possible, using them to explain the indistinguishability

of quantum states is implausible. The overlap bounds derived each also tend to zero in some limit, suggesting negligible ontic overlap in those limits (often in large dimensions).

These overlap bound ontology theorems are of most interest to this chapter. As numerical inequalities their conclusions can be more nuanced in ruling out a range of ontological models. They are also likely to be more amenable to experimental investigation [NMS⁺15, RDB⁺15, Kne16]. This contrasts especially with results based on the Kochen-Specker theorem, which has finite-precision loopholes [Mey99, Ken99, CK00, BK04].

The above is not an exhaustive list, nor a thorough discussion, of the current state of ontology theorems. Only those directly relevant to the approach taken in this chapter were included. The majority of other overlap theorems, including Refs. [CR13, CR12, CR11, CR16, Har04, Mon08, Har13, PPM13, ABCL13], are more specialised, often making extra assumptions on top of the ontological models formalism. Several are reviewed at length in Ref. [Lei14a]. However, no previous result has yet tackled the ontology of quantum superposition states directly.

Before discussing some shortcomings of these results, it is worth noting that one will often hear that ontological models for quantum systems must be preparation contextual [Spe05, LM13, Lei14a]. However, such results only rule out preparation non-contextuality for *mixed* quantum states—*viz.* they prove that there must exist mixed quantum states where the ontological description of the preparation can vary. Preparation non-contextuality for *pure* quantum states is certainly possible [LJBR12, ABCL13, Lei14a] and often assumed without question. This is likely because the contextuality/non-contextuality distinction for pure state preparations is often irrelevant [Lei14a]. However, as will be seen in Sec. 2.4, this will not be the case here and discussion of pure state preparation contextuality will be required.

2.1.5 Limitations and Loopholes

Of the ontology theorems mentioned in Sec. 2.1.4, those which bound ontic overlaps for $d > 2$ are probably the greatest threat to the epistemic realist view. References [Mar12, BCLM14, Bra14, Lei14b, Bal14] each prove a bound of this type: a set of quantum states is constructed and an inequality on ontological overlaps is proved to hold for at least one pair from the set, bounding their ontic overlaps to be less than the corresponding Born rule overlap. Trivially, this implies that the ontological model for the system cannot be maximally ψ -epistemic. The authors further argue that when these upper bounds become small (typically in large dimensions) then explaining indistinguishability of quantum states by ontological uncertainty becomes implausible. Indeed, if the epistemic realist is hoping to use the ontic overlap to explain quantum features, then it seems unlikely that a very

small overlap could have much explanatory power. All of these theorems, however, share at least the following two shortcomings.

First, the proofs are non-constructive proofs of existence. That is, they conclude that there is *some* pair of quantum states with a small ontic overlap, but no guarantee is made as to *which* pair. This may not seem problematic, but consider the following. If a theorem only requires that I incorporate at least one pair of quantum states with small ontic overlap, then I am not prevented from postulating a model where exactly one pair has small ontic overlap, but all other pairs overlap maximally. Such a model would be indistinguishable from a maximally ψ -epistemic model, since the probability of encountering exactly that pair of quantum states is zero. Such theorems, as stated, are therefore very weak restrictions on the types of ontological model that are compatible with quantum theory.

The above is a caricature of the loophole and by examining the proofs of the theorems in Refs. [Mar12, BCLM14, Bra14, Lei14b, Bal14] it seems possible to use the unitary symmetry of quantum theory to extend them and obtain conclusions that require more pairs of quantum state to have less-than-maximal ontic overlap. However, this has not been done rigorously so the loophole remains, making them much weaker than they seem at face value. Even if such improvements were made, they would still not guarantee that any particular pair of quantum states has less-than-maximal overlap, which may leave significant room for the epistemic realist to explain quantum experiments using ontological uncertainty simply by assuming that the quantum states used in the experiment do have large ontic overlap.

The second shortcoming is less of a loophole and more of a limitation. The overlap bounds proved in Refs. [Mar12, BCLM14, Bra14, Lei14b, Bal14] each approach zero in some limit (usually as $d \rightarrow \infty$). However, as noted in Ref. [Lei14a] as one approaches these limits, the sets of states considered also approach orthogonality. Orthogonal quantum states are perfectly distinguishable and therefore trivially must have zero ontic overlap. This casts doubt on these results having any meaning at all in those limits where they at first appear to be most powerful. Even being close to the limit makes the results seem less impressive, it is easier to believe that quantum states that are close to orthogonality have very small ontic overlap than those that are nearly collinear.

These theorems are therefore less powerful than they initially may seem and far from conclusive. In particular, it would be much more convincing to have a theorem that could identify particular pairs of quantum states which have bounded ontic overlap and/or a theorem that identified ontic overlaps approaching zero for quantum states with fixed inner product.

2.1.6 Chapter Overview

In this chapter the epistemic realist view of quantum states will be interrogated, with a particular focus on the ontology of superpositions.

In Sec. 2.2 the asymmetric overlap and anti-distinguishable quantum states will be introduced. This is necessary mathematical groundwork for the results that follow. This will include proving several lemmas at a level of rigour that has not yet been achieved in the literature.

Section 2.3 will discuss the ontology of quantum superposition states. A theorem will be presented showing that for a $d > 3$ dimensional quantum system, almost all quantum superpositions with respect to any given orthonormal basis must be ontic. That is, the epistemic realist must include superpositions in their ontology explicitly.

The techniques used will then be adapted in Sec. 2.4 to address the shortcomings of current ontology theorems noted in Sec. 2.1.5. In particular, theorems will be proved bounding ontic overlap for large numbers of specific pairs of quantum states and demonstrating bounds that approach zero overlap in the large- d limit, even while the quantum states have fixed inner product. In order to adapt the method of the previous theorem, a very mild form of preparation non-contextuality will be assumed, which will be defended as a natural assumption in Sec. 2.4.4. The feasibility of extending these results to be error-tolerant will be tackled in Sec. 2.4.3.

It is fairly common for ontology theorems to inspire results in quantum information and communication. In Sec. 2.5 this example will be followed and a simple method for exponentially bounding the ability of classical systems to simulate quantum communication will be demonstrated.

Chapter 3 will continue directly from this work, applying the same techniques to the problem of “macro-realism” in quantum theory. As a result, Sec. 3.4 will present an error-tolerant variant theorem that can rule out epistemic superpositions. This result will have to be logically weaker than the theorem presented in this chapter to gain error tolerance. The results of both chapters will then be discussed together in Sec. 3.5.

2.2 The Asymmetric Overlap

As discussed in Sec. 2.1.4, many ontology theorems proceed by bounding the ontic overlap between quantum states. When a quantum state is prepared, some ontic states can obtain and some cannot. Loosely speaking, the ontic overlap between two quantum states is made up of those ontic states that can obtain by preparing either. To precisely discuss and derive ontology theorems, it is necessary to quantify this notion in some way.

In quantum theory, the overlap between any pair of $|\psi\rangle, |\phi\rangle \in \mathcal{P}(\mathcal{H})$ is quantified by the Born rule probability $|\langle\phi|\psi\rangle|^2$. That is, for a system prepared in state $|\psi\rangle$ the probability for it to behave (for all intents and purposes) like it was prepared in state $|\phi\rangle$ is $|\langle\phi|\psi\rangle|^2$.

Adapting this logic to an ontological model for the quantum system, consider the probability that a system prepared according to measure μ behaves like it was prepared according to ν . That is, the probability that the ontic state obtained from μ could also have been obtained from ν . This quantity is called the *asymmetric overlap* and is mathematically defined [Bal14, LM13, Mar12]

$$\varpi(\nu | \mu) \stackrel{\text{def}}{=} \inf\{\mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) = 1\}, \quad (2.5)$$

recalling that the infimum of a subset of real numbers is the greatest lower bound of that set. This is because a preparation of ν has unit probability of producing a λ from each measurable subset $\Omega \subseteq \Lambda$ that satisfies $\nu(\Omega) = 1$. Therefore minimising $\mu(\Omega)$ with respect to Ω gives the desired probability. The notation used here deliberately borrows from conditional probabilities and $\varpi(\nu | \mu)$ may be read “asymmetric overlap with ν given a preparation of μ ”.

It is convenient to slightly overload the terminology and notation and define the asymmetric overlap for more general quantities. First, define the asymmetric overlap between a preparation measure μ and some *quantum state* $|\phi\rangle$ as the probability that preparing μ will produce a λ obtainable by preparing $|\phi\rangle$. This corresponds to

$$\varpi(|\phi\rangle | \mu) \stackrel{\text{def}}{=} \inf\{\mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) = 1, \forall \nu \in \Delta_{|\phi}\}\}. \quad (2.6)$$

The next useful generalisation is the asymmetric overlap of some preparation measure μ with *two* quantum states $|0\rangle, |\phi\rangle$. This can be thought of as the *union* of the overlaps expressed by $\varpi(|\phi\rangle | \mu)$ and $\varpi(|0\rangle | \mu)$ and is mathematically defined as

$$\varpi(|0\rangle, |\phi\rangle | \mu) \stackrel{\text{def}}{=} \inf\{\mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) = \chi(\Omega) = 1, \forall \nu \in \Delta_{|\phi}, \chi \in \Delta_{|0}\}\}. \quad (2.7)$$

So $\varpi(|0\rangle, |\phi\rangle | \mu)$ is the probability that sampling from μ produces a λ obtainable by preparing either $|0\rangle$ or $|\phi\rangle$.

Finally, this can be extended to the asymmetric overlap of μ with a *set* of quantum states $\mathcal{S} \subseteq \mathcal{P}(\mathcal{H})$ given preparation of μ in the obvious way

$$\varpi(\mathcal{S} | \mu) \stackrel{\text{def}}{=} \inf\{\mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) = 1, \forall \nu \in \Delta_{|\phi}, \forall |\phi\rangle \in \mathcal{S}\}\}. \quad (2.8)$$

Clearly, the asymmetric overlap is not the only sensible way to quantify ontic overlaps. Its main advantages are a natural motivation and ontological interpretation and it will be

used extensively in what follows. Another popular choice is the symmetric overlap, which will be used briefly in Sec. 2.4.3.

The remainder of this section will flesh out the asymmetric overlap's properties. This will unfortunately be somewhat dry and mathematical but, once in place, will significantly reduce the complexity of the proofs later in the chapter.

2.2.1 Properties of the Asymmetric Overlap

Many of the properties of the asymmetric overlap derived here will not greatly surprise those familiar with using ontological models. However, they have not yet been achieved at this level of rigour using the measure-theoretic approach to ontological models, necessitating full proofs.

While this rigour is important, it can often obscure the more intuitive reasons that the results hold. To offset this, most proofs presented below will start with a rough-but-intuitive argument followed by the more cumbersome-but-correct measure theory.

It is convenient to begin with the following definition to simplify the notation for the rest of this section.

Definition 2.1. For any measurable function $g : \Lambda \rightarrow [0, 1]$ let

$$\bar{k}(g) \stackrel{\text{def}}{=} \ker(1 - g) = \{\lambda \in \Lambda : g(\lambda) = 1\} \in \Sigma. \quad (2.9)$$

Next, this technical lemma forms the basis of proofs of many of the following properties.

Lemma 2.1. Given any measurable function $f : \Lambda \rightarrow [0, 1]$ and preparation measure ν satisfying $\int_{\Lambda} d\nu(\lambda) f(\lambda) = 1$, then $\nu(\bar{k}(f)) = 1$.

Proof. Roughly, this lemma simply affirms that if the average $f(\lambda)$ according to ν is unity, then the probability that $f(\lambda) = 1$ according to ν is also unity.

If $\int_{\Lambda} d\nu(\lambda) f(\lambda) = 1$ then

$$1 = \int_{\bar{k}(f)} d\nu(\lambda) f(\lambda) + \int_{\Lambda \setminus \bar{k}(f)} d\nu(\lambda) f(\lambda) \quad (2.10)$$

$$= \nu(\ker \bar{f}) + \int_{\Lambda \setminus \bar{k}(f)} d\nu(\lambda) f(\lambda) \quad (2.11)$$

since if $\lambda \in \bar{k}(f)$ then $f(\lambda) = 1$.

Suppose that the the second term in Eq. (2.11) is non-zero. Since $f(\lambda) < 1$ for all $\lambda \in \Lambda \setminus \bar{k}(f)$, then $\int_{\Lambda \setminus \bar{k}(f)} d\nu(\lambda) f(\lambda) < \nu(\Lambda \setminus \bar{k}(f))$. This further implies

$$\nu(\bar{k}(f)) + \nu(\Lambda \setminus \bar{k}(f)) > 1 \quad (2.12)$$

$$\nu(\Lambda) > 1 \quad (2.13)$$

which is a contradiction as $\nu(\Lambda) = 1$ by definition. Therefore the second term in Eq. (2.11) must be zero and Eq. (2.11) implies $1 = \nu(\bar{k}(f))$ as desired. \square

Having established this technical background, the following shows a fundamental property of the asymmetric overlap: it is upper-bounded by the Born rule probability.

Lemma 2.2. *For any pair of pure quantum states $|\psi\rangle, |\phi\rangle \in \mathcal{P}(\mathcal{H})$ and for any preparation $\mu \in \Delta_{|\psi\rangle}$, then*

$$\varpi(|\phi\rangle | \mu) \leq |\langle\phi|\psi\rangle|^2. \quad (2.14)$$

From this it immediately follows that $\varpi(\nu | \mu) \leq |\langle\phi|\psi\rangle|^2$ for every $\nu \in \Delta_{|\phi\rangle}$.

Proof. Roughly, this holds because almost all ontic states obtained by preparing $|\phi\rangle$ must also return $|\phi\rangle$ in any measurement where that is an option. $|\langle\phi|\psi\rangle|^2$ is the probability of getting outcome $|\phi\rangle$ when $|\psi\rangle$ has been prepared. This must therefore occur at least as often as obtaining an ontic state accessible by $|\phi\rangle$ when preparing $|\psi\rangle$. This will now be properly proved.

Consider preparing $|\phi\rangle$ via any $\nu \in \Delta_{|\phi\rangle}$ and then performing some quantum measurement $M_\phi \ni |\phi\rangle$. Since the ontological model reproduces quantum probabilities [Eq. (2.3)] then

$$\int_{\Lambda} d\nu(\lambda) \mathbb{P}_{M_\phi}(|\phi\rangle | \lambda) = 1. \quad (2.15)$$

Letting $g(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_{M_\phi}(|\phi\rangle | \lambda)$, Lem. 2.1 shows that $\nu(\bar{k}(g)) = 1$ for every $\nu \in \Delta_{|\phi\rangle}$.

Now consider preparing $|\psi\rangle$ via any $\mu \in \Delta_{|\psi\rangle}$ and then measuring with the same M_ϕ . By Eq. (2.3)

$$|\langle\phi|\psi\rangle|^2 = \int_{\Lambda} d\mu(\lambda) g(\lambda) \geq \int_{\bar{k}(g)} d\mu(\lambda) g(\lambda) = \mu(\bar{k}(g)). \quad (2.16)$$

having used that $g(\lambda \in \bar{k}(g)) = 1$ in the final step. Recalling that $\nu(\bar{k}(g)) = 1$ for every $\nu \in \Delta_{|\phi\rangle}$, then $\varpi(|\phi\rangle | \mu) \leq \mu(\bar{k}(g))$ by definition. Combining these inequalities gives the desired general result. Applying these inequalities to any particular $\nu \in \Delta_{|\phi\rangle}$ gives the specific case $\varpi(\nu | \mu) \leq |\langle\phi|\psi\rangle|^2$. \square

The next property demonstrates that the asymmetric overlap is non-increasing under transformations.

Lemma 2.3. *Let unitary U satisfy $U|0\rangle = |\phi\rangle$ and $\mu' \xrightarrow{\gamma} \mu$ for some $\gamma \in \Gamma_U$, then*

$$\varpi(|\phi\rangle | \mu) \geq \varpi(|0\rangle | \mu'). \quad (2.17)$$

Proof. Roughly, this is because ontic states preparable by μ' will be mapped onto ontic states preparable by μ by γ (since preparing μ' then applying γ is equivalent to preparing μ). Similarly, γ maps ontic states preparable by $|0\rangle$ onto ontic states preparable by $|\phi\rangle$. Thus, ontic states in the overlap of μ' and $|0\rangle$ will be mapped onto the overlap of μ and $|\phi\rangle$, implying that the overlap as measured by ϖ cannot decrease under the action of γ .

It suffices to prove that for any measurable $\Omega \in \Sigma$ satisfying $\nu(\Omega) = 1, \forall \nu \in \Delta_{|\phi\rangle}$ there exists some measurable $\Omega' \in \Sigma$ such that $\chi(\Omega') = 1, \forall \chi \in \Delta_{|0\rangle}$ and $\mu(\Omega) \geq \mu'(\Omega')$.

For any $\chi \in \Delta_{|0\rangle}$ there is some $\nu \in \Delta_{|\phi\rangle}$ such that $\chi \overset{\gamma}{\rightsquigarrow} \nu$. So for any such Ω Eq. (2.3) gives

$$1 = \nu(\Omega) = \int_{\Lambda} d\chi(\lambda) \gamma(\Omega | \lambda). \quad (2.18)$$

Letting $g(\lambda) \stackrel{\text{def}}{=} \gamma(\Omega | \lambda)$ then by Lem. 2.1 this implies $\chi(\bar{k}(g)) = 1$ for all $\chi \in \Delta_{|0\rangle}$. Therefore $\Omega' \stackrel{\text{def}}{=} \bar{k}(g)$ is a valid choice of Ω' .

Similarly, therefore

$$\mu(\Omega) = \int_{\Lambda} d\mu'(\lambda) \gamma(\Omega | \lambda) = \int_{\Lambda} d\mu'(\lambda) g(\lambda) \geq \int_{\Omega'} d\mu'(\lambda) g(\lambda) = \mu'(\Omega') \quad (2.19)$$

recalling that $g(\lambda \in \bar{k}(g)) = 1$ by definition. \square

The following property is perhaps the easiest to both understand and prove. It relates arbitrary multipartite overlaps, Eq. (2.8), to the underlying individual overlaps.

Lemma 2.4. *For any finite set $\mathcal{S} \subset \mathcal{P}(\mathcal{H})$ of quantum states and any preparation measure μ*

$$\sum_{|i\rangle \in \mathcal{S}} \varpi(|i\rangle | \mu) \geq \varpi(\mathcal{S} | \mu). \quad (2.20)$$

Proof. The overlap $\varpi(\mathcal{S} | \mu)$ is the probability of a disjunction of events and each $\varpi(|i\rangle | \mu)$ is a probability of one of those events. The result is therefore a simple application of Boole's inequality. \square

This last property relates tripartite asymmetric overlaps to certain quantum measurements. It is a little arbitrary, but will be repeatedly used in the theorems that follow and is therefore useful to prove separately here.

Lemma 2.5. *Consider quantum states $|\psi\rangle, |\phi\rangle, |0\rangle$ and orthonormal basis $\mathcal{B} \supset \{|a\rangle, |b\rangle\}$. Suppose that, when written as superpositions over \mathcal{B} , $|\psi\rangle$ and $|\phi\rangle$ only have common support on $|a\rangle$ and $|b\rangle$ (viz. they are orthogonal except for components in the linear subspace spanned*

by $|a\rangle, |b\rangle$). Similarly, suppose that $|\psi\rangle$ and $|0\rangle$ only have common support on $|a\rangle$ and $|b\rangle$. Then for every $\mu \in \Delta_{|\psi\rangle}$

$$\varpi(|\phi\rangle, |0\rangle | \mu) \leq \mathbb{P}_{\mathcal{B}}(|a\rangle \vee |b\rangle | |\psi\rangle) \quad (2.21)$$

for some basis measurement in \mathcal{B} .

Proof. The gist of the proof is that if ontic state λ is accessible by preparing two quantum states then it may only return measurement results that are compatible with both preparations. So if λ is accessible by preparing both $|\psi\rangle$ and $|\phi\rangle$, then it may only return $|a\rangle$ or $|b\rangle$ in the measurement of \mathcal{B} . Similarly, if λ is accessible by preparing both $|\psi\rangle$ and $|0\rangle$, then it may only return $|a\rangle$ or $|b\rangle$ for a similar measurement. Thus, if one of these λ s is obtained in a preparation of $|\psi\rangle$ then the measurement result is necessarily either $|a\rangle$ or $|b\rangle$. This will now be fully fleshed out.

Consider a measurement $M_{\mathcal{B}}$ of the basis \mathcal{B} and let $f(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_{\mathcal{B}}(|a\rangle | \lambda) + \mathbb{P}_{\mathcal{B}}(|b\rangle | \lambda)$. Let $A \subset \mathcal{B}$ be the set of remaining basis states on which $|\phi\rangle$ has support (i.e. $|\phi\rangle$ is a superposition over $|a\rangle, |b\rangle$, and A) and define $g_{\phi}(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_{\mathcal{B}}(A | \lambda)$. Similarly define $g_0(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_{\mathcal{B}}(B | \lambda)$ for $B \subset \mathcal{B}$ the remaining basis states over which $|0\rangle$ has support.

By these definitions, the following quantum probabilities are known for any $\nu \in \Delta_{|\phi\rangle}$ and $\chi \in \Delta_{|0\rangle}$ by Eq. (2.3)

$$\int_{\Lambda} d\nu(\lambda) (f(\lambda) + g_{\phi}(\lambda)) = 1, \quad (2.22)$$

$$\int_{\Lambda} d\chi(\lambda) (f(\lambda) + g_0(\lambda)) = 1. \quad (2.23)$$

Therefore by Lem. 2.1 it follows that $\nu(\bar{k}(f + g_{\phi})) = \chi(\bar{k}(f + g_0)) = 1$. Since $\bar{k}(f + g_{\phi}) \subseteq \bar{k}(f + g_{\phi} + g_0) \supseteq \bar{k}(f + g_0)$ it follows that $\Omega \stackrel{\text{def}}{=} \bar{k}(f + g_{\phi} + g_0)$ is a measurable subset of Λ for which $\nu(\Omega) = \chi(\Omega) = 1, \forall \nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$.

The desired quantum probability is given by

$$\mathbb{P}_{\mathcal{B}}(|a\rangle, |b\rangle | |\psi\rangle) = \int_{\Lambda} d\mu(\lambda) f(\lambda). \quad (2.24)$$

Since $|a\rangle$ and $|b\rangle$ are the *only* basis states where $|\psi\rangle$ and $|\phi\rangle$ have common support, then $\int_{\Lambda} d\mu(\lambda) g_{\phi}(\lambda) = 0$. Similarly, $\int_{\Lambda} d\mu(\lambda) g_0(\lambda) = 0$ and so it follows that

$$\mathbb{P}_{\mathcal{B}}(|a\rangle, |b\rangle | |\psi\rangle) = \int_{\Lambda} d\mu(\lambda) (f(\lambda) + g_{\phi}(\lambda) + g_0(\lambda)) \quad (2.25)$$

$$\geq \int_{\bar{k}(f+g_{\phi}+g_0)} d\mu(\lambda) (f(\lambda) + g_{\phi}(\lambda) + g_0(\lambda)) \quad (2.26)$$

$$= \mu(\bar{k}(f + g_{\phi} + g_0)) = \mu(\Omega). \quad (2.27)$$

Since, by definition $\mu(\Omega)$ upper bounds $\varpi(|\phi\rangle, |0\rangle | \mu)$ this completes the proof. \square

This does not quite conclude the properties of the asymmetric overlap required for this chapter. Before continuing to the main result it will be necessary to consider how asymmetric overlaps interact with so-called anti-distinguishable quantum states.

2.2.2 Anti-distinguishability

Quantum states are perfectly distinguishable—that is, there is a measurement telling them apart with certainty—if and only if they are mutually orthogonal. Distinguishable states must be also ontologically distinct in order to satisfy Eq. (2.1): if the preparation measures non-trivially overlapped, then any ontic states in that overlap would fail to return *any* consistent outcome in a distinguishing measurement. This makes distinguishable states too restrictive to be very helpful in ontology theorems. The opposite and more subtle concept of *anti-distinguishability* is much more useful in discussions of ontic overlaps³.

A finite set of quantum states $\{|\psi\rangle, |\phi\rangle, \dots\} \subset \mathcal{P}(\mathcal{H})$ is anti-distinguishable if and only if there exists a measurement $M = \{E_{\neg\psi}, E_{\neg\phi}, \dots\}$ such that

$$\langle\psi|E_{\neg\psi}|\psi\rangle = \langle\phi|E_{\neg\phi}|\phi\rangle = \dots = 0, \quad (2.28)$$

i.e. the measurement can tell, with certainty, one state from the set that was *not* prepared. This is the “opposite” to distinguishable since for distinguishable sets there is a measurement that can tell, with certainty, one state which *was* prepared. It has been proved⁴ that if some inner products $a = |\langle\phi|\psi\rangle|^2$, $b = |\langle 0|\psi\rangle|^2$, $c = |\langle 0|\phi\rangle|^2$ satisfy

$$a + b + c < 1, \quad (1 - a - b - c)^2 \geq 4abc, \quad (2.29)$$

then the triple $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ must be anti-distinguishable by a projective measurement.

For the purposes of ontological models, the main utility of anti-distinguishable sets is that they exclude intersections of certain overlaps. This can also be viewed in relation to Lem. 2.4: anti-distinguishable triples guarantee that Eq. (2.20) holds with equality. This will now be proved, together with a rough-but-intuitive argument.

Lemma 2.6. *For any anti-distinguishable triple of quantum states $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ Lem. 2.4 holds with equality, that is for all $\mu \in \Delta_{|\psi\rangle}$*

$$\varpi(|0\rangle, |\phi\rangle | \mu) = \varpi(|0\rangle | \mu) + \varpi(|\phi\rangle | \mu). \quad (2.30)$$

³Anti-distinguishability was introduced in Ref. [CFS02a] under the name “PP-incompatibility” and was given the more informative name of anti-distinguishability in Ref. [Lei14a].

⁴This result was proved in Ref. [CFS02a] but Ref. [BCLM14] points out and corrects a typographical error in their result (the original had the second inequality as a strict inequality, which is incorrect).

Proof. Roughly speaking, this result is fairly easy to see. Probabilities $\varpi(|\phi\rangle | \mu)$ and $\varpi(|0\rangle | \mu)$ correspond to the events of preparing $|\psi\rangle$ via μ and getting an ontic state compatible with preparing $|\phi\rangle$ and $|0\rangle$ respectively. Suppose that these events are not mutually exclusive, *viz.* suppose that μ can prepare an ontic state compatible with preparing *both* $|\phi\rangle$ and $|0\rangle$ at the same time. Such an ontic state would not be able to return a consistent outcome for the anti-distinguishing measurement: being compatible with preparations of all three states rules out all measurement outcomes. This is a contradiction, so the events must indeed be mutually exclusive and therefore the probabilities sum to the probability of their disjunction, which is exactly what Eq. (2.30) requires. This will now be rigorously proved.

Recall that $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ is an anti-distinguishable triple if and only if there is some quantum measurement M with three outcomes $E_{-\psi}, E_{-\phi}, E_{-0}$ such that the outcome of getting $E_{-\psi}$ from a system prepared in state $|\psi\rangle$ is zero and similarly for the other state/outcome pairs. By Eq. (2.3) it therefore follows that for all $\mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$

$$\int_{\Lambda} d\mu(\lambda) \mathbb{P}_M(E_{-\psi} | \lambda) = 0, \quad (2.31)$$

$$\int_{\Lambda} d\nu(\lambda) \mathbb{P}_M(E_{-\phi} | \lambda) = 0, \quad (2.32)$$

$$\int_{\Lambda} d\chi(\lambda) \mathbb{P}_M(E_{-0} | \lambda) = 0. \quad (2.33)$$

To prove that Lem. 2.4 holds with equality it suffices to show that given any $\Omega \in \Sigma$ for which $\nu(\Omega) = \chi(\Omega) = 1$ for all $\nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$, there exists some $\Omega', \Omega'' \in \Sigma$ for which $\nu(\Omega') = \chi(\Omega'') = 1$ for all $\nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$ and

$$\mu(\Omega) \geq \mu(\Omega') + \mu(\Omega''). \quad (2.34)$$

This, together with Lem. 2.4 itself, would prove the desired result since the right-hand side bounds $\varpi(|\phi\rangle | \mu) + \varpi(|0\rangle | \mu)$ from above.

To prove that Eq. (2.34) holds define the following measurable functions from Λ to $[0, 1]$

$$g_{\psi}(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_M(E_{-\phi} | \lambda) + \mathbb{P}_M(E_{-0} | \lambda), \quad (2.35)$$

$$g_{\phi}(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_M(E_{-\psi} | \lambda) + \mathbb{P}_M(E_{-0} | \lambda), \quad (2.36)$$

$$g_0(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_M(E_{-\psi} | \lambda) + \mathbb{P}_M(E_{-\phi} | \lambda). \quad (2.37)$$

Using the fact that, for any $\lambda \in \Lambda$, the sum of probabilities of outcomes for any measurement must be unity it follows that $\mathbb{P}_M(E_{-\psi} | \lambda) = 1 - g_{\psi}(\lambda)$ and similarly for $|\phi\rangle$ and $|0\rangle$.

Therefore Eqs. (2.31, 2.32, 2.33) are equivalent to

$$\int_{\Lambda} d\mu(\lambda) g_{\psi}(\lambda) = 1, \quad (2.38)$$

$$\int_{\Lambda} d\nu(\lambda) g_{\phi}(\lambda) = 1, \quad (2.39)$$

$$\int_{\Lambda} d\chi(\lambda) g_0(\lambda) = 1. \quad (2.40)$$

By Lem. 2.1 it immediately follows that $\nu(\bar{k}(g_{\phi})) = \chi(\bar{k}(g_0)) = 1$ where, recall, ν and χ are arbitrary measures from $\Delta_{|\phi\rangle}$ and $\Delta_{|0\rangle}$ respectively.

With these definitions, consider $\mu(\Omega)$ for any $\Omega \in \Sigma$ satisfying $\nu(\Omega) = \chi(\Omega) = 1$ for all $\nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$.

$$\mu(\Omega) = \int_{\Omega} d\mu(\lambda) \quad (2.41)$$

$$= \int_{\Omega} d\mu(\lambda) \left(g_{\phi}(\lambda) + g_0(\lambda) - \mathbb{P}_M(E_{\neg\psi} | \lambda) \right) \quad (2.42)$$

follows by definition of $g_{\phi,0}$. The last term vanishes by Eq. (2.31), so

$$\mu(\Omega) = \int_{\Omega} d\mu(\lambda) g_{\phi}(\lambda) + \int_{\Omega} d\mu(\lambda) g_0(\lambda). \quad (2.43)$$

By restricting the domain of integration

$$\mu(\Omega) \geq \int_{\Omega \cap \bar{k}(g_{\phi})} d\mu(\lambda) g_{\phi}(\lambda) + \int_{\Omega \cap \bar{k}(g_0)} d\mu(\lambda) g_0(\lambda) \quad (2.44)$$

$$= \mu(\Omega \cap \bar{k}(g_{\phi})) + \mu(\Omega \cap \bar{k}(g_0)) \quad (2.45)$$

recalling that $\forall \lambda \in \bar{k}(g_{\phi}), g_{\phi}(\lambda) = 1$ (and similarly for g_0). Note that as both Ω and $\bar{k}(g_{\phi})$ are measure-one according to any $\nu \in \Delta_{|\phi\rangle}$, it follows that their intersection also satisfies $\nu(\Omega \cap \bar{k}(g_{\phi})) = 1$. Similarly, $\chi(\Omega \cap \bar{k}(g_0)) = 1$. Thus what has been proved is that given any $\Omega \in \Sigma$ such that $\nu(\Omega) = \chi(\Omega) = 1$ for all $\nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$ there exist measurable sets $\Omega' \stackrel{\text{def}}{=} \Omega \cap \bar{k}(g_{\phi})$ and $\Omega'' \stackrel{\text{def}}{=} \Omega \cap \bar{k}(g_0)$ satisfying $\nu(\Omega') = \chi(\Omega'') = 1$ for all $\nu \in \Delta_{|\phi\rangle}, \chi \in \Delta_{|0\rangle}$ and

$$\mu(\Omega) \geq \mu(\Omega') + \mu(\Omega''). \quad (2.46)$$

This is exactly what was to be proved. \square

Together, Lems. 2.2–2.6 form the mathematical scaffolding necessary to prove the main results of this chapter that follow.

2.3 Must Superpositions be Real?

With the technical background of asymmetric overlaps and anti-distinguishable sets it is now possible to properly address the main question of this chapter. Is it possible to have an ontology for quantum systems where superpositions are simply statistical effects—where they are not real?

2.3.1 Distinguishing Ontic Superpositions

In order to even precisely formulate this question, one first needs to carefully define what is meant by “real” superpositions. This was touched on in Sec. 2.1.3 but will now be done precisely.

Before even getting to the ontology, it is important to note that superpositions are defined with respect to some orthonormal basis. To be exact, $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ is a *superposition with respect to orthonormal basis \mathcal{B}* if and only if $|\psi\rangle$ has non-zero inner product with more than one state in \mathcal{B} . This is equivalent to saying that $|\psi\rangle \notin \mathcal{B}$.

Now, suppose that Bob believes superpositions with respect to some particular basis \mathcal{B} are just a statistical effect. That is, he believes that superpositions have no ontology of their own but are emergent effects from the statistics. In terms of ontological models, this means that an ontic state obtained by preparing $|\psi\rangle \notin \mathcal{B}$ cannot be independent of \mathcal{B} . That is, Bob believes that there is no new ontology required to describe $|\psi\rangle$ that was not already present when describing \mathcal{B} . In this case, he is compelled to say that whenever he prepares $|\psi\rangle$ the probability of getting an ontic state obtainable by preparing some state from \mathcal{B} is unity:

$$\varpi(\mathcal{B} | \mu) = 1, \quad \forall \mu \in \Delta_{|\psi\rangle}. \quad (2.47)$$

To be consistent with the language of Sec. 2.1.3, call $|\psi\rangle$ an *epistemic superposition with respect to \mathcal{B}* if and only if Eq. (2.47) holds. Conversely, if Eq. (2.47) is violated then there is a finite probability that preparing $|\psi\rangle$ results in novel ontic states not accounted for by \mathcal{B} . Therefore call $|\psi\rangle$ an *ontic superposition with respect to \mathcal{B}* if and only if it is not epistemic.

2.3.2 Almost All Superpositions are Real

Superpositions epistemic with respect to some basis would have considerable explanatory power. Any seemingly bizarre quantum effect based on using such a superposition would have a neat underlying explanation in terms of the ontology of basis states. So, are such things possible?

Perhaps unfortunately, no. Indeed, the no-go result proved below is just about the strongest rejection of epistemic superpositions one could think of. It is relatively easy, for example, to prove that “not every quantum superposition is epistemic”. What is proved in the following two theorems is the much stronger statement that almost every superposition with respect to any given orthonormal basis is ontic.

This result is presented in two theorems. The first contains the main technical argument.

Theorem 2.1. *Consider a $d > 3$ dimensional quantum system and any orthonormal basis \mathcal{B} of \mathcal{H} . For any $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ such that there is a $|0\rangle \in \mathcal{B}$ satisfying $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$, then there exists some other $|i\rangle \in \mathcal{B}$ and $\mu \in \Delta_{|\psi\rangle}$ such that*

$$\varpi(|i\rangle | \mu) \neq |\langle i|\psi\rangle|^2. \quad (2.48)$$

Proof. The proof proceeds by contradiction. To that end, assume that

$$\varpi(|j\rangle | \mu) = |\langle j|\psi\rangle|^2, \quad \forall |j\rangle \in \mathcal{B}, \forall \mu \in \Delta_{|\psi\rangle}. \quad (2.49)$$

Consider $|0\rangle \in \mathcal{B}$ as described in the theorem’s statement. By choosing the global phase of $|\psi\rangle$ appropriately, another orthonormal basis $\mathcal{B}' = \{|0\rangle\} \cup \{|i'\rangle\}_{i=1}^{d-1}$ can be defined such that

$$|\psi\rangle = \alpha|0\rangle + \beta|1'\rangle + \tau|2'\rangle, \quad (2.50)$$

where $\alpha \in (0, \frac{1}{\sqrt{2}})$ and $\beta \stackrel{\text{def}}{=} \sqrt{2}\alpha^2$. Note that $\alpha = \langle 0|\phi\rangle$ can be taken to be real and positive without loss of generality, as states in $\mathcal{P}(\mathcal{H})$ are equivalent up to global phase factors. This is always possible since $|\alpha|^2 + |\beta|^2 = \alpha^2(1 + 2\alpha^2) < 1$ for every such α . With respect to the same \mathcal{B}' define

$$|\phi\rangle \stackrel{\text{def}}{=} \delta|0\rangle + \eta|1'\rangle + \kappa|3'\rangle, \quad (2.51)$$

where $\delta \stackrel{\text{def}}{=} 1 - 2\alpha^2$ and $\eta \stackrel{\text{def}}{=} \sqrt{2}\alpha$. This is always possible since $|\delta|^2 + |\eta|^2 = (1 - 2\alpha^2)^2 + 2\alpha^2 < 1$ for all α and thus an appropriate κ always exists.

The above construction has been chosen such that

- $|\langle 0|\psi\rangle|^2 = \alpha^2 = |\langle \phi|\psi\rangle|^2$ so that there exists a unitary operator U for which $U|0\rangle = |\phi\rangle$ and $U|\psi\rangle = |\psi\rangle$; and
- the inner products $|\langle 0|\psi\rangle|^2$, $|\langle \phi|\psi\rangle|^2$, $|\langle 0|\phi\rangle|^2$ satisfy Eq. (2.29) and, therefore, the triple $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ is anti-distinguishable.

Choose any preparation measure $\mu' \in \Delta_{|\psi\rangle}$ and any stochastic map $\gamma \in \Gamma_U$. Let $\mu \in \Delta_{U|\psi\rangle} = \Delta_{|\psi\rangle}$ be the preparation measure such that $\mu' \xrightarrow{\gamma} \mu$. Then

$$\varpi(|\phi\rangle, |0\rangle | \mu) = \varpi(|\phi\rangle | \mu) + \varpi(|0\rangle | \mu) \quad (2.52)$$

$$\geq \varpi(|0\rangle | \mu') + \varpi(|0\rangle | \mu) \quad (2.53)$$

$$= 2|\langle 0|\psi\rangle|^2 = 2\alpha^2 \quad (2.54)$$

where the first line follows from anti-distinguishability and Lem. 2.6, the second from Lem. 2.3, and the third from Eq. (2.49).

Now consider that in any basis measurement M of \mathcal{B}' , $|0\rangle$ and $|1'\rangle$ are the only outcomes compatible with preparations of either both $|\psi\rangle$ & $|\phi\rangle$ or both $|\psi\rangle$ & $|0\rangle$. So by Lem. 2.5, $\mathbb{P}_M(|0\rangle \vee |1'\rangle | |\psi\rangle) \geq \varpi(|0\rangle, |\phi\rangle | \mu)$. Therefore, in order to reproduce quantum predictions and satisfy Eq. (2.3),

$$\varpi(|0\rangle, |\phi\rangle | \mu) \leq |\langle 0|\phi\rangle|^2 + |\langle 1'|\psi\rangle|^2 = \alpha^2 + 2\alpha^4. \quad (2.55)$$

Combining Eqs. (2.54, 2.55) one finds $\alpha \geq \frac{1}{\sqrt{2}}$. This is a contradiction, since $\alpha \in (0, \frac{1}{\sqrt{2}})$ by construction. Therefore, Eq. (2.49) must be false, implying that there exists some $|i\rangle \in \mathcal{B}$ for which $\varpi(|i\rangle | \mu) \neq |\langle i|\psi\rangle|^2$ for some $\mu \in \Delta_{|\psi\rangle}$. \square

Theorem 2.1 is the central idea that allows the no-go proof for epistemic superpositions. Comparing it to Lem. 2.2, Thm. 2.1 states that the ontic overlap between $|i\rangle \in \mathcal{B}$ and $|\psi\rangle \notin \mathcal{B}$ cannot be maximal. That is, like many of the results discussed in Sec. 2.1.4, this is a bound on ontic overlaps at heart.

Using Thm. 2.1 the main result of this chapter can be proved fairly easily.

Theorem 2.2. *Consider a quantum system of dimension $d > 3$ and define superpositions with respect to any orthonormal basis \mathcal{B} . Almost all quantum superposition states $|\psi\rangle \notin \mathcal{B}$ are ontic.*

Proof. Let $|\psi\rangle \notin \mathcal{B}$ be any superposition state with respect to \mathcal{B} , such that $\exists |0\rangle \in \mathcal{B}$ for which $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$. By Thm 2.1 it follows that $\varpi(|i\rangle | \mu) \neq |\langle i|\psi\rangle|^2$ for some $|i\rangle \in \mathcal{B}$ and some $\mu \in \Delta_{|\psi\rangle}$ and further by Lem. 2.2 that $\varpi(|i\rangle | \mu) < |\langle i|\psi\rangle|^2$. Using this, with Lem. 2.2 again, also gives

$$\sum_{|j\rangle \in \mathcal{B}} \varpi(|j\rangle | \mu) < \sum_{|j\rangle \neq |i\rangle} |\langle j|\psi\rangle|^2 + |\langle i|\psi\rangle|^2 = 1. \quad (2.56)$$

Finally, note that Lem. 2.4 requires that $\varpi(\mathcal{B} | \mu) \leq \sum_{|j\rangle \in \mathcal{B}} \varpi(|j\rangle | \mu)$ and therefore it is found that

$$\varpi(\mathcal{B} | \mu) < 1. \quad (2.57)$$

This directly contradicts Eq. (2.47) and therefore shows that $|\psi\rangle$ is an ontic superposition with respect to \mathcal{B} . However, $|\psi\rangle$ could be any superposition so long as $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$ for some $|0\rangle \in \mathcal{B}$. This is true for all that are not exact 50 : 50 superpositions over two states of \mathcal{B} . Indeed, it is true of almost all states in $\mathcal{P}(\mathcal{H})$. This completes the proof. \square

This is a powerful result. It applies to any superposition state over any basis, so long as it is not an exact 50 : 50 superposition and $d > 3$. It therefore shows that quantum theory is grossly logically incompatible with epistemic superpositions.

This is the primary result of this chapter. In the sections that follow, it will be seen how variations on this theorem form other strong restrictions on the character of ontology of quantum systems, as well as how they powerfully constrain the ability of classical resources to simulate quantum channels.

2.4 Must any Specific States be Real?

2.4.1 State-Specific Ontology

The primary categories of realist ontologies for quantum systems were briefly introduced in Sec. 2.1.3. These are: ψ -ontic, ψ -epistemic, and maximally ψ -epistemic (a small subset of ψ -epistemic). Now with the machinery of asymmetric overlaps, these can be defined formally.

ψ -ontic ontological models were introduced by saying that the ontic state uniquely identifies the quantum state that was prepared—there is no ontic overlap. Clearly, this means that when preparing any $|\psi\rangle$ there is zero probability of getting an ontic state obtainable by preparing any other $|\phi\rangle \neq |\psi\rangle$. In the language of asymmetric overlaps, ψ -ontic models can therefore be precisely defined as those satisfying

$$\varpi(\nu | \mu) = 0, \quad \forall \nu \in \Delta_{|\phi\rangle}, \mu \in \Delta_{|\psi\rangle}, |\psi\rangle \neq |\phi\rangle, \quad (2.58)$$

or equivalently $\varpi(|\phi\rangle | \mu) = 0$ for all $\mu \in \Delta_{|\psi\rangle}$ and $|\psi\rangle \neq |\phi\rangle$. Since ψ -epistemic models are precisely those which are not ψ -ontic, this also mathematically defines ψ -epistemic models.

Maximally ψ -epistemic models can be formalised similarly. Section 2.1.3 introduced maximally ψ -epistemic models as those where the ontic overlap accounts for all of the Born rule overlap. As noted in Sec. 2.2, the Born rule overlap $|\langle \phi|\psi\rangle|^2$ measures the probability that $|\psi\rangle$ will act like $|\phi\rangle$ and the asymmetric overlap does the same for ontological preparations. Clearly then, a model is maximally ψ -epistemic if and only if

$$\varpi(\nu | \mu) = |\langle \phi|\psi\rangle|^2, \quad \forall \mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|\phi\rangle}. \quad (2.59)$$

By Lem. 2.2, this means that all asymmetric overlaps must be maximal, hence the name.

In Secs. 2.1.3–2.1.5 it was discussed how these mark the extremes of possible ontological models for quantum systems. On the face of it, one might expect *some* ψ -epistemic ontological model (such as in Ref. [BCLM14]) to exist, since Eq. (2.58) is such a strong condition. Similarly, it is relatively easy to rule out all maximally ψ -epistemic models for $d > 2$ since Eq. (2.59) is so restrictive. The ontology theorems that bound overlaps—discussed in Secs. 2.1.4, 2.1.5—attempt to bridge the gap between these extremes by deriving quantitative bounds on overlap measures like the asymmetric overlap. But different papers often use different measures and they all suffer from shortcomings discussed in Sec. 2.1.5.

Another way to reach a more nuanced discussion is to apply the ontic/epistemic dichotomy not just to the ontological models as a whole, but to the quantum states themselves too. That is, a *quantum state* $|\psi\rangle$ is ψ -ontic if and only if it has no ontic overlap with any other quantum state $|\phi\rangle \neq |\psi\rangle$: $\varpi(|\phi\rangle | \mu) = 0$ for all $\mu \in \Delta_{|\psi\rangle}$. Correspondingly, $|\psi\rangle$ is ψ -epistemic if and only if it is not ψ -ontic. Clearly, these are much more fine-grained requirements than the equivalents for the entire ontological model.

A specific state can similarly be defined as maximally ψ -epistemic in two ways. First, $|\psi\rangle$ is *maximally ψ -epistemic with respect to* $|\phi\rangle$ if and only if $\varpi(\nu | \mu) = |\langle\phi|\psi\rangle|^2$ for all $\nu \in \Delta_{|\phi\rangle}$ and $\mu \in \Delta_{|\psi\rangle}$. This is very fine-grained, identifying exactly which pair of quantum states overlap maximally. Second, $|\psi\rangle$ is *maximally ψ -epistemic in the ontological model* if and only if it is maximally ψ -epistemic with respect to all states in $\mathcal{P}(\mathcal{H})$. This is clearly less fine-grained, identifying a point between a maximally ψ -epistemic model and a state that is maximally ψ -epistemic with respect to only one other.

Contextuality was also briefly introduced in Sec. 2.1.3. Reference [Spe05] shows how many types of contextuality can be precisely defined using the ontological models framework. For this thesis, only preparation contextuality needs to be considered in detail.

Preparation contextual models are those that describe operationally equivalent preparations differently. Therefore, for a quantum system, preparation non-contextuality requires that all preparations for the same quantum state $|\psi\rangle$ (all of which are operationally equivalent) are described with the same preparation measure; *i.e.* every $\Delta_{|\psi\rangle}$ is a singleton. Recall that preparation non-contextuality is often simply (and implicitly) assumed for pure states, while preparation contextuality is necessary for mixed states [Sec. 2.1.4].

Similarly to the ψ -ontic models, the definition of preparation non-contextual models is very strong but can be used as a starting point for more subtle forms of contextuality. For example, one can consider preparations that are (non)-contextual with respect to certain sets of preparations. One such form of preparation contextuality is pure state *preparation (non)-contextuality with respect to stabiliser unitaries* of some given state $|\psi\rangle$.

For any quantum state $|\psi\rangle \in \mathcal{P}(\mathcal{H})$, the set of *stabiliser unitaries*, $\mathcal{S}_{|\psi\rangle}$, is formed of those U on \mathcal{H} which leave $|\psi\rangle$ unaffected: $U|\psi\rangle = |\psi\rangle$. That is, $\mathcal{S}_{|\psi\rangle}$ is the stabiliser subgroup of the unitary group with respect to $|\psi\rangle$.

Preparations of $|\psi\rangle$ are non-contextual with respect to stabiliser unitaries of $|\psi\rangle$ if and only if preparations that differ only by the action of a $U \in \mathcal{S}_{|\psi\rangle}$ are identical. In other words, $\mu \overset{\gamma}{\rightsquigarrow} \mu$ for every $\mu \in \Delta_{|\psi\rangle}$ and $\gamma \in \Gamma_U$. Therefore, to be preparation non-contextual with respect to stabiliser unitaries simply requires that these unitaries cannot affect the distribution over Λ for a preparation of $|\psi\rangle$. This may seem an oddly specific thing to assume. However Sec. 2.4.4 will argue that it is quite a natural assumption. Moreover, pure state preparation non-contextuality is often assumed wholesale without question and this is a much weaker assumption.

These definitions start to fill in the ψ -epistemic gulf between ψ -ontic models and maximally ψ -epistemic models. Using them as waypoints facilitates a cleaner discussion of exactly which types of ψ -epistemic models are and are not possible. In particular, the theorems that follow will use them to address the shortcomings discussed in Sec. 2.1.5.

2.4.2 No States Can Be ψ -Epistemic

The first shortcoming of previous ontology theorems, as noted in Sec. 2.1.5, is that they only prove the *existence* of a single pair of quantum states that are not maximally ψ -epistemic, without being able to identify those states. Even if the theorems were extended to prove that more than one pair is not maximally ψ -epistemic, they would still not identify which states fail to be maximally ψ -epistemic. This is a loophole for the epistemic realist, who can still postulate ontological models where the vast majority of states of interest are maximally ψ -epistemic without violating these theorems.

The following theorem addresses this loophole directly. By making the additional assumption of preparation non-contextuality with respect to stabiliser unitaries it proves that any given $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ is not maximally ψ -epistemic with respect to very many other states. In particular, no individual state can be maximally ψ -epistemic at all.

Theorem 2.3. *Consider a $d > 3$ dimensional quantum system and any pair of quantum states $|\psi\rangle, |0\rangle \in \mathcal{P}(\mathcal{H})$ that satisfy $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$. Assume that preparations of $|\psi\rangle$ are non-contextual with respect to stabiliser unitaries of $|\psi\rangle$. For any preparation measure $\mu \in \Delta_{|\psi\rangle}$, the asymmetric overlap is bounded by*

$$\varpi(|0\rangle | \mu) \leq |\langle 0|\psi\rangle|^2 \left(\frac{1}{2} + |\langle 0|\psi\rangle|^2 \right) < |\langle 0|\psi\rangle|^2. \quad (2.60)$$

In particular, this implies that the asymmetric overlap $\varpi(|0\rangle | \mu)$ is strictly less than maximal. Therefore $|\psi\rangle$ is not maximally ψ -epistemic with respect to $|0\rangle$ for any such pair of states in $d > 3$ dimensions satisfying $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$.

The proof closely follows that of Thm. 2.1 though it must be proved separately due to the difference in assumptions. The full proof can be found in Appendix A. An immediate corollary is that no quantum state $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ can be ψ -epistemic for $d > 3$.

The epistemic realist must therefore either accept preparation contextuality with respect to stabiliser unitaries, or accept that reality is a long way from being maximally ψ -epistemic. The extra non-contextuality assumption is very mild. As noted above, it is much weaker than general pure-state preparation contextuality, which is very often assumed. Moreover, Sec. 2.4.4 will argue that it is a natural assumption for any minimally realistic ontological model.

Theorem 2.3 improves upon the previous results in as far as ruling out maximally ψ -epistemic ontologies and closing loopholes for ontologies that are close to maximally ψ -epistemic. However, as a bound on the ontic overlap, Eq. (2.60) is rather weak. The next theorem aims instead to establish a restrictive quantitative bound on ontic overlaps by a similar method.

Theorem 2.4. Consider a $d > 3$ dimensional quantum system and any pair $|\psi\rangle, |0\rangle \in \mathcal{P}(\mathcal{H})$ that satisfy $\alpha \stackrel{\text{def}}{=} |\langle 0|\psi\rangle| \in (0, \frac{1}{4})$. Assume that preparations of $|\psi\rangle$ are non-contextual with respect to stabiliser unitaries of $|\psi\rangle$. For any preparation measure $\mu \in \Delta_{|\psi\rangle}$, the asymmetric overlap must satisfy

$$\varpi(|0\rangle | \mu) \leq \alpha^2 \left(\frac{1 + 2\alpha}{d - 2} \right) \quad (2.61)$$

$$\lim_{d \rightarrow \infty} \varpi(|0\rangle | \mu) = 0 \quad (2.62)$$

and so becomes arbitrarily small as d increases even as α is held constant.

The proof strategy here is similar to that of Thm. 2.3 but modified to make use of the higher available dimensions. It can also be found in Appendix A. Note that this modification necessarily weakens the bound compared to Thm. 2.3 in low dimensions.

While the $1/d$ scaling of Eq. (2.61) is relatively weak compared to previous overlap bounds (some of which scale exponentially [Lei14b, Bra14]), Thm. 2.4 overcomes both limitations noted in Sec. 2.1.5. First, Thm. 2.4 bounds the overlap for many specific pairs of quantum states, just like Thm. 2.3. Second, as the dimension increases, the bound tightens even while the inner products of the states remains constant. In particular then, this implies that in large-dimensional systems many specific pairs of states can only barely overlap at all.

2.4.3 Towards Error-Tolerance and Experiments

Thus far, it has been assumed that quantum statistics must be exactly reproduced by a valid ontological model [Eq. (2.3)]. However, it is impossible to exactly verify this. At most, experiments can demonstrate that quantum probabilities hold to within some finite additive error $\epsilon \in (0, 1]$, as in Eq. (2.4). It is therefore necessary to consider *error-tolerant* versions of the above theorems.

Unfortunately, the asymmetric overlap is an error-intolerant quantity. That is, for any quantum system there is an ontological model that satisfies Eq. (2.4) for any given $\epsilon \in (0, 1]$ but for which every asymmetric overlap is unity.

One can construct such an approximate model by modifying the Beltrametti-Bugajski model [BB95] (that is, quantum theory re-phrased as an ontological model, Sec. 2.1.2). Simply adjust the preparation measures so that each has probability ϵ of preparing a completely random state (according to, for example, the Haar measure [ZS01]) and probability $1 - \epsilon$ of acting as usual. Thus, the model will differ from quantum predictions with probability at most ϵ . However, now every preparation measure can prepare any state with finite probability, so the asymmetric overlaps are all unity.

There is an alternative overlap measure, the symmetric overlap $\omega(|\psi\rangle, |\phi\rangle)$ [Mar12, BCLM14, Lei14b, Bra14, Lei14a], that does not have this problem—it is robust to small errors. The symmetric overlap is based on distinguishability of measures.

Suppose you are given some $\lambda \in \Lambda$ obtained by sampling from either μ or ν (each with equal *a priori* probability). Consider using the optimal strategy to guess which of μ, ν was used. The *symmetric overlap of the measures μ and ν* is defined as twice the probability of guessing incorrectly (and therefore takes values in $[0, 1]$). This is known to correspond to [Mar12, BCLM14, Lei14a]

$$\omega(\mu, \nu) \stackrel{\text{def}}{=} \inf \{ \mu(\Omega) + \nu(\Lambda \setminus \Omega) : \Omega \in \Sigma \}. \quad (2.63)$$

As the names suggest, $\omega(\mu, \nu)$ is necessarily symmetric in its arguments and $\varpi(\nu | \mu)$ is not.

Just as with the asymmetric overlap, it is useful to slightly overload the notation and define the *symmetric overlap of the states $|\psi\rangle$ and $|\phi\rangle$* as

$$\omega(|\psi\rangle, |\phi\rangle) \stackrel{\text{def}}{=} \sup_{\mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|\phi\rangle}} \omega(\mu, \nu). \quad (2.64)$$

Quantum theory provides an upper bound on the symmetric overlap, since any quantum procedure for distinguishing $|\psi\rangle, |\phi\rangle$ is also a method for distinguishing $\mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|\phi\rangle}$ in an ontological model. As $\frac{1}{2} \left(1 - \sqrt{1 - |\langle \phi | \psi \rangle|^2} \right)$ is the minimum average error

probability when distinguishing $|\psi\rangle, |\phi\rangle$ within quantum theory⁵ it follows that $\omega(\mu, \nu) \leq 1 - \sqrt{1 - |\langle\phi|\psi\rangle|^2}$ for every $\mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|\phi\rangle}$ and so

$$\omega(|\psi\rangle, |\phi\rangle) \leq 1 - \sqrt{1 - |\langle\phi|\psi\rangle|^2}. \quad (2.65)$$

This is analogous to Lem. 2.2 for the asymmetric overlap.

With this new machinery, Thm. 2.4 can be modified to bound the symmetric overlap while only assuming that quantum probabilities are reproduced to within some finite additive error.

Theorem 2.5. *Consider the assumptions of Thm. 2.4, but only assume that the ontological model reproduces quantum probabilities to within $\pm\epsilon$ for some $\epsilon \in (0, 1]$, as in Eq. (2.4). The symmetric overlap must satisfy*

$$\omega(|0\rangle, |\psi\rangle) \leq \alpha^2 \left(\frac{1 + 2\alpha}{d - 2} \right) + \frac{(3d^2 - 7d)}{2(d - 2)} \epsilon. \quad (2.66)$$

This bound is tighter than Eq. (2.65) for $d > 5$ and sufficiently small ϵ .

The proof strategy of Thm. 2.4 is closely related to the properties of the asymmetric overlap. Since Thm. 2.5 adapts this method to the symmetric overlap, the fit between method and overlap measure is much less comfortable. This also makes the proof itself—provided in Appendix B—unfortunately long, ugly⁶, and resulting in a looser bound [Eq. (2.66)] than its asymmetric counterpart Eq. (2.61).

The significance of Thm. 2.5 is twofold. First, it allows many of the same conclusions as Thm. 2.4 but in a context with finite error: it demonstrates that many specific pairs of quantum states cannot be maximally ψ -epistemic and must have small ontic overlap in the large- d limit (without those states approaching orthogonality) for sufficiently small error. Being error-tolerant, it opens this conclusion up to experimental investigation. Second, it is as a first-step and proof-of-concept for error-tolerance for Thms. 2.2, 2.3. It does not immediately imply that almost all superpositions are real, but by demonstrating how Thm. 2.4’s arguments can be made robust against small error it suggests that error-tolerant versions of the other theorems of this chapter should also be possible. It also proves that no pure state can be individually maximally ψ -epistemic in an error-tolerant way.

Even so, an error-tolerant version of Thm. 2.2 would require the definition of “ontic superposition” to be modified, since it is currently defined in terms of the asymmetric

⁵By using the Helstrom measurement [WDN⁺12, BCLM14].

⁶It should be noted that the proof in Appendix B uses the simpler non-measure-theoretic version of ontological models in order to keep it as short and legible as possible, the measure-theoretic version would be even longer and uglier.

overlap. This is tackled in Sec. 3.4.1, where such a re-definition is provided and error-tolerant variations on Thms. 2.2, 3.1 are presented.

Theorem 2.5 is probably most valuable as a proof-of-concept for experimental applicability rather than forming the basis of a concrete experimental proposal itself. The mismatch between the proof strategy and the symmetric overlap probably makes the result very non-optimal. For example, take the case where $\alpha = 0.245$ and $d = 6$. This makes the bound of Eq. (2.66) $\approx 0.0224 + 8.25\epsilon$, while the bound of Eq. (2.65) ≈ 0.0305 . In this case, Eq. (2.66) is an improved bound for $\epsilon \lesssim 0.0009$. Such experimental accuracy does not seem completely infeasible with current technology [RDB⁺15], but this is nevertheless a challenging experiment to simply improve on the easy bound of Eq. (2.65). Of course, as d increases precise experiments become more challenging, meaning that the error term of Eq. (2.66) is likely to increase super-linearly with d .

2.4.4 Justifying Preparation Contextuality for Stabiliser Unitaries

Theorems 2.3–2.5 all assume that preparations are non-contextual with respect to stabiliser unitaries for given states, as defined in Sec. 2.4.1. This is an extra assumption beyond the bare ontological models framework. Can such an assumption be justified? What follows is a heuristic argument aiming to do exactly that.

The ontological models framework combines postulated fundamental objective ontology with operational notions. The fundamental ontology is reflected in the idea that ontic states represent actual states of affairs, independently of any other theories an observer might use to describe the same system. On the other hand, the only way to reason about this largely-unspecified ontological level is operationally: how does it respond to preparations, transformations, and measurements that can actually be performed?

An assumption of non-contextuality is an assumption about these operational bridges between our capabilities and the ontology. With this perspective, extra assumptions of non-contextuality can be justified by arguing that they are part of any sensible operational understanding of ontological models.

Any specific operational method for preparing some state $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ may be thought of as a black box which the system is fed into. When the system is fed out of the box, it is promised that the box has prepared the system in state $|\psi\rangle$ according to some specific method. In terms of ontological models, the preparation method corresponds to a measure $\mu \in \Delta_{|\psi\rangle}$ and any such method μ can be considered in terms of such a box.

Suppose you design some experiment which involves preparing $|\psi\rangle$ via a method corresponding to μ . Scientists implementing that experiment would acquire the corresponding

black box to be sure that the method is indeed used. Once prepared, the system will need to be presented to other pieces of apparatus. However, there will always be variation in how the system is treated between preparation and the action of any other apparatus, any amount of motion or passage of time or other (seemingly innocuous) treatment amounts to applying some unitary U to the system. Each scientist will, no doubt, be careful to ensure that the system is not disturbed from its preparation state, so it can be safely assumed that any such U is a stabiliser unitary $U \in \mathcal{S}_{|\psi\rangle}$. However, the point remains that some unknown $U \in \mathcal{S}_{|\psi\rangle}$ is inevitably applied to the system after preparation via μ , and this can never be perfectly accounted for.

Therefore, to analyse the result of the experiment, you have to allow for some unknown $U \in \mathcal{S}_{|\psi\rangle}$ to be applied (via some unknown $\gamma \in \Gamma_U$) after preparation of $|\psi\rangle$ via μ . As a result, on this minimally realistic operational level, an arbitrary preparation distribution μ can never be prepared unscathed; you have to account for the inevitable, unknown, subsequent stabiliser unitary. It is therefore prudent to have the effective preparation measure that you use to describe the experiment be one that is non-contextual with respect to such transformations, allowing the experiment to still be analysed despite the application of an unknown $U \in \mathcal{S}_{|\psi\rangle}$.

One must be careful to only consider operational features that are not, even in principle, impossible to reliably perform. Since the sets of preparation measures for any given quantum states are, in the end, operational in character, one may safely restrict to preparation measures that satisfy certain sensible realistic requirements. The above heuristic argument aims to establish pure state preparation non-contextuality with respect to stabiliser unitaries as such a realistic requirement. It is, however, only a heuristic argument and is not rigorous. In particular, no strong reason has been given for including *all* stabiliser unitaries.

2.5 Communication Bounds from Ontology Results

Foundational ontology results often suggest broader implications in many areas. The most obvious example is, of course, Bell's theorem, which has found applications across all areas of quantum theory as well as inspiring many related results in disparate places [BZ02, BZ17]. Not all foundational results can claim quite such outstanding success, of course. Most commonly they can find applications in quantum information, particularly with relation to communications tasks [Mon12, BHK05, PJO15, LPZ⁺16, Mon13, Mon15, Mon16b, Lei14a]. In this section that pattern will be repeated, with the techniques used above to find ontology results applied to a communication task in quantum information.




Figure 2.3: Schematic representation of the finite communication (FC) protocol for simulating a quantum channel with classical resources. Alice receives a description of quantum state $|\psi\rangle$ and Bob receives a description of quantum measurement M . They share a random value r from a set of possibilities \mathcal{R} . Alice sends a classical message $c \in \mathcal{C}$ to Bob through a noiseless channel, who then outputs measurement outcome $E \in M$ based on c and r .

2.5.1 Relating Ontology and Communication

The framework of ontological models has a very natural link to communication tasks in quantum information [PJO15, Mon15]. Indeed, in certain cases direct parallels are known between ontological models and communication protocols [Mon12]. When seeking to apply the techniques of new ontology results to quantum information, communication tasks are therefore a natural place to start.

One of the simplest communication tasks with links to ontological models is the *finite communication* (FC) protocol where Alice and Bob simulate a noiseless n_q -qubit quantum channel with a finite n_c -bit noiseless classical channel. In this task, Alice is given a description of some quantum state $|\psi\rangle$ from a d -dimensional system and Bob is given a description of a quantum measurement M on the same system. The task is for Bob to output some outcomes $E \in M$ with probabilities compatible with quantum theory over many runs, with possibly different $|\psi\rangle$ and M each time.

Clearly, this task can be achieved with one-way communication from Alice to Bob using a noiseless *quantum* channel of enough qubits n_q —Alice sends $|\psi\rangle$ to Bob, who simply measures it. To simulate this with *classical* resources, Alice and Bob are given access to a shared random value $r \in \mathcal{R}$ (taking a new value each run according to some probability measure $\varrho(r)$) and noiseless channel with which Alice can send classical messages $c \in \mathcal{C}$ to Bob from a finite set $|\mathcal{C}| < \infty$. The question is: what quantities of classical resources $|\mathcal{C}|$ and $|\mathcal{R}|$ are required for exact simulation of the n_q -qubit quantum channel? Equivalently, how many classical bits $n_c \geq \log |\mathcal{C}|$, $n_r \geq \log |\mathcal{R}|$ are required to simulate a given number of qubits $n_q \geq \log d$? This protocol is schematically illustrated in Fig. 2.3.

It is easy to see that such a protocol implies a basic ontological model for the same

quantum system. The ontic state space is $\Lambda = \mathcal{C} \times \mathcal{R}$ with ontic states $\lambda = (c, r)$. If Alice sends message c with probability given by $\mathbb{P}(c | |\psi\rangle, r)$ then the corresponding preparation measure for $|\psi\rangle$ is $\mu_{|\psi\rangle}(\lambda) = \mathbb{P}(c | |\psi\rangle, r) \varrho(r)$. If Bob outputs measurement outcome $E \in M$ with probability $\mathbb{P}(E | c, r)$ then this is exactly the corresponding response function $\mathbb{P}_M(E | \lambda)$. This ontological model is missing a description of transformations, which could be provided by extending the FC protocol to include another agent who performs a transformation, but such an extension will not be required here.

A couple of properties of FC protocols can be stated immediately. First, the excess baggage theorem requires ontological models (even those missing transformations) to have infinite ontic state spaces in order to exactly reproduce quantum predictions [Har04]. Therefore, there must be an infinite amount of shared random data $|\mathcal{R}| \geq \infty$ for any such protocol. As a result, the focus is normally on how many classical bits $n_c \geq \log |\mathcal{C}|$ are required to exactly simulate an n_q -qubit quantum system given arbitrary shared random data \mathcal{R} . This approach to FC protocols will be used here.

Given this simplification, one can always assume that all randomness in the protocol comes from the shared random data. That is, any FC protocol for n_q qubits with n_c bits where Alice and Bob can act stochastically implies the existence of an equivalent FC protocol for the same n_q and n_c where Alice and Bob act deterministically given $|\psi\rangle$, M , and $r \in \mathcal{R}$. This is because Alice and Bob can simply obtain any required randomness from r , as \mathcal{R} is arbitrarily large.

Alice's task can therefore be quite simple in an FC protocol. She simply assigns quantum states to classical messages $c \in \mathcal{C}$ based on r in a deterministic way. Her strategy becomes that of an infinite look-up table. Bob's strategy can be similarly described.

Such simulation tasks are not just interesting from a foundational view, but are also important in quantum information [BCMW10]. Considering distributed computing, for example, it is useful to know what sort of advantages quantum channels can offer over classical channels. More generally, an optimal classical FC protocol would provide a natural measure of the power of quantum channels [Mon11].

The best known example of an FC protocol is from Ref. [TB03]. There, an explicit FC protocol is given that simulates a $n_q = 1$ qubit channel with subsequent projective measurement using exactly $n_c = 2$ bits of classical communication (of course, with infinite shared random data). However, no similarly general protocols are known for any quantum dimension $d > 2$. Some partial protocols are known to exist—such as in Ref. [Mon13], which demonstrates how to construct an FC protocol for arbitrary d if Bob can only perform two-outcome projective measurements—but it is not known whether such a protocol is possible for arbitrary projective (or, more generally POVM measurements) even for $d = 3$.

What is known, however, are certain lower-bounds for the required n_c as n_q increases. These will be briefly reviewed in the next section, before a new bound is given with a simple proof based on the results from earlier in the chapter.

2.5.2 A Simple Exponential Bound

Several lower-bounds on the classical bits n_c required to simulate an n_q -qubit quantum channel are known. These require n_c to scale exponentially with n_q and may therefore be seen as “anti-Holevo” results: while the Holevo bound states that a n_q -qubit quantum state can store at most n_q classical bits, these results effectively show that to store n_q qubits one requires at least $\mathcal{O}(2^{n_q})$ classical bits⁷.

The first example is from Ref. [BCT99], following the work of Ref. [BCW98], where it was proved that $n_c \geq c2^{n_q}$ classical bits are required to exactly simulate n_q qubits for some constant $c \approx 0.01$. This was improved in Ref. [Mon11], where the same asymptotic lower bound was obtained with $c = 0.293$ and a postulated improvement (based on a plausible but unproved conjecture) was presented for a $2^{n_q} - 1$ lower bound.

Taking a somewhat different approach, Ref. [Mon16a] proved an $\mathcal{O}(2^{n_q})$ asymptotic lower bound. This has the key advantage of bounding approximate, as well as exact, simulations and also applying to two-way classical communication between Alice and Bob. No precise constant factors for the bound are given, however.

The following theorem shows how the methods used to prove the ontology results in this chapter can also provide a comparable lower bound for the exact FC protocol. A key advantage of this theorem is its simplicity. Even a cursory glance at Refs. [BCT99, Mon11, Mon16a] will show that their proofs are often very mathematically involved. It is comparatively simple to prove the following exponential bound.

Theorem 2.6. *For any $n \in \mathbb{Z}^+$, there is a quantum system of $n_q = \mathcal{O}(\log n)$ qubits such that any FC protocol exactly simulating an n_q -qubit channel requires at least n_c classical bits of communication bounded by*

$$n_c \geq 2^{n_q + \mathcal{O}(1)} - 1. \quad (2.67)$$

Proof. This result follows quickly by borrowing a result from the study of quantum fingerprinting [BCWW01]. In quantum fingerprinting, the aim is to relate n -bit classical bit strings to quantum states that have bounded Born rule overlap in such a way that n scales exponentially in the number of qubits required.

⁷This observation is taken from Ref. [Mon16a], where the “anti-Holevo” moniker is attributed to Tony Short.

The exact result required here is from Ref. [BCWW01, Thm. 2]. This establishes that, for any $n \in \mathbb{Z}^+$ there is a set of $N \stackrel{\text{def}}{=} 2^n$ quantum states $\mathcal{F} = \{|F_x\rangle\}_x$ of $n_q = \log n + \mathcal{O}(1)$ qubits such that $|\langle F_x | F_y \rangle|^2 \leq \frac{1}{4}$ whenever $x \neq y$. It is simple to see that these Born rule overlaps ensure that every triple $\{|F_x\rangle, |F_y\rangle, |F_z\rangle\}$ of unequal states from this set is anti-distinguishable by Eq. (2.29).

Consider an FC protocol for this system of $\log n + \mathcal{O}(1)$ qubits. For any value $r \in \mathcal{R}$, Alice can assign a maximum of two states from \mathcal{F} to each message $c \in \mathcal{C}$. If she were to assign three states from \mathcal{F} to the same message $c \in \mathcal{C}$ then Bob would not be able to correctly simulate the corresponding anti-distinguishable measurement. Therefore, the number of messages required for the protocol must be $|\mathcal{C}| \geq N/2$.

Since $N = 2^n$ is a power of two, the number of required classical bits is $n_c \geq \log |\mathcal{C}| = n - 1$. It therefore immediately follows that $n_c \geq 2^{n_q + \mathcal{O}(1)} - 1$. \square

It is clear that this bound comes from an easily-understood property of quantum states. That is, there exist sets of quantum states with bounded Born-rule overlaps that are exponential in the Hilbert space dimension (that is, \mathcal{F} from the theorem). It is this fact that directly prevents storing quantum states in a small number of classical bits even with shared random data.

This proof also has methodological advantages over the previous bounds found in Refs. [BCT99, Mon11, Mon16a], which will now be discussed.

2.5.3 Generating Further Bounds

The proof of Thm. 2.6 is very easy to understand compared to its peers, while giving a comparable bound. But more than this, careful examination of the proof reveals a general method for deriving such bounds from classical error correction codes. This suggests that the proof of Thm. 2.6 may be more important than the result itself, as simply by finding appropriate classical codes this proof will generate more precise bounds.

As noted above, the key component in the proof is a set \mathcal{F} of quantum states such that $N = |\mathcal{F}|$ is exponential in the Hilbert space dimension d and yet all pairs of states from \mathcal{F} have bounded inner product. Existence of an appropriate such set was taken from quantum fingerprinting, specifically Ref. [BCWW01, Thm. 2].

This points to the first way this proof can generate more bounds. If one finds such a quantum fingerprinting set \mathcal{F} where Born-rule overlaps are bounded $\leq \frac{1}{4}$ and with good scaling of N with d then this yields a bound exactly as in Thm. 2.6. I am, however, unaware of any better sets for this purpose than those used above.

The second way that the proof can generate more bounds is perhaps more interesting. Existence of the \mathcal{F} used in Thm. 2.6 is known due the existence of certain classical error-correction codes. Specifically, any code $E : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with $m = \mathcal{O}(n)$ such that the Hamming distance between unequal code words is greater than $m/4$ yields a quantum fingerprinting set \mathcal{F} with the properties required by Thm. 2.6. The states of this set are in a $(d = m)$ -dimensional Hilbert space and of the form

$$\frac{1}{\sqrt{m}} \sum_{\alpha=1}^m (-1)^{E_{\alpha}(x)} |\alpha\rangle, \quad \forall x \in \{0, 1\}^n \quad (2.68)$$

where $\{|\alpha\rangle\}_{\alpha}$ is an arbitrary orthonormal basis of the system and $E_{\alpha}(x)$ is the α th bit of the code word $E(x)$ [BCWW01]. Clearly, there are $2^n = \mathcal{O}(2^m) = \mathcal{O}(2^d)$ states in this set and they can be verified to have Born rule overlap $\leq \frac{1}{4}$.

Any classical error correcting code of this form (or a sufficiently similar form) will generate an exponential communication bound via the proof of Thm. 2.6. Reference [BCWW01] notes an existence proof for such codes, but specific examples would give specific numerical FC protocol bounds.

This means that Thm. 2.6 has the potential to address one of the shortcomings of previous exponential bounds for FC protocols. That is, those bounds are not particularly useful in low-dimensional settings. For example, the $n_c \geq 0.293 \times 2^{n_q}$ bound only starts to exceed the trivial bound of $n_c \geq n_q$ at $d = 2^{n_q} = 16$ dimensions. Such low-dimensional cases are important for potential experiments, where it is currently unfeasible to test very large-dimensional systems.

So although Thm. 2.6 does not currently provide an exact bound in any dimension, by finding appropriate classical error-correction codes exact bounds will immediately follow. This is in marked contrast to the proofs of existing bounds. In Ref. [Mon16a] only an asymptotic bound is given without a precise form, while in Refs. [BCT99, Mon11] the proofs do not appear to admit easy modification to provide tighter bounds.

2.6 Summary

This chapter tackled the ontology of quantum states directly. By considering what types of realist ontologies are possible, one is naturally led to using the ontological models framework, introduced in Sec. 2.1.1. Broadly, there are two types of realist ontology for quantum states: those where the ontology is similar to the description in quantum theory and those where ontological uncertainty enables potentially more elegant ontologies. The second type is preferred by the “epistemic realist” [Sec. 2.1.2] and raises the question as

to what exactly can be gained through ontological uncertainty in terms of elegance and explanatory power.

The family of results that attempt to answer these questions are known as “ontology theorems” [Sec. 2.1.4]. Due to the restrictions from Bell’s theorem and the PBR theorem, modern ontology theorems typically aim to apply to any single-system ontological model, without specifying how individuals combine into multipartite systems. Within this scope, ψ -epistemic models that reproduce quantum theory have been exhibited but maximally ψ -epistemic models have been proved to be impossible for $d > 2$ dimensions. This leaves the question as to exactly how close to maximally ψ -epistemic ontological models can get before they need to violate quantum predictions. The “overlap theorems” [Sec. 2.1.4] that aim to restrict epistemic ontological models severely all share some shortcomings which appear to leave the epistemic realist plenty of room to consider ontologies that are close to maximally ψ -epistemic, though not exactly [Sec. 2.1.5].

One aspect of quantum state ontology that had not been considered in detail is the ontology of superposition states. Since quantum superpositions inherit all of their properties from their underlying basis states, it is natural to ask whether their ontology can similarly only depend on ontic states accessible to the basis. This question was considered in detail in Sec. 2.3, where Thm. 2.2 proves that, for $d > 3$, almost every superposition is necessarily ontic. That is, any epistemic realist account of quantum theory must include ontic features corresponding to superposition states and the unfortunate cat cannot be put out of its misery.

By adapting the methods of Thm. 2.2, it was possible to obtain more general overlap ontology theorems in Sec. 2.4. By making a very mild extra non-contextuality assumption, Thms. 2.3, 2.4 were proved—between them addressing the shortcomings of other overlap theorems noted in Sec. 2.1.5. In making explicit use of an assumption beyond the bare ontological models framework, they are technically weaker results than some previous ontology theorems. However, the assumption is weak and also arguably natural [Sec. 2.4.4]. In particular, it is much weaker than an assumption that is very often implicitly made. Between them, Thms. 2.3, 2.4 prove that quantum theory is incompatible with very many pairs of states being maximally ψ -epistemic (in particular, any given quantum state cannot be individually maximally ψ -epistemic) and that in large-dimensional systems many pairs of states cannot have any substantial ontic overlap at all.

Theorems 2.2–2.4 were all based on the asymmetric overlap, introduced in detail in Sec. 2.2. This way of quantifying ontic overlaps is easily understood and has a clear motivation, but is unfortunately intolerant to error. The problem of how to adapt these results to be robust to experimental error was therefore considered in Sec. 2.4.3 and

Thm. 2.5, which adapts the methods of Thm. 2.4 to apply instead to the error-tolerant symmetric overlap. The resulting bound is somewhat weaker, due to the mismatch between methodology and overlap measure, but gains error tolerance. It still, however, rules out any given quantum state from being maximally ψ -epistemic in $d > 5$ dimensions and provides a bound on ontic overlaps that approaches zero in large dimensions without the quantum states needing to approach orthogonality (for small enough error).

Often, new results in the foundations of quantum theory can be used to obtain parallel results in quantum information. In Sec. 2.5, the question of how the methods used in this chapter might affect communication abilities was considered. The result, Thm. 2.6, was a simple argument proving that an exponentially large classical channel is required to perfectly simulate a quantum channel, even when arbitrary pre-shared random data is available. While this result does not yet improve on the best known such bounds [Mon16a, Mon11] (asymptotically, they are equivalent), it does provide a general recipe from which these bounds can be generated, given appropriate classical error-correction channels.

This chapter proved some powerful restrictions on the types of ontological model that can reproduce quantum predictions, paying particular attention to the reality of quantum superpositions. In the next, these ideas will be applied to the concept of macro-realism, which will also lead to an error tolerant variation on Thm. 2.2 in Sec. 3.4. A full discussion of the meaning and impact of these results will therefore be deferred until Sec. 3.5.

3

Ontology and Macro-realism

Chapter 2 began to address the ontology of quantum states using ontological models. The aim was to derive ontology theorems in the tradition of Bell and PBR—exactly the task the ontological models framework was designed for (being a refinement of Bell’s original approach to “hidden variable” models). This chapter will use those same methods to tackle the issue of macro-realism, which (despite the name) is not normally considered using ontological models. The main result of this will be a no-go theorem for macro-realism in quantum theory that closes loopholes in the original approach due to Leggett and Garg.

3.1 The Meaning of Macro-realism

The concept of macro-realism was introduced to the study of quantum theory by Leggett & Garg alongside their eponymous inequalities [LG85]. The *Leggett-Garg inequalities* (LGIs) are inequalities on observed measurement statistics that are derived by assuming a particular form of macro-realism and can be violated by measurements on quantum systems. The purpose of the LGIs is therefore to prove that quantum theory and macro-realism are incompatible. However, since its introduction the exact meaning of “macro-realism” has been the subject of debate [Bal87, LG87, Leg88, Leg02a, Leg02b, KB13, MT17]. The purpose of this section is to clarify the meaning of macro-realism, though for the sake of brevity some details will have to be omitted. A more thorough account can be found in Ref. [MT17].

3.1.1 Introducing Macro-realism

Macro-realism is an ontological position. Loosely, macro-realism is the philosophical requirement that certain “macroscopic” quantities always possess definite values. As such, macro-realism is a great candidate for analysis with ontological models.

By using ontological models it is possible to illuminate and classify various definitions of macro-realism precisely (Sec. 3.2.2, following Ref. [MT17]). This analysis will reveal some fundamental loopholes in the Leggett-Garg argument for the incompatibility of quantum theory and macro-realism [Sec. 3.2.4]. In particular, it will show that violation of the LGIs serves only to rule out one sub-category of macro-realist models and that there are other macro-realist models of quantum theory which are compatible with the LGIs. These loopholes are not experimental but logical; the only way to close them is to fundamentally change the argument.

The main result of this chapter is a stronger theorem for the incompatibility between quantum theory and macro-realism [Sec. 3.3]. This theorem closes a loophole in the Leggett-Garg argument by establishing that quantum theory is incompatible with a larger subset of macro-realist models. It does not prove incompatibility of quantum theory with *all* macro-realist models since that is impossible due to existing counter-examples [Sec. 3.2.4]. The theorem proceeds in a very different manner than the Leggett-Garg argument and is related to Thm. 2.2. It thereby circumvents many of the controversies of the original Leggett-Garg approach.

Macro-realism is of interest to experimentalists as well as theorists. There has been a surge of recent work on experimental verification of LGI violation [WKZ⁺17, KKY⁺16, HM16, ZHLG15, GCKGM⁺12] and in particular on noise-tolerance and closing experimental loopholes. At face value, the main theorem presented in Sec. 3.3 will not be suitable for experimental investigation, but Sec. 3.4 will follow one route to error-tolerance for experiments. This will also enable an error-tolerant variation of Thm. 2.2. Further discussion of the experimental relevance of these results will be deferred until Sec. 3.5.

It should be noted that mathematically there is no meaning to the stipulation that macro-realism is about “macroscopic” quantities, as opposed to other physical quantities that aren’t “macroscopic”. Philosophically, however, it is easy to understand the desire for macro-realism applying to “macroscopic” quantities. The types of physical quantity that humans experience are all considered macroscopic and they certainly appear to possess definite values. On the other hand, it is much easier to imagine that microscopic quantities that aren’t directly observed behave in radically different ways. So while there is nothing in the structure of quantum theory to pick-out “macroscopic” versus “microscopic”,

the motivation for considering macro-realism does come from considering macroscopic quantities, hence the name.

3.1.2 Defining Macro-realism

Exactly what is meant by “macro-realism” has been a subject of contention ever since its introduction alongside the LGIs. This controversy has fed into more recent work on understanding the violation of the LGIs [CK15, CK16, Bac15]. In Ref. [MT17], uses of the term “macro-realism” are analysed and the concept is illuminated using ontological models. One result of that paper is that the “macro-realism” intended by Leggett and Garg, as well as many subsequent authors, can be made precise in a reasonable way with the definition:

“A macroscopically observable property with two or more distinguishable values available to it will at all times determinately possess one or other of those values.” [MT17]

Throughout this chapter, “distinguishable” will be taken to mean “in principle perfectly distinguishable by a single measurement in the noiseless case”. Note that macro-realism is defined with respect to some specific property Q . A macro-realist model generally will be macro-realist for some properties and not others. This property will have values $\{q\}$ and to be “observable” must correspond to at least one measurement M_Q with corresponding outcomes E_q which faithfully reveals the underlying macro-realist value q .

Reference [MT17] fleshes out this definition using ontological models and as a result describes three sub-categories of macro-realism. In order to discuss these it will be necessary to first define an *operational eigenstate* in ontological models.

An operational eigenstate Q_q of any value q of an observable property Q is a set of preparation procedures $\{P_q\}$. This set is defined so that immediately following any P_q with any measurement of the quantity Q will result in the outcome E_q with certainty. In other words, an operational eigenstate is simply an extension of the concept of a quantum eigenstate to ontological models: the preparations which, when appropriately measured, always return a particular value of a particular property. Note that if two values q, q' have operational eigenstates then they can sensibly be called “distinguishable”, since any system prepared in a corresponding operational eigenstate can be identified to have one value and not the other with certainty.

3.2 The Leggett-Garg Inequalities

The LGIs are inequalities on the outcomes of certain experiments. The aim of the Leggett-Garg argument is to derive them by assuming macro-realism so that if measurements in quantum theory violate these inequalities then quantum theory must be incompatible with macro-realism. Quantum theory certainly predicts measurements that violate the LGIs and so must be incompatible with at least one of the assumptions needed to derive them. Whether or not the Leggett-Garg argument proves the incompatibility of quantum theory and macro-realism therefore rests on exactly what assumptions are required to derive the LGIs.

3.2.1 Outline of the Argument

The Leggett-Garg argument has a similar structure to Bell's theorem and the LGIs themselves also bear striking resemblance to some Bell inequalities. Operationally, however, the approaches of Bell and Leggett-Garg are quite different. The LGIs only require a single system that is measured several times in sequence. A thorough and complete discussion of LGI derivations is inappropriate here. What follows is rather a sketch sufficient to give context to the remarks later in the chapter. A more extensive discussion can be found in, *e.g.*, Refs. [MT17, ELN14].

Consider measuring a two-valued property Q for a system at three times $t_1 < t_2 < t_3$ in sequence. Label the outcome values $Q_{1,2,3}^{[123]} \in \{-1, +1\}$ for the first, second, and third measurements respectively. The superscript [123] labels this as an experiment where measurements are performed at all three times. On any run of this experiment, it is simple to verify that $Q_1^{[123]}Q_2^{[123]} + Q_1^{[123]}Q_3^{[123]} + Q_2^{[123]}Q_3^{[123]}$ can only equal -1 or 3 . Clearly, taking the average over many runs gives

$$-1 \leq \langle Q_1^{[123]}Q_2^{[123]} \rangle + \langle Q_1^{[123]}Q_3^{[123]} \rangle + \langle Q_2^{[123]}Q_3^{[123]} \rangle \leq 3. \quad (3.1)$$

Now assume that Q is a macro-realist quantity. This means that Q always has some value, whether or not it is measured. The same argument can therefore also be run where no measurements are actually made. Letting $Q_{1,2,3}^*$ be the underlying values of Q in an experiment where no measurements are made, then $Q_1^*Q_2^* + Q_1^*Q_3^* + Q_2^*Q_3^*$ can only equal -1 or 3 .

Since the underlying macro-realist values of Q can always be revealed faithfully by a measurement, it follows that $Q_1^* = Q_1^{[123]}$. However, it does not immediately follow that $Q_2^* = Q_2^{[123]}$, while it does follow that $Q_2^* = Q_2^{[23]}$ (*i.e.*, where no measurement is performed at t_1). The reason for this is simple, the very act of measuring at t_1 could change the

underlying value of Q at t_2 compared to not having measured at t_1 . Similar comments hold for t_3 . The underlying value of Q revealed by a measurement will generally depend on whether any measurements have occurred before.

So consider a final assumption: *non-invasive measurability*. Suppose that measurements of Q do not affect subsequent underlying values of Q . Now it follows, for example, that $Q_2^* = Q_2^{[12]}$ since the measurement, or not, at t_1 has been assumed not to affect the underlying value at t_2 revealed by $Q_2^{[12]}$. With this non-invasive measurability, one obtains the LGIs

$$-1 \leq \langle Q_1^{[12]} Q_2^{[12]} \rangle + \langle Q_1^{[13]} Q_3^{[13]} \rangle + \langle Q_2^{[23]} Q_3^{[23]} \rangle \leq 3. \quad (3.2)$$

To recap: the LGIs of Eq. (3.2) are inequalities on the outcomes of three different types of experiments labelled [12], [13], and [23]. For experiment [13], measurements of Q are made at t_1 and t_3 only and similarly for [12] and [23]. Equation (3.2) has been derived by assuming both macro-realism and non-invasive measurability of Q . There are many quantum experiments of this form that can violate Eq. (3.2) [LG85] so quantum theory must be incompatible with either macro-realism or non-invasive measurability for Q .

That non-invasive measurability is required to derive LGIs has been known since their introduction [LG85]. Some go so far as to include non-invasive measurability as part of their definition of macro-realism to avoid having to deal with it explicitly, relegating definitions like that in Sec. 3.1.2 to “macro-realism *per se*”. The final part of the Leggett-Garg argument has always been to contend that this non-invasive measurability is a necessary consequence of macro-realism.

While this concludes the sketch of the traditional Leggett-Garg argument, the problem of deriving non-invasive measurability from macro-realism will be returned to in Sec. 3.2.4.

3.2.2 Sub-classes of Macro-realism

To clearly discuss exactly what is ruled out by the Leggett-Garg argument—and how it might be improved—it is necessary to understand three sub-categories of macro-realism. These were identified in Ref. [MT17] by considering the definition of Sec. 3.1.2 in terms of ontological models for the system. The resulting sub-categories are then categories of ontological models with particular properties.

The three sub-categories of macro-realism for some quantity Q are:

1. *Operational eigenstate mixture macro-realism (EMMR)* – The only preparations in the model are operational eigenstates of Q or statistical mixtures of operational eigenstates. That is, every preparation measure can be written in the form $\nu = \sum_q \sum_i c_{q,i} \mu_{q,i}$ where each $\mu_{q,i}$ is a preparation measure for an operational eigenstate

for q and $\{c_{q,i}\}$ are positive reals summing to unity. Note that this means the space of ontic states Λ need only include ontic states accessible by preparing some operational eigenstate of Q , as no other ontic states can ever be prepared.

2. *Operational eigenstate support macro-realism (ESMR)* – Like EMMR, every ontic state $\lambda \in \Lambda$ is accessible by preparing some operational eigenstate but, unlike EMMR, there are preparation measures in the model that are not statistical mixtures of operational eigenstate preparations for Q . That is, if $\Omega \in \Sigma$ satisfies $\mu_q(\Omega) = 1$ for every operational eigenstate preparation μ_q of Q , then every preparation measure ν in the model also satisfies $\nu(\Omega) = 1$. Moreover, the model has at least one preparation measure not in the mixture form required by EMMR. In other words, if you’re certain to prepare an ontic state from some subset Ω when preparing operational eigenstates, then you’re also certain to prepare an ontic state from Ω from any other preparation measure in the model.
3. *Supra eigenstate support macro-realism (SSMR)* – Every ontic state λ in the model will produce some specific value q_λ of Q when a measurement of Q is made, but some of those ontic states are not accessible by preparing any operational eigenstates of Q . That is, for every $\lambda \in \Lambda$ there is some value q_λ of Q such that $\mathbb{P}_M(E_{q_\lambda} | \lambda) = 1$ whenever Q is measured. Moreover, there exists some $\Omega \in \Sigma$ and preparation measure ν such that $\nu(\Omega) > 0$ while $\mu_q(\Omega) = 0$ for every operational eigenstate preparation measure μ_q .

To help unpack these definitions, they are illustrated in Fig. 3.1.

In each of these cases, every ontic state λ (up to possible measure-zero sets of exceptions) is associated with a specific value q_λ of Q , such that it can be sensibly said that λ “possesses” q_λ . This is why they are all considered types of macro-realism. Consider this for each case in turn.

In an EMMR model, every preparation can be read as a probabilistic choice between operational eigenstate preparations. Depending on which preparation is chosen, the resulting ontic state λ therefore “possesses” the value q for the operational eigenstate.

In an ESMR model, every ontic state λ can be prepared by an operational eigenstate of exactly one value of Q (up to measure-zero sets of exceptions). Similarly, therefore, each ontic state “possesses” the corresponding value of Q .

In SSMR models the link between each λ and the corresponding q_λ is explicit. Each λ “possesses” the value q_λ for which $\mathbb{P}_M(E_{q_\lambda} | \lambda) = 1$, as required by the definition. That is, λ “possesses” the value which it must return with certainty in any appropriate measurement.

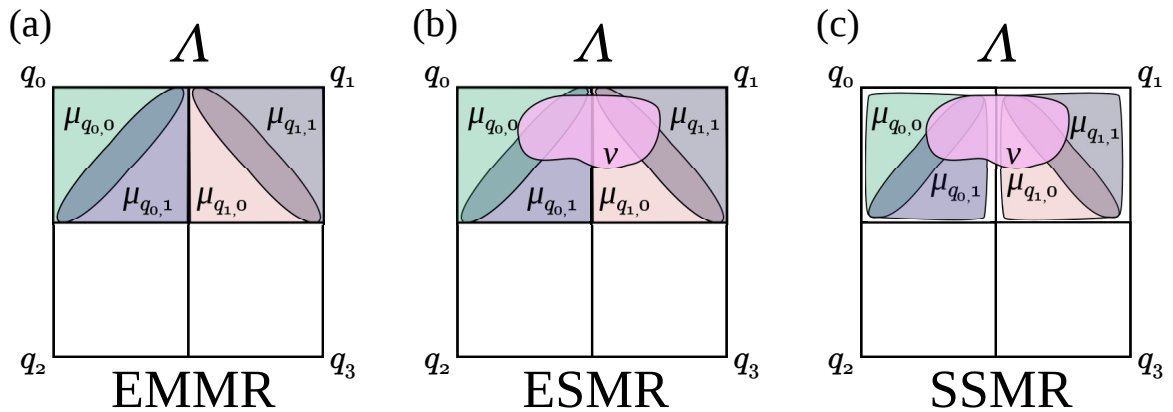


Figure 3.1: Illustration of the three sub-categories of macro-realism as defined in the text. In each case the large square represents the whole ontic state space Λ , the four smaller squares indicate those subspaces of ontic states associated with each value $q_{0..3}$ of some quantity Q , and the shaded regions represent those ontic states accessible by preparing some select preparation measures. Only a few preparation measures are shown while more would exist in a real model, in particular measures have been left out of the q_2 and q_3 boxes to avoid over-cluttering the figures.

(a) illustrates EMMR, where the squares for each q_i contain all ontic states preparable via some operational eigenstate preparation $\mu_{q_i,j}$ and all other allowed preparation measures are simply statistical mixtures of these, *e.g.* $\nu = \frac{1}{3}(\mu_{q_0,0} + \mu_{q_0,1} + \mu_{q_1,0})$ is permissible.

(b) illustrates ESMR, where the state space is exactly as in EMMR, but now more general preparation measures, such as the ν illustrated, are permitted.

(c) illustrates SSMR, where now every λ in the box for q_i must produce outcome q_i in any appropriate measurement of Q , but the operational eigenstates no longer fill these boxes. That is, there are ontic states that lie outside the preparations for operational eigenstates. General preparation measures over the boxes, like ν , are still permitted.

Note that these three sub-categories of macro-realism are defined such that they are mutually exclusive, but they still have a natural hierarchy to them. ESMR can be seen as a less restrictive variation on EMMR, since you can make an EMMR model into an ESMR model simply by including a single preparation measure that is not a statistical mixture of operational eigenstate preparations (the ontic state space and everything else can remain unchanged). Similarly, SSMR can be seen as a less restrictive variation on ESMR. In ESMR, every ontic state λ can be obtained by preparing an operational eigenstate preparation for a value of Q . By definition of operational eigenstate it follows that a measurement of Q will therefore return some specific value for each ontic state (up to measure-zero sets of exceptions), which is the primary requirement on the ontic states for SSMR.

3.2.3 Macro-realism for Quantum Systems

Just like ontological models, the concept of macro-realism is—and always should be—logically independent of quantum theory. So far this section has presented macro-realism in this general way. In order to proceed to a precise discussion of the loopholes in the Leggett-Garg argument, it is now necessary to bring macro-realism, ontological models, and quantum theory together.

To do this, consider what can count as a “macroscopically observable” quantity Q . To be observable Q must correspond to some quantum measurement M_Q . Therefore, there is some orthonormal basis \mathcal{B}_Q so that for each value q of Q the corresponding outcome of M_Q is a state in \mathcal{B}_Q . In order to make sense of the definitions Q must also have operational eigenstates for each value q of Q . Fortunately this is straightforward in quantum theory: the states in \mathcal{B}_Q are exactly the operational eigenstates of Q . Moreover, because the elements of \mathcal{B}_Q are orthogonal it follows that preparations corresponding to different values q, q' of Q are distinguishable.

So in quantum theory, macro-realism for quantity Q simply means that there is a basis \mathcal{B}_Q each element of which is an (operational) eigenstate of Q . This is gained simply by applying the definition of macro-realism from Sec. 3.1.2 to quantum theory, the three different sub-categories of macro-realism noted in Sec. 3.2.2 correspondingly define three possible sub-categories of macro-realism in quantum theory.

3.2.4 Loopholes in the Leggett-Garg Argument

The aim of the LGIs has always been to rule out macro-realist ontologies for quantum theory, where the inequalities are violated. However, in light of the above precise definition of macro-realism some loopholes in the argument can be identified. These loopholes all

have to do with the attempt to derive non-invasive measurability from macro-realism as noted in Sec. 3.2.1.

The first loophole is that violation of the LGIs cannot rule out SSMR models of quantum systems. Indeed, no argument that rests on compatibility with quantum predictions can completely rule out SSMR models since there exists a well-known SSMR model for quantum systems that reproduces all quantum predictions: Bohmian mechanics [Boh52a, Boh52b, Bro09].

To see that Bohmian mechanics implies an SSMR ontological model consider, for example, the Bohmian description of a single spinless point particle in three-dimensional space (the argument for more general systems is analogous). Bohmian mechanics has the ontic state as a pair $\lambda = (\vec{r}, |\psi\rangle) \in \mathbb{R}^3 \times \mathcal{P}(\mathcal{H})$ where \vec{r} is the actual position of the particle and $|\psi\rangle$ is the quantum state (or “pilot wave”). Note that the quantum state is part of the ontology here. The “macroscopically observable property” is the position of the particle, \vec{r} , and any sharp measurement of position will reveal the true value of \vec{r} with certainty. Thus, for any ontic state λ there is some value of the macroscopically observable property (that is, \vec{r}) which is obtained with certainty from any appropriate measurement. Thus, Bohmian mechanics provides an SSMR ontological model.

The second loophole is that LGI violation is unable to rule out ESMR ontological models. This also has a counter-example in the form of the Kochen-Specker model for the qubit [KS67], which is an ontological model satisfying ESMR¹. The Kochen-Specker model exactly reproduces quantum predictions for $d = 2$ dimensional systems. As the LGIs can be written in $d = 2$ the Kochen-Specker model must therefore violate them.

The key point is *why* these counter-examples evade the Leggett-Garg argument. As discussed in Sec. 3.2.1, to derive the LGIs one needs to assume non-invasive measurability. If it is not possible to derive or demonstrate non-invasive measurability as a consequence of some type of macro-realism, then it is not possible to derive the LGIs from that type of macro-realism and the Leggett-Garg argument does not apply. It turns out that there is no way to do this by assuming either SSMR or ESMR and therefore Leggett-Garg arguments generically have loopholes for these types of macro-realism [MT17]. Bohmian mechanics and the Kochen-Specker model are both examples: they contain measurement disturbance that violates the non-invasive measurability assumption, while still satisfying SSMR and ESMR respectively. So the crux is that both SSMR and ESMR models can include measurements that don’t disturb the distribution over Λ if the system is prepared in an

¹Strictly speaking, the Kochen-Specker model was not defined with a post-measurement update rule and so cannot deal with sequences of measurements (and therefore Leggett-Garg experiments). However, it is simple to append the obvious update rule “prepare a new state corresponding to the measurement outcome” and this fixes the issue.

operational eigenstate, but still disturb the distribution over Λ for systems prepared in other ways.

On the other hand, EMMR requires that all preparations are statistical mixtures of operational eigenstates. If one can demonstrate experimentally that operational eigenstates are not disturbed by some measurement, then no preparations can be disturbed by that measurement in an EMMR model. It is this which prevents EMMR models from violating the LGIs. A more extensive discussion of this point can be found in Ref. [MT17].

Recent experiments [KKY⁺16, HM16] following Ref. [WM12] have sought to address the “clumsiness” loophole in the Leggett-Garg argument by dropping the assumption of non-invasive measurability, replacing it with control experiments that experimentally serve the same purpose. These approaches follow the Leggett-Garg argument quite closely and show that the disturbance on a general preparation cannot be explained in terms of the disturbances on a statistical mixture of operational eigenstates. As a result, they are still only capable of ruling out EMMR models, as noted in Ref. [KKY⁺16].

So the Leggett-Garg proof only rules out EMMR macro-realism and leaves loopholes for SSMR and ESMR. Moreover, the loophole for SSMR models cannot be fully plugged by any proof because Bohmian mechanics exists as a counter-example. Similarly, the loophole for ESMR cannot be fully plugged in $d = 2$ dimensions since the Kochen-Specker model exists as a counter-example. This leaves a clear question: can the ESMR loophole can be closed by another theorem for any $d > 2$? Answering this question requires a different approach to the Leggett-Garg argument, one that doesn’t use any assumptions about measurement disturbance.

3.3 A Stronger Theorem Against Macro-realism

The Leggett-Garg argument, in its usual form, does not make use of anything like ontological models. Macro-realism is, however, an ontological position. The success of ontological models in investigating other such positions (such as excess baggage, locality, and ψ -ontology) suggests that it may also be useful when investigating macro-realism. Moreover, the three sub-categories of macro-realism of Sec. 3.2.2 were identified in Ref. [MT17] by using an ontological model approach. Therefore, this section will take the ontological techniques developed in Chap. 2 and apply them to macro-realism to obtain a no-go result in quantum theory that is stronger than the traditional Leggett-Garg argument.

3.3.1 Theorem

As discussed in Sec. 3.2.4, the Leggett-Garg argument can rule out EMMR macro-realist models for quantities Q with $n > 1$ distinguishable values. The following theorem improves on this by ruling out both EMMR and ESMR macro-realist models for $n > 3$.

Theorem 3.1. *Quantum theory is incompatible with ESMR or EMMR macro-realist models for quantities Q with $n > 3$ distinguishable values.*

Proof. In ESMR and EMMR ontological models, every ontic state in the model can be obtained through preparing some operational eigenstate of Q . From this it follows that, when preparing any quantum state $|\psi\rangle$ through any $\mu \in \Delta_{|\psi\rangle}$, the probability of getting an ontic state that could have been prepared by preparing a state in \mathcal{B}_Q is unity. Mathematically, $\varpi(\mathcal{B}_Q | \mu) = 1$. Comparing this to Eq. (2.47) it is seen that every $|\psi\rangle \notin \mathcal{B}_Q$ must therefore be an epistemic superposition over \mathcal{B}_Q .

Clearly, if Q has $n > 3$ distinguishable values then $d = |\mathcal{B}_Q| > 3$. So by using Thm. 2.2, one reaches a contradiction.

The assumptions that went in to reaching this contradiction were:

- (a) the ontology satisfies ESMR or EMMR;
- (b) the “macroscopically observable quantity” Q has $n > 3$ distinguishable values;
- (c) the states preparations and measurements used in the proof of Thm. 2.2 are possible;
and
- (d) the ontological model reproduces quantum measurement predictions, as in Eq. (2.3).

Assumptions (a–b) are about the underlying ontological model, whereas assumptions (c–d) are implications of standard quantum theory. The conclusion must therefore be that both ESMR and EMMR ontologies for $n > 3$ are impossible, or quantum theory is incorrect. \square

Since the Leggett-Garg argument is only able to prove incompatibility between EMMR and quantum theory, this is a strict improvement for $n > 3$.

3.3.2 Relation to Other Ontology Results

The proof of this stronger no-go theorem directly used Thm. 2.2 to get its result. Some comments are therefore in order about the exact relationship between Thms. 2.2, 3.1.

The obvious difference between the two is in the quantifiers. To prove Thm. 2.2 a contradiction is required for almost all superposition states. Comparing this to the proof

of Thm. 3.1, one sees that the latter only requires a contradiction for a single superposition state. From this point of view, Thm. 3.1 is weaker than Thm. 2.2 which is why the latter is used to prove the former rather than the other way around. Indeed, once it has been established that ESMR/EMMR ontologies imply epistemic superpositions, it is almost a corollary.

Clearly, the two theorems, in their current forms, are getting at a similar point. It even appears that ESMR/EMMR macro-realism might be simply be subsumed under the banner of epistemic superpositions. But there are also divergences. Foremost of these is that “macro-realism” can (and should) be defined independently from quantum theory (as in Secs. 3.1.2–3.2.2) while the very quantum notion of “superpositions” is required to even talk about “epistemic superpositions”. The two concepts are of very different character.

This difference in character becomes very important when considering extensions beyond Thms. 2.2, 3.1, especially to statements amenable to experimental tests. Both results in their current forms concentrate on properties of pure states. This is entirely natural for Thm. 2.2, since superpositions are pure states. But as macro-realism is defined independently from quantum theory, it would be appropriate to extend Thm. 3.1 using mixed state preparations and more general quantum measurements. More discussion on possible appropriate extensions to each of the theorems will be in Sec. 3.5.

Neither of these theorems are intended to be final, but rather starting points that point to new lines of research. The most obvious next direction is establishing error-tolerant variations that allow for experimental investigation. One method for doing this will follow in the next section.

3.4 An Error-Tolerant Argument

Theorems 2.2, 3.1 are the main results of their respective chapters. However, both proofs rely on requiring ontological models to exactly reproduce quantum theory, that is Eq. (2.3). This means that, in their current form, neither is amenable to experimental investigation. There is always noise and error in experimental results. Experimental results against ESMR/EMMR or epistemic superpositions would carry much more weight, so it is prudent to seek error-tolerant variations of these theorems.

In Sec. 2.4.3 an error-tolerant variation on Thm. 2.4 was presented and used as a proof-of-concept for error-tolerance of the other theorems of Chap. 2. This is a valuable result (and provides scope for experimental tests in its own right) but by too closely following the same proof strategy it does not address the core hurdle to error tolerance for the main results.

As noted in Sec. 2.4.3, the primary reason that these results are not error-tolerant is that they are based on the asymmetric overlap, which becomes useless when there is finite error in reproducing quantum statistics. Worse, the mathematical formalisation of “epistemic superposition”, Eq. (2.47), is the key fact used in proofs of both Thms. 2.2 and 3.1 and is itself in terms of the asymmetric overlap. Thus, to properly find error-tolerant variations on Thms. 2.2 and 3.1, this formalisation needs to be generalised to be error-tolerant.

3.4.1 Error-Tolerance and the ϵ -Asymmetric Overlap

In order to generalise the formalisation of epistemic superpositions and ESMR/EMMR macro-realism, it is necessary to first generalise the asymmetric overlap.

Consider the definitions of the asymmetric overlap in Eqs. (2.5, 2.6, 2.8). Each definition extremises over measurable subsets $\Omega \in \Sigma$ that satisfy $\nu(\Omega) = 1$ for some preparation measure(s) ν . This condition can be relaxed by considering ϵ -typical subsets. For any $\epsilon \in [0, 1)$, $\Omega \in \Sigma$ is ϵ -typical for preparation measure ν if and only if $\nu(\Omega) > 1 - \epsilon$. That is, the probability of ν preparing an ontic state *outside* any given ϵ -typical subset Ω is ϵ . Despite the name, a given ϵ -typical subset Ω does not identify the ontic states inside it as typical as such, but rather it identifies the ontic states *outside* it as *atypical* (where (a)typicality is quantified by ϵ).

ϵ -typical subsets can be used to generalise the asymmetric overlap ϖ to the ϵ -asymmetric overlap ϖ_ϵ , defined for preparation measures as

$$\varpi_\epsilon(\nu | \mu) \stackrel{\text{def}}{=} \inf\{\mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) > 1 - \epsilon\}. \quad (3.3)$$

That is, $\varpi_\epsilon(\nu | \mu)$ is a lower bound probability that μ will produce an ontic state that is in some ϵ -typical subset for ν . Clearly, by taking $\epsilon = 0$ the 0-typical asymmetric overlap is the asymmetric overlap of Eq. (2.5).

How should ϖ_ϵ be interpreted? Recalling that all ontic states outside an ϵ -typical subset Ω are atypical for ν , one sees that $1 - \mu(\Omega)$ is a probability that μ will produce an ontic state that is definitely atypical for ν . Therefore, $1 - \varpi_\epsilon(\nu | \mu)$ is the upper bound probability that μ will produce an ontic state that is definitely atypical for ν . This double-negative interpretation of the ϵ -asymmetric overlap is the precise one, but is also clumsy. In effect, it says that $\varpi_\epsilon(\nu | \mu)$ is the probability that μ will produce an ontic state that is typical for ν (where typicality is quantified by ϵ). However, the nature of measure theory means that this must be understood as a double-negative.

When describing a quantum system, Eq. (2.6) defines the ordinary asymmetric overlap with a quantum state. The same can be done for the ϵ -asymmetric overlap:

$$\varpi_\epsilon(|\phi\rangle | \mu) \stackrel{\text{def}}{=} \inf\{\mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) > 1 - \epsilon, \forall \nu \in \Delta_{|\phi}\}. \quad (3.4)$$

Finally, the definition can also be extended to overlapping with a *set* of quantum states $\mathcal{S} \subseteq \mathcal{P}(\mathcal{H})$, as in Eq. (2.8),

$$\varpi_\epsilon(\mathcal{S} | \mu) \stackrel{\text{def}}{=} \inf \{ \mu(\Omega) : \Omega \in \Sigma, \nu(\Omega) > 1 - \epsilon, \forall \nu \in \Delta_{|\phi\rangle}, \forall |\phi\rangle \in \mathcal{S} \}. \quad (3.5)$$

Now defined, the ϵ -asymmetric overlap can be used to identify an error-tolerant formalisation of epistemic superpositions.

As discussed in Sec. 2.3.1, an epistemic superposition is one that is just a statistical effect, one where the ontology is all accounted for by the underlying basis. That is, somebody who believes $|\psi\rangle$ is an epistemic superposition over \mathcal{B} should believe that any ontic state obtained by preparing $|\psi\rangle$ could also be obtained by preparing some state from \mathcal{B} . In the case of finite error this is clearly an inconsequential belief as there is always some finite probability of preparing any outlandish ontic state in any preparation. An appropriate modified belief would be that any preparation of $|\psi\rangle$ will, with high probability, produce an ontic state that is *typical* for at least one state in \mathcal{B} . That is, the vast majority of the time the resulting ontic state will be one that can typically be obtained by at least one state in \mathcal{B} . Mathematically, this can be characterised with the ϵ -asymmetric overlap:

$$\varpi_\eta(\mathcal{B} | \mu) > 1 - \tau, \quad \forall \mu \in \Delta_{|\psi\rangle} \quad (3.6)$$

for small values η and τ which quantify the strength of the belief. Clearly, taking $\tau = \eta = 0$ this reduces to the noiseless case of Eq. (2.47).

A similar argument can be made for ESMR/EMMR macro-realism for a quantity Q , generalising the proof of Thm. 3.1. The ESMR/EMMR macro-realist believes that every ontic state preparable by any $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ can also be obtained by preparing a state from a basis \mathcal{B}_Q . In the case of finite error this is a trivial belief, since there's a finite probability of preparing any ontic state from any quantum state. A reasonable macro-realist would therefore rather believe that when preparing any $|\psi\rangle$ there's a high probability that the ontic state obtained will be *typical* for at least one state in \mathcal{B}_Q . Mathematically, this implies Eq. (3.6) again, with $\mathcal{B} = \mathcal{B}_Q$.

3.4.2 Properties of the ϵ -Asymmetric Overlap

Just as with the asymmetric overlap, it will be necessary to prove some properties of the ϵ -asymmetric overlap before proceeding to the main theorem. The properties here are mostly generalisations of those proved in Sec. 2.2.1. The reader may find it easier to first read the theorem in the next section, referring back here only when the properties are referenced.

For notational convenience, it is useful to generalise Def. 2.1 in the following way.

Definition 3.1. For any measurable function $g : \Lambda \rightarrow [0, 1]$ let

$$\bar{k}_\epsilon(g) \stackrel{\text{def}}{=} \{\lambda \in \Lambda : g(\lambda) \geq 1 - \epsilon\} \in \Sigma \quad (3.7)$$

for any $\epsilon \in [0, 1)$.

Next, this lemma generalises Lem. 2.1 and will similarly provide a link between integrals over a measure and measurable subsets.

Lemma 3.1. Given any measurable function $f : \Lambda \rightarrow [0, 1]$ and preparation measure ν satisfying $\int_\Lambda d\nu(\lambda) f(\lambda) \geq 1 - \delta$ then

$$\nu(\bar{k}_\kappa(f)) > 1 - \frac{\delta}{\kappa} \quad (3.8)$$

for any $0 \leq \delta \leq \kappa \leq 1$.

Proof. For any $\kappa \in [0, 1)$ then by assumption

$$1 - \delta \leq \int_{\bar{k}_\kappa(f)} d\nu(\lambda) f(\lambda) + \int_{\Lambda \setminus \bar{k}_\kappa(f)} d\nu(\lambda) f(\lambda). \quad (3.9)$$

Considering the first term, clearly $\int_{\bar{k}_\kappa(f)} d\nu(\lambda) f(\lambda) \leq \nu(\bar{k}_\kappa(f))$ as $f(\lambda) \leq 1$. For the second term, note $f(\lambda) < 1 - \kappa$ when $\lambda \in \Lambda \setminus \bar{k}_\kappa(f)$, giving

$$1 - \delta < \nu(\bar{k}_\kappa(f)) + (1 - \kappa)\nu(\Lambda \setminus \bar{k}_\kappa(f)) \quad (3.10)$$

which, recalling $\nu(\Lambda) = 1$, implies the desired result. \square

This first property does not have an analogue for the asymmetric overlap. It simply notes that some ϵ -asymmetric overlaps are more restrictive than others (depending on the values of ϵ).

Lemma 3.2. For any set \mathcal{S} of quantum states and any preparation measure μ ,

$$\varpi_p(\mathcal{S} | \mu) \geq \varpi_q(\mathcal{S} | \mu) \quad (3.11)$$

for any $0 \leq p \leq q < 1$.

Proof. For any Ω_p that is p -typical for \mathcal{S} then

$$\nu_i(\Omega_p) > 1 - p \geq 1 - q, \quad \forall \nu_i \in \Delta_{|i\rangle}, \forall |i\rangle \in \mathcal{S} \quad (3.12)$$

so Ω_p is also q -typical for \mathcal{S} and by definition $\varpi_p(\mathcal{S} | \mu) \geq \varpi_q(\mathcal{S} | \mu)$. \square

With this technical background, the following property generalises Lem. 2.2, showing how measurement probabilities simply bound the ϵ -asymmetric overlap.

Lemma 3.3. *For any state $|\phi\rangle \in \mathcal{P}(\mathcal{H})$ and measurement M with $|\phi\rangle$ as an outcome, if $\mathbb{P}_M(|\phi\rangle | \nu) \geq 1 - \sigma$ for some $\sigma \in [0, 1)$ and all $\nu \in \Delta_{|\phi\rangle}$ then*

$$\varpi_\eta(|\phi\rangle | \mu) \leq \frac{\mathbb{P}_M(|\phi\rangle | \mu)}{1 - \sigma/\eta}, \quad \forall \eta \in (\sigma, 1) \quad (3.13)$$

for any preparation measure μ , where $\mathbb{P}_M(|\phi\rangle | \mu)$ is obtained from the ontological model via Eq. (2.1).

Proof. For some $\eta \in (\sigma, 1)$, let $\Omega \stackrel{\text{def}}{=} \bar{k}_{\sigma/\eta}(\mathbb{P}_M(|\phi\rangle | \lambda))$. For all $\nu \in \Delta_{|\phi\rangle}$,

$$1 - \sigma \leq \mathbb{P}_M(|0\rangle | \nu) = \int_{\Lambda} d\nu(\lambda) \mathbb{P}_M(|\phi\rangle | \lambda) \quad (3.14)$$

implies that $\nu(\Omega) > 1 - \eta$ by Lem. 3.1. Therefore, Ω is η -typical for $|\phi\rangle$, giving

$$\varpi_\eta(|\phi\rangle | \mu) \leq \mu(\Omega) \quad (3.15)$$

$$\leq \frac{\int_{\Omega} d\mu(\lambda) \mathbb{P}(|\phi\rangle | \lambda)}{1 - \sigma/\eta} \quad (3.16)$$

$$\leq \frac{\mathbb{P}_M(|\phi\rangle | \mu)}{1 - \sigma/\eta}. \quad (3.17)$$

The second line follows as $\mathbb{P}_M(|\phi\rangle | \lambda) \geq (1 - \sigma/\eta)$ for $\lambda \in \Omega$ and the third by expanding the range of integration. \square

Lemma 2.3 establishes how the asymmetric overlap changes under a unitary transformation, this next lemma does the same for the ϵ -asymmetric overlap.

Lemma 3.4. *Let unitary U satisfy $U|0\rangle = |\phi\rangle$ and $\mu' \overset{\gamma}{\rightsquigarrow} \mu$ for some $\gamma \in \Gamma_U$, then*

$$\varpi_\epsilon(|\phi\rangle | \mu) \geq \left(1 - \frac{\epsilon}{\delta}\right) \varpi_\delta(|0\rangle | \mu') \quad (3.18)$$

for any $0 \leq \epsilon \leq \delta < 1$.

Proof. Since ϖ_ϵ is defined as a greatest lower bound, it suffices to prove that for every ϵ -typical $\Omega \in \Sigma$ for $|\phi\rangle$ there exists a δ -typical $\Omega' \in \Sigma$ for $|0\rangle$ such that

$$\mu(\Omega) \geq \left(1 - \frac{\epsilon}{\delta}\right) \mu'(\Omega'). \quad (3.19)$$

For any $\chi \in \Delta_{|0\rangle}$, consider the $\nu \in \Delta_{|\phi\rangle}$ such that $\chi \overset{\gamma}{\rightsquigarrow} \nu$. Then for any ϵ -typical $\Omega \in \Sigma$ for $|\phi\rangle$ it follows that

$$1 - \epsilon < \nu(\Omega) = \int_{\Lambda} d\chi(\lambda) \gamma(\Omega | \lambda). \quad (3.20)$$

Therefore, by Lem. 3.1, $\chi(\bar{k}_\kappa(\gamma(\Omega | \lambda))) > 1 - \frac{\epsilon}{\kappa}$ for any $\kappa \in [\epsilon, 1]$. Recalling that $\chi \in \Delta_{|0\rangle}$ is arbitrary, this means that $\Omega' \stackrel{\text{def}}{=} \bar{k}_\kappa(\gamma(\Omega | \lambda))$ is (ϵ/κ) -typical for $|0\rangle$.

Considering then $\mu(\Omega)$,

$$\mu(\Omega) = \int_{\Lambda} d\mu'(\lambda) \gamma(\Omega | \lambda) \geq \int_{\Omega'} d\mu'(\lambda) \gamma(\Omega | \lambda) \quad (3.21)$$

$$\geq (1 - \kappa)\mu'(\Omega') \quad (3.22)$$

having noted that $\gamma(\Omega | \lambda) \geq (1 - \kappa)$ for $\lambda \in \Omega'$. Letting $\delta \stackrel{\text{def}}{=} \epsilon/\kappa$ completes the proof. \square

This next property generalises Lem. 2.4 by providing a relationship between the overlap with a set of quantum states and the individual overlaps with each state.

Lemma 3.5. *For any finite set $\mathcal{S} \subset \mathcal{P}(\mathcal{H})$ of quantum states, any preparation measure μ , and any $\eta \in (0, 1)$*

$$\sum_{|i\rangle \in \mathcal{S}} \varpi_\eta(|i\rangle | \mu) \geq \varpi_\eta(\mathcal{S} | \mu). \quad (3.23)$$

Proof. For each $|i\rangle \in \mathcal{S}$ let $\Omega_i \in \Sigma$ be an η -typical subset for $|i\rangle$. Consider $\Omega \stackrel{\text{def}}{=} \cup_{|i\rangle \in \mathcal{S}} \Omega_i$ which must also be η -typical for each $|i\rangle \in \mathcal{S}$ as $\nu_i(\Omega) \geq \nu_i(\Omega_i)$. Therefore by definition $\varpi_\eta(\mathcal{S} | \mu) \leq \mu(\Omega)$. By Boole's inequality,

$$\varpi_\eta(\mathcal{S} | \mu) \leq \mu(\Omega) \leq \sum_{|i\rangle} \mu(\Omega_i) \quad (3.24)$$

from which the result follows by recalling that all Ω_i are arbitrary η -typical subsets. \square

Finally, in Sec. 2.2.2 anti-distinguishable sets of quantum states were defined. Lemma 2.6 then showed how anti-distinguishable triples affect asymmetric overlaps. In the following lemma, the corresponding statement is proved for the ϵ -asymmetric overlap and *approximately* anti-distinguishable sets (that is, the ‘‘anti-distinguishing’’ measurement must be accurate to within $\pm\epsilon$).

Lemma 3.6. *Consider a quantum measurement $M = \{E_{-\psi}, E_{-\phi}, E_{-0}\}$ such that for some $\epsilon \in [0, 1)$*

$$\int_{\Lambda} d\mu(\lambda) \mathbb{P}_M(E_{-\psi} | \lambda) \leq \epsilon, \quad \forall \mu \in \Delta_{|\psi\rangle} \quad (3.25)$$

and similarly for $|\phi\rangle$ and $|0\rangle$. Then for any $\mu \in \Delta_{|\psi\rangle}$ and $\kappa \in (2\epsilon, 1)$,

$$\varpi_{\epsilon/\kappa}(|0\rangle, |\phi\rangle | \mu) \geq (1 - \kappa) (\varpi_{2\epsilon/\kappa}(|0\rangle | \mu) + \varpi_{2\epsilon/\kappa}(|\phi\rangle | \mu)) - \epsilon. \quad (3.26)$$

Proof. Since ϖ_ϵ is defined as a greatest lower bound, it suffices to prove that for every $\Omega \in \Sigma$ that is (ϵ/κ) -typical for both $|0\rangle$ and $|\phi\rangle$ there exists $\Omega' \in \Sigma$ that is $(2\epsilon/\kappa)$ -typical for $|\phi\rangle$ and Ω'' that is $(2\epsilon/\kappa)$ -typical for $|0\rangle$ such that

$$\mu(\Omega) \geq (1 - \kappa) (\mu(\Omega') + \mu(\Omega'')) - \epsilon. \quad (3.27)$$

Start by defining the measurable function $g_\phi(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_M(E_{-\psi}|\lambda) + \mathbb{P}_M(E_{-0}|\lambda)$. Define g_0 similarly.

By assumption, $\int_\Lambda d\nu(\lambda) g_\phi(\lambda) \geq 1 - \epsilon$ for all $\nu \in \Delta_{|\phi\rangle}$ so, by Lem. 3.1, $\bar{k}_\kappa(g_\phi)$ is (ϵ/κ) -typical for $|\phi\rangle$. Similarly, $\bar{k}_\kappa(g_0)$ is (ϵ/κ) -typical for $|0\rangle$.

So consider any $\Omega \in \Sigma$ that is (ϵ/κ) -typical for both $|\phi\rangle$ and $|0\rangle$. For every $\nu \in \Delta_{|\phi\rangle}$ then

$$\nu(\Omega \cap \bar{k}_\kappa(g_\phi)) \geq \nu(\Omega) + \nu(\bar{k}_\kappa(g_\phi)) - 1 > 1 - \frac{2\epsilon}{\kappa}. \quad (3.28)$$

Therefore $\Omega' \stackrel{\text{def}}{=} \Omega \cap \bar{k}_\kappa(g_\phi)$ is $(2\epsilon/\kappa)$ -typical for $|\phi\rangle$. By a similar argument $\Omega'' \stackrel{\text{def}}{=} \Omega \cap \bar{k}_\kappa(g_0)$ is $(2\epsilon/\kappa)$ -typical for $|0\rangle$.

Finally, it therefore follows that

$$\mu(\Omega) = \int_\Omega d\mu(\lambda) (g_\phi(\lambda) + g_0(\lambda) - \mathbb{P}_M(E_{-\psi}|\lambda)) \quad (3.29)$$

$$\geq \int_{\Omega'} d\mu(\lambda) g_\phi(\lambda) + \int_{\Omega''} d\mu(\lambda) g_0(\lambda) - \epsilon \quad (3.30)$$

$$\geq (1 - \kappa) (\mu(\Omega') + \mu(\Omega'')) - \epsilon. \quad (3.31)$$

The second line follows from the upper bound on $\int_\Lambda d\mu(\lambda) \mathbb{P}_M(E_{-\psi}|\lambda)$ and by restricting the range of the other integrals. The third line follows by noting that g_ϕ and g_0 are lower-bounded in Ω' and Ω'' respectively. This completes the proof. \square

3.4.3 Error-Tolerant Results

With the technical background of ϵ -asymmetric overlaps, the proof of Thm. 2.1 can be adapted to obtain an error-tolerant variant. Since the proofs of Thm. 2.2 and Thm. 3.1 use Thm. 2.1, this will provide an error-tolerant argument against both epistemic superpositions and ESMR/EMMR macro-realism in quantum theory.

Theorem 3.2. *Consider a $d > 3$ dimensional quantum system described by some ontological model. Assume one can experimentally demonstrate quantum probabilities to within some $\pm\epsilon \in (0, 1]$ as in Eq. (2.4). Let \mathcal{B} be any orthonormal basis of \mathcal{H} and $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ be any pure state such that $|\langle 0|\psi\rangle| = \alpha \in (0, \frac{1}{\sqrt{2}})$ for some $|0\rangle \in \mathcal{B}$. For any choices of $0 < \tau \leq \eta < \frac{1}{16}$ there exist choices of $\epsilon > 0$ such that*

$$\varpi_\eta(\mathcal{B} | \mu) \leq 1 - \tau \quad (3.32)$$

for some $\mu \in \Delta_{|\psi\rangle}$, for a finite range of $\alpha \in (0, \frac{1}{\sqrt{2}})$.

Proof. The proof proceeds by contradiction. To this end, assume that

$$\varpi_\eta(\mathcal{B} | \mu) > 1 - \tau \quad (3.33)$$

for all $\mu \in \Delta_{|\psi\rangle}$ where $0 < \tau \leq \eta < \frac{1}{16}$.

Consider the construction of $|\psi\rangle, |\phi\rangle, |0\rangle, \mathcal{B}' \ni |0\rangle$, and U used in the proof of Thm. 2.1 [Eqs. (2.50, 2.51)]. Let $M_A = \{E_{-\psi}, E_{-\phi}, E_{-0}\}$ be a quantum anti-distinguishing measurement for $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ (i.e. the measurement that would be anti-distinguishing in the absence of error), M be a basis measurement for \mathcal{B} , and let M' be a basis measurement for \mathcal{B}' . Recall that $\langle 0 | \psi \rangle = \alpha \in (0, \frac{1}{\sqrt{2}})$.

The assumption that quantum probabilities are reproduced to within $\pm\epsilon$ variously implies the following.

$$\mathbb{P}_{M_A}(E_{-\psi} | \mu) \leq \epsilon, \quad \forall \mu \in \Delta_{|\psi\rangle} \quad (3.34)$$

$$\mathbb{P}_{M_A}(E_{-\phi} | \nu) \leq \epsilon, \quad \forall \nu \in \Delta_{|\phi\rangle} \quad (3.35)$$

$$\mathbb{P}_{M_A}(E_{-0} | \chi) \leq \epsilon, \quad \forall \chi \in \Delta_{|0\rangle} \quad (3.36)$$

$$\mathbb{P}_{M'}(|0\rangle \vee |1'\rangle | \chi) \geq 1 - \epsilon, \quad \forall \chi \in \Delta_{|0\rangle} \quad (3.37)$$

$$\mathbb{P}_{M'}(|0\rangle \vee |1'\rangle \vee |3'\rangle | \nu) \geq 1 - \epsilon, \quad \forall \nu \in \Delta_{|\phi\rangle} \quad (3.38)$$

$$\mathbb{P}_{M'}(|3'\rangle | \mu) \leq \epsilon, \quad \forall \mu \in \Delta_{|\psi\rangle} \quad (3.39)$$

$$\mathbb{P}_M(|i\rangle | \chi_i) \geq 1 - \epsilon, \quad \forall \chi_i \in \Delta_{|i\rangle}, \forall |i\rangle \in \mathcal{B} \quad (3.40)$$

Now proceed to derive the contradiction. Choosing any $\mu \in \Delta_{|\psi\rangle}$, Lem. 3.6 implies that

$$\varpi_{\epsilon/\kappa}(|0\rangle, |\phi\rangle | \mu) \geq (1 - \kappa) (\varpi_{2\epsilon/\kappa}(|\phi\rangle | \mu) + \varpi_{2\epsilon/\kappa}(|0\rangle | \mu)) - \epsilon \quad (3.41)$$

for any $\kappa \in (2\epsilon, 1)$. Let $\delta \in [2\epsilon/\kappa, 1)$, then Lem. 3.2 gives $\varpi_{2\epsilon/\kappa}(|0\rangle | \mu) \geq \varpi_\delta(|0\rangle | \mu)$. This, together with Lem. 3.4, gives

$$\varpi_{\epsilon/\kappa}(|0\rangle, |\phi\rangle | \mu) \geq (1 - \kappa) \left(\left(1 - \frac{2\epsilon}{\kappa\delta}\right) \varpi_\delta(|0\rangle | \mu') + \varpi_\delta(|0\rangle | \mu) \right) - \epsilon \quad (3.42)$$

where $\mu' \overset{\gamma}{\rightsquigarrow} \mu$ for some choice of $\gamma \in \Gamma_U$. Note that for the first term to contribute non-trivially, it is required that $\delta > 2\epsilon/\kappa$.

The next step is to relate the ϖ_δ terms on the right hand side to their quantum probabilities. For any $\bar{\mu} \in \Delta_{|\psi\rangle}$, Lem. 3.5 and Eq. (3.33) imply

$$\sum_{|i\rangle \in \mathcal{B}} \varpi_\eta(|i\rangle | \bar{\mu}) \geq \varpi_\eta(\mathcal{B} | \bar{\mu}) > 1 - \tau. \quad (3.43)$$

Using Lem. 3.3 and recalling that $|0\rangle \in \mathcal{B}$ this further implies that

$$\varpi_\eta(|0\rangle | \bar{\mu}) > 1 - \tau - \frac{\sum_{|i\rangle \neq |0\rangle} \mathbb{P}_M(|i\rangle | \bar{\mu})}{1 - \epsilon/\eta} \quad (3.44)$$

so long as $\eta > \epsilon$. Lastly, noting that $\alpha^2 - \epsilon \leq \mathbb{P}_M(|0\rangle | \bar{\mu}) = 1 - \sum_{|i\rangle \neq |0\rangle} \mathbb{P}_M(|i\rangle | \bar{\mu})$ gives

$$\varpi_\eta(|0\rangle | \bar{\mu}) > \frac{\alpha^2 - \epsilon - \tau - \frac{\epsilon}{\eta}(1 - \tau)}{1 - \epsilon/\eta} \quad (3.45)$$

for all $\bar{\mu} \in \Delta_{|\psi\rangle}$. In particular, this holds both when $\bar{\mu} = \mu$ and $\bar{\mu} = \mu'$ from Eq. (3.42).

The final step before putting everything together is to relate the bipartite (ϵ/κ) -asymmetric overlap (left hand side of Eq. (3.42)) to a corresponding quantum measurement probability. Let $f(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_{M'}(|0\rangle | \lambda) + \mathbb{P}_{M'}(|1'\rangle | \lambda)$ and $g(\lambda) \stackrel{\text{def}}{=} \mathbb{P}_{M'}(|3'\rangle | \lambda)$. From the above measurement probabilities, Lem. 3.1 gives that

$$\chi(\bar{k}_\kappa(f)) > 1 - \frac{\epsilon}{\kappa}, \quad \forall \chi \in \Delta_{|0\rangle} \quad (3.46)$$

$$\nu(\bar{k}_\kappa(f + g)) > 1 - \frac{\epsilon}{\kappa}, \quad \forall \nu \in \Delta_{|\phi\rangle} \quad (3.47)$$

since it is already required that $\kappa > 2\epsilon > \epsilon$. Because $\bar{k}_\kappa(f) \subseteq \bar{k}_\kappa(f + g)$ it follows that $\Omega \stackrel{\text{def}}{=} \bar{k}_\kappa(f + g)$ is a measurable subset of Λ which is (ϵ/κ) -typical for both $|\phi\rangle$ and $|0\rangle$. Consider then

$$\mathbb{P}_{M'}(|0\rangle \vee |1'\rangle | \mu) = \int_\Lambda d\mu(\lambda) f(\lambda). \quad (3.48)$$

By Eq. (3.39) it follows that $\int_\Lambda d\mu(\lambda) g(\lambda) \leq \epsilon$ so

$$\mathbb{P}_{M'}(|0\rangle \vee |1'\rangle | \mu) \geq \int_\Lambda d\mu(\lambda) (f(\lambda) + g(\lambda)) - \epsilon \quad (3.49)$$

$$\geq \int_\Omega d\mu(\lambda) (f(\lambda) + g(\lambda)) - \epsilon \quad (3.50)$$

$$\geq (1 - \sigma)\mu(\Omega) - \epsilon \quad (3.51)$$

$$\geq (1 - \sigma)\varpi_{\epsilon/\kappa}(|\phi\rangle, |0\rangle | \mu) - \epsilon. \quad (3.52)$$

Where the final line follows as Ω is (ϵ/κ) -typical for both $|\phi\rangle$ and $|0\rangle$. Using that quantum probabilities are reproduced to within $\pm\epsilon$ and noting the construction of Eq. (2.50) gives

$$\alpha^2(1 + 2\alpha^2) + 2\epsilon \geq (1 - \kappa)\varpi_{\epsilon/\kappa}(|0\rangle, |\phi\rangle | \mu). \quad (3.53)$$

Combining Eqs. (3.41, 3.45, 3.53) shows that the assumption-towards-contradiction Eq. (3.33) implies

$$\alpha^2(1 + 2\alpha^2) + \epsilon(3 - \kappa) > \frac{2(1 - \kappa)^2(1 - \frac{\epsilon}{\kappa\eta})}{1 - \epsilon/\eta} \left(\alpha^2 - \epsilon - \tau - \frac{\epsilon}{\eta}(1 - \tau) \right) \quad (3.54)$$

where $\delta = \eta$ has been taken. Recall that the derivation of this inequality required that $\eta > 2\epsilon/\kappa$ and $\eta > \epsilon$, and that $\kappa > 2\epsilon$ is chosen.

To obtain a contradiction and complete the proof, it must be demonstrated that the inequality of Eq. (3.54) is violated for appropriate choices of parameters. This is much easier if one first simplifies Eq. (3.54).

To simplify, first assume that $\eta \geq \tau$, so that Eq. (3.33) implies $\varpi_\eta(\mathcal{B} | \mu) > 1 - \eta$. Therefore, Eq. (3.54) holds with $\tau = \eta$. Another simplifying assumption comes by taking $\kappa = \sqrt{\epsilon}$. Under these assumptions, Eq. (3.54) becomes

$$\alpha^2(1 + 2\alpha^2) + \kappa^2(3 - \kappa) > \frac{2(1 - \kappa)^2(\eta - \kappa)}{\eta - \kappa^2} \left(\alpha^2 - \eta - \frac{\kappa^2}{\eta} \right) \quad (3.55)$$

where $\eta > 2\kappa$, $\eta > \kappa^2$ and $\kappa < \frac{1}{2}$.

Observe that for any fixed value of η , a value of $\epsilon = \kappa^2$ may be chosen to be arbitrarily small. In the limit of $\epsilon \rightarrow 0$ the inequality becomes

$$\alpha^2(2\alpha^2 - 1) + 2\eta > 0, \quad (3.56)$$

which is violated by a finite range of $\alpha^2 \in (0, \frac{1}{2})$ for every value of $\eta \in (0, \frac{1}{16})$.

Consider turning this logic around. For any value of $\eta \in (0, \frac{1}{16})$, if one assumes Eq. (3.33) with $\eta \geq \tau$ then there exists a value of $\epsilon > 0$ such that one can reach a contradiction by violating Eq. (3.54). This completes the proof. \square

As stated, Thm. 3.2 is somewhat formal and opaque. A little unpacking is required to see how it provides error-tolerant statements against both epistemic superpositions and ESMR/EMMR macro-realism.

First, consider Bob who believes that some superposition $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ should be epistemic with respect to some orthonormal basis $\mathcal{B} \not\ni |\psi\rangle$ of a $d > 3$ dimensional system. As discussed in Sec. 3.4.1, Bob should be compelled to make a statement along the lines of: “any preparation of $|\psi\rangle$ will produce an ontic state that is η -typical for at least one state in \mathcal{B} with probability at least $1 - \tau$ ”. The values of η and τ are left up to Bob and quantify exactly how macro-realist his belief is (relative to the noise in available experiments), but they should be small if $|\psi\rangle$ is to be approximately epistemic. In summary, Bob believes that

$$\varpi_\eta(\mathcal{B} | \mu) > 1 - \tau, \quad \forall \mu \in \Delta_{|\psi\rangle}. \quad (3.57)$$

Alice, who has access to a lab, can contradict Bob’s belief in the following way. Suppose that $0 < \tau \leq \eta < \frac{1}{16}$ and Alice can demonstrate that quantum probabilities are correct to within $\pm\epsilon \in (0, 1)$. By Thm. 3.2 then for sufficiently small ϵ there is some continuous region $\mathcal{R}_\epsilon \subset (0, 1)$ such that Bob is contradicted if $|\langle 0 | \psi \rangle|^2 \in \mathcal{R}_\epsilon$ for any $|0\rangle \in \mathcal{B}$.

In fact, given specific values for η , τ , and α , Alice can probably do better. Looking at the proof of Thm. 3.2, Alice simply needs to find an appropriate ϵ such that Eq. (3.54) can be violated (with valid choices for the other parameters in the inequality) and then demonstrate that quantum probabilities hold to within $\pm\epsilon$. Note that, in this more general case, Alice doesn't necessarily have to assume that $\tau \leq \eta$.

So why bother with the first, less general, statement if violation of Eq. (3.54) is the more powerful result? Simply put, it is not clear *a priori* that Eq. (3.54) can be violated for reasonable choices of the parameters. The first statement serves to demonstrate that the contradiction is possible, while the second is strictly more general.

Theorem 3.2 can be used in a similar way to rule out ESMR or EMMR ontologies in an error-tolerant way.

Consider Clare, who believes in ESMR or EMMR macro-realism for some quantity Q of a quantum system with $n > 3$ distinguishable values. Similarly to Bob—and as discussed in Sec. 3.4.1—Clare should be compelled to say “any preparation of any $|\psi\rangle \in \mathcal{P}(\mathcal{H})$ will produce an ontic state that is η -typical for at least one operational eigenstate of Q with probability at least $1 - \tau$.” As with Bob, the exact values of η and τ are up to Clare, but should be small if Clare seriously believes in ESMR/EMMR macro-realism.

As noted in Sec. 3.2.3, the operational eigenstates of Q in a quantum system form an orthonormal basis \mathcal{B}_Q . This means that Alice believes that

$$\varpi_\eta(\mathcal{B}_Q | \mu) > 1 - \tau, \quad \forall \mu \in \Delta_{|\psi\rangle}, \quad \forall |\psi\rangle \in \mathcal{P}(\mathcal{H}). \quad (3.58)$$

It is clear that Alice can contradict Clare's belief in a similar way to how she contradicted Bob. Specifically, if $0 < \tau \leq \eta < \frac{1}{16}$ then there is an $\epsilon \in (0, 1)$ such that if Alice can demonstrate that quantum probabilities hold to within $\pm\epsilon$ then Clare is proved wrong by Thm. 3.2. More generally given τ and η , if Alice can violate Eq. (3.54) for any valid values of α , κ , and ϵ , then demonstrating that quantum probabilities hold to within $\pm\epsilon$ will contradict Clare.

Note that, since Clare must believe Eq. (3.58) for all pure states $|\psi\rangle \in \mathcal{P}(\mathcal{H})$, it is easier for Alice to contradict Clare than Bob, whose choice of $|\psi\rangle$ is arbitrary.

In this way, Thm. 3.2 can be used as the basis for error-tolerant variations on both Thm. 2.2 and Thm. 3.1. This has been achieved by generalising the asymmetric overlap to the ϵ -asymmetric overlap in a natural way by using ϵ -typical subsets. The proof (while rather involved) is still arguably much simpler than that of Thm. 2.5, which established an error-tolerant version of Thm. 2.4 by using the symmetric overlap. This is because the asymmetric overlap has much more in common with the ϵ -asymmetric overlap than the symmetric overlap, so the proof strategy can more more easily adapted. The main disadvantage of

working with the ϵ -asymmetric overlap is a proliferation of free parameters, so rather than getting a clean contradiction, one finds an inequality Eq. (3.54) with five parameters which must be violated to obtain a contradiction. The proof of Thm. 3.2 demonstrates that this is possible at least in some regime—specifically, when $0 < \tau \leq \eta < \frac{1}{16}$ for some ranges of ϵ and α —but violations of Eq. (3.54) may also be possible outside of that regime.

3.5 Summary and Discussion

This chapter addressed the incompatibility of macro-realism and quantum theory by using ontological models and building on the work and techniques of Chap. 2. After a quick summary of this chapter, the results and outstanding questions from these two chapters will be discussed below.

The exact meaning of “macro-realism” has always been somewhat controversial. It was therefore appropriate to start by discussing what macro-realism means and identifying an appropriate specific definition in Sec. 3.1. The history of macro-realism is inseparable from the Leggett-Garg inequalities and their use in the Leggett-Garg argument for the incompatibility of quantum theory and macro-realism. This argument was outlined briefly in Sec. 3.2.1, with a focus on the assumptions required to derive the inequalities.

The definition of macro-realism used in Sec. 3.1.2 is from Ref. [MT17], where the history of “macro-realism” is analysed and the definition clarified using ontological models. This ontological model analysis led to the identification of three sub-categories of macro-realism presented in Sec. 3.2.2: EMMR, ESMR, and SSMR. By applying these definitions to quantum theory, one finds fundamental loopholes in the Leggett-Garg argument. In particular, Leggett-Garg is only able to prove incompatibility of EMMR models with quantum theory, leaving loopholes for ESMR and SSMR as discussed in Sec. 3.2.4. Moreover, the existence of counter-examples in the form of Bohmian mechanics and the Kochen-Specker model show that no theorem can rule out all SSMR models or all ESMR models for $d = 2$ dimensions respectively.

Papers on the Leggett-Garg argument, including those addressing the clumsiness loophole [WM12, KKY⁺16, HM16], have concentrated on $d = 2$ dimensional systems. As a result of closely following the Leggett-Garg assumptions, they are still unable to rule out any models outside of EMMR.

With this background, Sec. 3.3 proceeded to apply the methods of Chap. 2 to macro-realism, obtaining Thm. 3.1 as the main result of this chapter. This proved the incompatibility of quantum theory with all ESMR and EMMR models for observables with

$n > 3$ distinguishable values—a result which is therefore stronger than the Leggett-Garg argument.

Finally, Sec. 3.4 approached the problem of finding error-tolerant variations of Thms. 2.2, 3.1 with a view towards experimental investigations. Any such variation requires a generalisation of the mathematical formalisations of epistemic superpositions and ESMR/EMMR macro-realism. This was done by introducing the ϵ -asymmetric overlap: a generalisation of the asymmetric overlap that is error-tolerant. This enabled a proof of Thm. 3.2, which can be used to obtain error-tolerant variations of both Thm. 2.2 and Thm. 3.1.

Between Thm. 3.1 and the Leggett-Garg argument, the only possibilities for macro-realism compatible with quantum theory that remain are SSMR models (such as Bohmian mechanics) or ESMR models for $d = 2, 3$ dimensions (such as the Kochen-Specker model). One clear question for further investigation is whether it is possible to rule out any subsets of these remaining models. For example, Bohmian mechanics is a ψ -ontic theory and it may therefore be possible to prove the incompatibility of quantum theory and all SSMR ontologies that aren't ψ -ontic. This, together with the result presented here, would essentially say that to be macro-realist you must have an ontology consisting of the full quantum state plus extra information. Many would consider this a very strong argument against macro-realism.

Theorems 2.2–2.4 paint a similarly grim outlook for the epistemic realist, especially one who seeks to explain any superposition-based phenomena using ontological uncertainty. In particular, Thm. 2.2 proves that, for $d > 3$, almost all superpositions defined with respect to any given basis \mathcal{B} must be real. Therefore, any epistemic realist account of quantum theory must include ontic features corresponding to superposition states and the unfortunate cat cannot be put out of its misery. Theorems 2.3, 2.4 then proceed to provide new bound on ontic overlaps for largely arbitrary pairs of quantum states. In particular from Thm. 2.3 one can conclude that almost no quantum states can be ψ -epistemic. From Thm. 2.4, one can conclude that in any moderately large system a large number of pairs of non-orthogonal states cannot overlap significantly, making it unlikely that such overlaps can satisfactorily explain quantum features. These are both much stronger statements than previous results in the same vein have been able to achieve, circumventing the shortcomings noted in Sec. 3.2.4. However, it should be noted that Thms. 2.3, 2.4 do depend on an extra assumption, though one that is mild.

The next stage for all of these results will be to develop experimental tests, which require detailed error-tolerant analyses. Theorem 2.5 begins to develop such an analysis based on Thm. 2.4. An experiment based on this would require demonstration of small errors in probabilities for a wide range of measurements on a $d > 5$ dimensional system.

The error-tolerant variations on Thms. 2.2, 3.1 provided by Thm. 3.2 also open the door for experimental investigation. While this variation uses the ϵ -asymmetric overlap, there are potentially other routes to error-tolerant variations, including one currently in development [HM17]. This demonstrates how Thms. 2.2, 3.1 are just the first steps towards new ways of investigating epistemic superpositions and macro-realism respectively.

It is interesting to note that experiments based on this result will be an entirely new avenue for tests of macro-realism. Experimental tests based on the Leggett-Garg argument will always have certain features and difficulties in common (such as the clumsiness loophole briefly mentioned in section 3.2.4). However, since the approach of this chapter is so different in character one can expect the resulting experiments to be similarly different, hopefully avoiding many of the difficulties common to Leggett-Garg while requiring challenging new high-precision tests of quantum theory in $d > 3$ Hilbert spaces.

One should note that in this chapter the “macro” quantity Q was taken to correspond to a measurement in basis \mathcal{B}_Q in the quantum case. A more general approach might allow Q to correspond to a POVM measurement instead. That is, for each value q of Q there would be some POVM element E_q and the operational eigenstates $|\psi\rangle$ of q would be those satisfying $\langle\psi|E_q|\psi\rangle = 1$. It does not seem unlikely that the results presented here could be fairly directly extended to such a case and this would be another interesting avenue for further work. Such an extension would likely add significant complexity to the proofs without changing the fundamental ideas, however.

Extensions of Thm. 3.1 along these lines would likely start to further separate out the work on macro-realism from that on epistemic superpositions. As noted in Sec. 3.3.2, superpositions are pure states, so extensions looking for new results there will likely stay within the realm of pure-state quantum theory. This is in contrast to macro-realism, which is defined independently from quantum theory, so extensions in that direction may well lead to mixed quantum theory and POVM measurements.

The methods used for the foundational results of Chap. 2 were also adapted to prove communication bounds in quantum information in Sec. 2.5. The result was Thm. 2.6 which provides a bound on the ability for classical resources to perfectly simulate quantum channels which asymptotically matches the best known results from the literature [Mon16a, Mon11, BCT99] while using a substantially simpler argument. In particular, it was proved that at least $2^{n_q + \mathcal{O}(1)} - 1$ bits of classical communication is required for the simulation of a noiseless n_q -qubit quantum channel, even when using arbitrary shared random data. Beyond its simplicity, the proof of Thm. 2.6 has two key advantages over other proofs for similar bounds. First, it identifies an underlying reason for the result: it is a consequence

of the existence of exponentially large sets of quantum states that have small mutual Born-rule overlaps. Second, the proof identifies a general strategy for deriving such bounds. If one found different sets of fingerprinting states and/or classical error-correction codes one could simply re-run the proof to potentially produce more powerful results [Sec. 2.5.3].

It is worth noting that Thm. 2.6 was developed independently from the exponential bound for the same problem presented in Ref. [Mon16a] and preceded the publication of that bound. However, as it stands the bound of Ref. [Mon16a] has two advantages over Thm. 2.6. First, it applies to approximate simulations as well as perfect simulations of quantum channels. Second, it also bounds simulations using two-way classical communication. The problem of extending the proof of Thm. 2.6 to apply to approximate simulations is a topic of ongoing research which is currently unready for publication but remains very promising nonetheless. Simulations using two-way communication are less related to ontological models and therefore of less relevance to this thesis, but it is an interesting open problem as to whether the techniques of Thm. 2.6 could also be applied to this more general setting.

As a final note on Thm. 2.6, one can understand the simplicity and power of the proof as coming from the combination of two facts. First, the existence of exponentially large sets of quantum states with bounded Born-rule overlap discussed above. Second, the property of anti-distinguishability and its implications for ontological models and simulations [Sec. 2.2.2]. The result is the existence of exponentially large sets of quantum states such that all triple subsets are anti-distinguishable. It seems likely that this fact can be fruitfully used to prove results in many areas of quantum information, potentially including quantum cryptography and quantum algorithms.

Finally, it has been assumed throughout Chaps. 2, 3 that all quantum systems are finite-dimensional for clarity. If such extensions are needed, it should not be conceptually difficult to extend any of these results to the infinite case. In particular, Thm. 2.4 strongly suggests that in the infinite case one would find zero overlap between many pairs of quantum states, which would be a very powerful restriction on many types of ontological models.

4

Quantum Conditional Independence and Quantum Common Causes

4.1 The Need for Quantum Conditional Independence

The ontology of quantum theory is normally discussed in terms of quantum states, as in Chaps. 2, 3. However, scientists also typically consider causal influences to be real. If quantum theory challenges our notions of ontology of the states of systems, then it is natural to ask how it might also challenge our causal notions. The simplest scenario where this becomes apparent is in the case of complete common causes and in the related concept of conditional independence. This chapter analyses these fully in a quantum universe. Much of the work presented here has been published in Ref. [ABH⁺17].

4.1.1 The Reality of Causal Influences

It is quite typical to think of causal influences as being “real” (without wishing to go into specifics as to what that might mean). Indeed, it is common to hear that the business of science is primarily to discover what these real causal influences are [Pea09, SGS01]. In particular, the whole science of causal discovery algorithms is predicated on the idea that there exist objective causes that can be discovered. If even the possibility of ontological causal influences is taken seriously in a quantum universe, then there needs to be a framework in which to discuss quantum causal influences.

It is important, especially for this chapter and the next, to distinguish causal notions from probabilistic or statistical ones. Causation is about influences between events, while probability and statistics describe if and how often events (or combinations of events) occur.

Reichenbach's principle [Rei56] might be viewed as a cornerstone of causal reasoning. It is often paraphrased as “no correlation without causation”. Thus providing a link between the probabilistic/statistical idea of “correlation” and that of causation. Reichenbach's principle is so well ingrained in both everyday and scientific thought that it is rarely specifically discussed outside of philosophy.

To illustrate, if the dog barks every time the post arrives people naturally assume that it is the arrival of the post (the sound it makes, the smell of the deliverer, *etc.*) that causes the dog to bark, rather than a mere coincidence. Similarly, if every time the cat runs away the dog is scolded it is likely that both events are caused by the dog barking rather than either a coincidence or one event being a cause for the other. More seriously, without some general commitment to Reichenbach's principle (or something sufficiently similar) it would be impossible to regard two experiments as independent due to a lack of causal mechanism between them. This would be catastrophic for science in general.

So Reichenbach's principle calls for causal links between any correlated events A and B . What are these links? They must be at least one of the following: A is a cause for B , B is a cause for A , or there is some common cause for both A and B . The last of these is the most interesting. Reichenbach's principle requires a quantitative restriction on the probabilities if the link is solely via a common cause: conditional independence. A full statement and discussion of Reichenbach's principle is deferred until Sec. 4.2.

4.1.2 Causation in Bell Experiments

Despite the central role of causal explanations in science and life, certain quantum experiments seem to elude such explanations. The best known of these are Bell experiments—*viz.* experiments of the type seen in Bell's theorem [Bel66]. In such experiments, Alice prepares a pair of systems that are distributed to Bob and Clare who perform independent measurements on those systems. These experiments are purposefully constructed to suggest the only reasonable causal link between the measurement outcomes of Bob and Clare is via a common cause (presumably involving Alice). Bell's theorem proceeds to show that certain quantum experiments of this form can produce statistics which violate Reichenbach's principle for any conceivable common cause. The findings of Bell's theorem have been thoroughly experimentally verified [HBD⁺15, SMSC⁺15, GVW⁺15].

Bell’s theorem is most often discussed in terms of non-locality, however recent interpretations have put its causal implications front-and-centre [WS15, CL14, HLP14, PB15]. In particular, Ref. [WS15] considers the statistics observed in Bell experiments in terms of classical causal models. It is shown that no classical causal model—even those with exotic properties such as retrocausality and superluminal signalling—can account for these statistics without undesirable fine tuning. Fine tuning is typically considered enough to rule out a proposed causal explanation [Pea09]. Note that, at least from a causal perspective, this is strictly stronger than Bell’s theorem. Bell’s theorem assumes that all causal links except common causes are ruled out by the experimental structure, while in Ref. [WS15] all causal structures are considered.

The conclusion must be that quantum theory demands a revision of our ideas of causality—either to allow and explain fine tuning or to revise what is permitted as a common cause. At very least, Bell experiments demand a causal explanation. It has been suggested [CL14, WS15] that an appropriate revision could preserve a form of Reichenbach’s principle while rejecting fine tuning.

In this chapter, such a revised quantum Reichenbach’s principle is motivated and presented. Specifically, this chapter mostly deals with the case where quantum systems B and C are in the causal future of system A . The aim is to characterise the quantum channels from A to BC that can occur when A is a complete common cause for B and C . Since Reichenbach’s principle uses conditional independence to characterise common causes [Sec. 4.2], this is done by generalising to a natural quantum conditional independence. Because of this limited scope, the quantum conditional independence thus defined only needs to make sense when A is in the causal past of BC . Compare this with classical conditional independence, which can hold between random variables with any causal relationships. Quantum conditional independence can be defined in several equivalent ways, each of which naturally generalises a corresponding expression for classical conditional independence. Moreover, both quantum conditional independence and quantum Reichenbach presented here reduce to the classical cases in appropriate limits. These properties strongly suggest that the definitions here are the correct way of generalising these classical concepts to natural quantum analogues.

The equivalence of four ways to define quantum conditional independence will be proved in Sec. 4.3.4. A generalisation to channels from A to $k > 2$ outputs will then be given in Sec. 4.3.7.

Of course, there is much more to the classical Reichenbach’s principle and classical conditional independence than channels from A to BC . Indeed, there is a whole framework of classical causal models that generalises these concepts. The problem of making

the corresponding generalisations for quantum Reichenbach’s principle and quantum conditional independence will be tackled in Chap. 5.

For clarity of presentation in this chapter and the next, all random variables and graphs will be taken to be finite and all quantum systems finite-dimensional. It is not anticipated that any conceptual changes would be needed to extend the results to the infinite cases.

4.1.3 Metaphysically-Neutral Quantum Conditional Independence

Naturally, the subject matter of this chapter and the next involves some discussion of metaphysical concepts. However, the results themselves can be viewed somewhat independently from the metaphysics. Indeed, an aim in writing these chapters is to make them as metaphysically-neutral as the results allow.

As in the rest of this thesis, it is suspected that the results presented here can inform discussions of causation and ontology in metaphysics, but this thesis does not pretend to contain those discussions. In particular, no metaphysical positions need to be subscribed to in order to derive and understand the results presented here.

For example, the above discussion used the assumption that causal influences are fundamentally real to motivate the work in this chapter. However, it is likely that a philosopher who argues against this position would still find value in the definition of quantum conditional independence and quantum common cause derived here. Even for that philosopher, Bell’s theorem provides strong reasons for quantum common causes.

In short, the metaphysical content of this chapter and the next (while necessary) is always secondary to the physical and mathematical content.

4.2 Classical Conditional Independence and Reichenbach’s Principle

4.2.1 Reichenbach’s Principle in Two Parts

Reichenbach’s principle [Rei56] can be usefully thought of as comprising two parts (as identified in Ref. [CL14]). The first, the *qualitative part*, is the claim “no correlation without causation” most usually associated with Reichenbach. This is supplemented by the *quantitative part* which characterises common causes and provides the link to the probabilistic notion of conditional independence.

In full, the qualitative part of Reichenbach’s principle is: if two physical variables Y and Z are statistically dependent, then there should be a causal explanation of this fact such that at least one of the following is true:

- (a) Y is a cause of Z ;
- (b) Z is a cause of Y ; or
- (c) there is a third variable X which is a common cause of both Y and Z .

These causal influences may be indirect (mediated by other variables) and each of these cases may be true many times over (multiple chains of causation or multiple common causes). If none of these causal influences exist, then Y and Z are *ancestrally independent*: they have no causal ancestor in common. In this way, an alternative phrasing for the qualitative part is: ancestral independence of Y and Z implies statistical independence $\mathbb{P}(Y, Z) = \mathbb{P}(Y)\mathbb{P}(Z)$.

The quantitative part of Reichenbach’s principle applies to the case where X and Y share a common cause but neither is a direct cause of the other (that is, only case (c) from the qualitative part is true). X is called the *complete common cause* for Y and Z when X is the union of all common causes and there is no cause (direct or indirect) from Y to Z or vice versa. The quantitative part states that in such a case Y and Z are conditionally independent given their complete common cause X , that is

$$\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X). \quad (4.1)$$

Note that Reichenbach’s principle therefore forms a two-way link between causal and statistical notions. First, if a statistical condition is met (statistical dependence), then at least one causal condition must be met (the qualitative part). Second, in the case of a particular causal condition being met (simple common cause), a further statistical condition must be met for consistency (conditional independence).

4.2.2 Justifying the Quantitative Part

Unsurprisingly, justification for Reichenbach’s principle is a delicate problem in philosophy. Basic questions from across the metaphysics of causation and probability play into this issue. Nonetheless, it is useful to present one way to justify the quantitative part given the qualitative part, as this can then be used to motivate its quantum counterpart. Really, this “justification” is just for illustrative and motivational purposes. It will involve temporarily making a quite strong metaphysical assumption. However, Reichenbach’s principle (and its quantum generalisation in Sec. 4.3) stand apart from such assumptions.




Figure 4.1: A causal structure represented as a directed acyclic graph depicting that X is the complete common cause of Y and Z .

For the sake of definiteness, suppose the position of a determinist is taken. So uncertainty is understood as arising from ignorance about deterministic dynamics or from initial conditions. The task is to use this assumption and the qualitative part of Reichenbach's principle to show that if X is the complete common cause for Y and Z then $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$.

A classical channel describing the influence of random variable X on Y is given by a probability distribution $\mathbb{P}(Y|X)$. The determinist will always view this as the result of a deterministic *function* $f : X \times \Lambda \rightarrow Y$ for some unknown variable Λ . Any such channel can always be dilated in this way.

Definition 4.1. A classical deterministic dilation of a classical channel $\mathbb{P}(Y|X)$ is given by some function $f : X \times \Lambda \rightarrow Y$ for random variable Λ taking values λ with probability distribution $\mathbb{P}(\Lambda)$ such that

$$\mathbb{P}(Y|X = x) = \sum_{\lambda} \mathbb{P}(\lambda)\delta(Y, f(x, \lambda)) \quad (4.2)$$

where $\delta(Y, y') = 1$ for $Y = y'$ and 0 otherwise.

Consider the case where X is a complete common cause for Y and Z , illustrated in Fig. 4.1. Let $f = (f_Y, f_Z)$ be a classical dilation of $\mathbb{P}(Y, Z|X)$ for functions $f_Y : \Lambda \times X \rightarrow Y$ and $f_Z : X \times \Lambda \rightarrow Z$. Since X is a *complete* common cause, Λ must split into a pair of variables $\Lambda = \Lambda_Y \times \Lambda_Z$ such that Λ_Y only influences Y and Λ_Z only influences Z . If Λ were not of this form then it would be a new common cause for Y and Z which is not screened through X , violating the assumption that X is a complete common cause. Therefore, the dilation takes the form of $f_Y : \Lambda_Y \times X \rightarrow Y$ and $f_Z : X \times \Lambda_Z \rightarrow Z$ such that

$$\mathbb{P}(Y, Z|X = x) = \sum_{\lambda_Y, \lambda_Z} \mathbb{P}(\lambda_Y, \lambda_Z)\delta(Y, f_Y(\lambda_Y, x))\delta(Z, f_Z(x, \lambda_Z)) \quad (4.3)$$

as illustrated in Fig. 4.2.




Figure 4.2: The causal structure of Fig. 4.1, expanded so that Y and Z each has a latent variable as a causal parent in addition to X so that both Y and Z can be made to depend functionally on their parents.

Using the qualitative part of Reichenbach's principle, ancestral independence of Λ_Y and Λ_Z implies their statistical independence $\mathbb{P}(\Lambda_Y, \Lambda_Z) = \mathbb{P}(\Lambda_Y)\mathbb{P}(\Lambda_Z)$. Equation (4.3) then immediately implies $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$.

A well-known converse statement is also worth noting: any classical channel $\mathbb{P}(Y, Z|X)$ satisfying $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$ admits a deterministic dilation where X is the complete common cause of Y and Z [Pea09].

This completes the illustrative justification for the quantitative part of Reichenbach's principle. Note that the mathematics of classical deterministic dilations is independent of the metaphysical assumptions. In fact, just by acknowledging the logical possibility of a deterministic dilation of $\mathbb{P}(Y, Z|X)$ one can say the following.

Definition 4.2. A classical channel $\mathbb{P}(Y, Z|X)$ is compatible with X being the deterministic common cause for Y and Z if and only if there is a deterministic dilation of the channel of the form

$$\mathbb{P}(Y, Z|X = x) = \sum_{\lambda_Y, \lambda_Z} \mathbb{P}(\lambda_Y)\mathbb{P}(\lambda_Z)\delta(Y, f_Y(\lambda_Y, x))\delta(Z, f_Z(x, \lambda_Z)). \quad (4.4)$$

With this definition, the following can immediately be stated.

Theorem 4.1. Given a classical channel $\mathbb{P}(Y, Z|X)$, the following are equivalent:

1. $\mathbb{P}(Y, Z|X)$ is compatible with X being the deterministic common cause for Y and Z .
2. $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$.

That (1) \Rightarrow (2) follows immediately from Eq. (4.4). That (2) \Rightarrow (1) is proved in Ref. [Pea09], as noted above.

The reverse implication, (2) \Rightarrow (1), shows that common cause is a possible causal explanation of conditional independence. It is important to note that common cause is only one *possible* causal explanation, not the *only* causal explanation. However, in the case where Y and Z are in the causal future of X the other causal explanations will involve

fine tuning, suggesting that a complete common cause explanation is probably the best explanation.

Theorem 4.1 can be read in two ways: as a summary of the justification presented above, or simply as a pair of equivalent definitions for conditional independence. The reading depends on how condition (1) is understood. If condition (1) is taken as a causal statement, then Thm. 4.1 summarises the above justification. However, if condition (1) is taken as simply a formal probabilistic statement it can be read as a definition of conditional independence.

4.3 Quantum Conditional Independence and Quantum Reichenbach

Having thoroughly introduced Reichenbach's principle classically, these ideas can be used to motivate a natural quantum generalisation. Before proceeding to the results some useful notation for quantum channels should be introduced.

4.3.1 The Choi-Jamiołkowski Isomorphism

A quantum channel from system A to B is normally given as a completely positive trace preserving (CPTP) map from density operators on A to density operators on B , written as $\mathcal{E}_{B|A} : \mathcal{D}(\mathcal{H}_A) \rightarrow \mathcal{D}(\mathcal{H}_B)$. Alternatively, the same channel can equivalently be expressed as a density operator in $\mathcal{D}(\mathcal{H}_B \otimes \mathcal{H}_A)$ using the Choi-Jamiołkowski isomorphism [Jam72, Cho75]. There are several variations on the exact mathematical conventions of this isomorphism [LS13, Lei11], though morally they're more-or-less interchangeable. The notation convention used in this chapter and the next is

$$\rho_{B|A} \stackrel{\text{def}}{=} \sum_{i,j} \mathcal{E}_{B|A} (|i\rangle_A \langle j|) \otimes |i\rangle_{A^*} \langle j| \quad (4.5)$$

where $\{|i\rangle_A\}_i$ is some orthonormal basis for \mathcal{H}_A and $\{|i\rangle_{A^*}\}_i$ is the dual basis in the dual space \mathcal{H}_{A^*} . Strictly, $\rho_{B|A}$ is therefore an operator on $\mathcal{H}_B \otimes \mathcal{H}_{A^*}$ in this particular version of the isomorphism. This choice of definition has two key advantages over the alternatives: $\rho_{B|A}$ is a positive operator and the same regardless of the choice of basis used to define it. Note also that it is normalised such that $\text{Tr}_B \rho_{B|A} = \mathbb{1}_{A^*}$ and $\text{Tr} \rho_{B|A} = d_A$.

As $\rho_{B|A}$ is an equivalent way to express the channel $\mathcal{E}_{B|A}$, it follows that there must be a way to express $\rho_B = \mathcal{E}_{B|A}(\rho_A)$ in terms of $\rho_{B|A}$ and ρ_A . This is done by defining the *linking operator*

$$\tau_A^{\text{id}} \stackrel{\text{def}}{=} \sum_{i,j} |i\rangle_{A^*} \langle j| \otimes |i\rangle_A \langle j| \quad (4.6)$$

on $\mathcal{H}_{A^*} \otimes \mathcal{H}_A$ where, again, $\{|i\rangle_A\}_i$ is an arbitrary orthonormal basis and $\{|i\rangle_{A^*}\}_i$ is its dual. This will also be an important operator for general quantum causal models [Sec. 5.3.3]. Using this, it is easy to verify that

$$\rho_B = \mathcal{E}_{B|A}(\rho_A) = \text{Tr}_{A^*A}(\rho_{B|A} \tau_A^{\text{id}} \rho_A). \quad (4.7)$$

Note the similarities between $\rho_{B|A}$ for a quantum channel and $\mathbb{P}(Y|X)$ for a classical channel [LS13]. In particular, note that in a classical decohering limit where the quantum systems and channels are all diagonal in some particular choices of bases, then the main diagonal of $\rho_{B|A}$ is exactly $\mathbb{P}(Y|X)$ and all expressions reduce to their classical counterparts.

As in the rest of the thesis, a density operator written with missing system subscripts indicates the partial trace over those systems. So, for example, given a channel from AB to CD , $\rho_{CD|} \stackrel{\text{def}}{=} \text{Tr}_{AB} \rho_{CD|AB}$ and, by definition, $\rho_{\cdot|AB} = \mathbb{1}_{AB}$. It will also often be convenient to renormalise a channel $\rho_{B|A}$ to unity. This is written with a caret, $\hat{\rho}_{B|A} \stackrel{\text{def}}{=} \rho_{B|A}/d_A$.

4.3.2 Justifying Quantum Reichenbach

This subsection follows the classical example of Sec. 4.2.2 to develop a natural quantum version of Reichenbach’s principle via a definition for quantum conditional independence.

Following the example of Ref. [CL14], the qualitative part of Reichenbach’s principle can be applied to quantum theory almost unchanged: if quantum systems B and C are correlated then this implies a causal connection of at least one of the three forms listed in Sec. 4.2.1. The only difference is to clarify that “correlated” here means that there exist independent bipartite measurements on B and C that have correlated statistics. Alternatively, this can be phrased as: ancestral independence of B and C implies no correlated local statistics.

The challenge, as noted in Sec. 4.1.2, is to identify quantum conditional independence in the $A \rightarrow BC$ scenario and use it to define a quantum quantitative part for Reichenbach’s principle. That is, if A is a complete common cause for B and C , as illustrated in Fig. 4.3, what restrictions should hold for the channel $\rho_{BC|A}$, corresponding to $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$?

It may be tempting to first seek a quantum analogue of the joint distribution $\mathbb{P}(X, Y, Z)$. However, textbook quantum theory does not provide such an analogue when some systems are causally dependent on others and there are serious reasons to believe that such an analogue might be impossible [HHP⁺16]. This is the reason for concentrating on the channels $\mathbb{P}(Y, Z|X)$ and $\rho_{BC|A}$ and seeking conditions that apply to them directly.

In Sec. 4.2.2 a justification of the quantitative part of Reichenbach’s principle from the qualitative part was given by temporarily assuming that classical dynamics is fundamentally




Figure 4.3: A causal structure depicting three quantum systems with A the complete common cause of B and C .

deterministic. Suppose, by analogy, that quantum dynamics is taken to be fundamentally unitary. What characterisation of $\rho_{BC|A}$ in Fig. 4.3 follows from that? Just as in Sec. 4.2.2, this is a temporary assumption in order to motivate the final result, which will stand apart from it.

It is well known that any quantum channel $\rho_{B|A}$ (equivalently, $\mathcal{E}_{B|A}$) can be viewed as arising from underlying unitary dynamics via the Stinespring dilation [NC00].

Definition 4.3. A unitary dilation of a quantum channel $\mathcal{E}_{B|A}$ is given by some unitary U on $\mathcal{H}_B \otimes \mathcal{H}_F \cong \mathcal{H}_A \otimes \mathcal{H}_L$ for some ancillary system L in state ρ_L such that

$$\mathcal{E}_{B|A}(\cdot) = \text{Tr}_F (U(\cdot \otimes \rho_L)U^\dagger), \quad (4.8)$$

where F is some system of dimension $d_F = d_A d_L / d_B$.

Applying this to the common cause situation of Fig. 4.3 gives

$$\rho_{BC|A} = \text{Tr}_{FLL^*} (\rho_{BFC|AL}^U \tau_L^{\text{id}} \rho_L) \quad (4.9)$$

where $\rho_{BFC|AL}^U$ is the Choi-Jamiołkowski operator for the unitary and τ_L is defined in Eq. (4.6). Note that F is required so that the input and output dimensions match, but is unimportant and will normally be traced out.

Classically, in Sec. 4.2.2, it was argued that the ancilla variable Λ must split into two independent variables to preserve X as a complete common cause. This implicitly uses the obvious idea that if a deterministic causal relationship $f : X \times \bar{X} \rightarrow Y$ cannot be equivalently written as some other function $f' : X \rightarrow Y$, then \bar{X} must have a causal influence on Y . How could it not? Its value non-trivially affects Y .

In unitary quantum theory the corresponding condition is slightly less obvious, but can be made precise as follows.

Definition 4.4. For a unitary channel $\rho_{B\bar{B}|A\bar{A}}$, system \bar{A} has no causal influence on system B if and only if the marginal output state at B is independent of any extra operations applied to the input \bar{A} system before applying $\rho_{B\bar{B}|A\bar{A}}$.




Figure 4.4: The causal structure of Fig. 4.3, dilated so that B and C each has a latent system as a causal parent in addition to A . B and C (together with F , not shown) depend unitarily on their parents L_B , A , and L_C .

This definition captures the idea that interventions on \bar{A} (including preparing some specific input state) cannot affect the local output at B and therefore is a sensible notion for “no causal influence”. Similar properties of unitary channels have been studied before in different contexts, in particular in Ref. [SW05] under the guise of “non-signalling” and in Refs. [BGNP01, ESW02] for “semi-causal” unitaries. An equivalent definition, more directly useful to the treatment here, is that \bar{A} has no causal influence on B if and only if the partial trace of the channel satisfies¹ $\rho_{B|A\bar{A}} = \rho_{B|A} \otimes \mathbb{1}_{\bar{A}^*}$.

Just as in Sec. 4.2.2, the assumption that A is a *complete* common cause for B and C therefore implies that L must factorise into ancestrally independent L_B and L_C where L_B has no causal influence on C and L_C has no causal influence on B . It follows that the unitary channel U (followed by tracing out F) is of the form illustrated in Fig. 4.4. Using the qualitative part of Reichenbach’s principle, ancestral independence of L_B , A , and L_C implies that the input state to $\rho_{BFC|L_B A L_C}^U$ factorises as $\rho_{L_B A L_C} = \rho_{L_B} \rho_A \rho_{L_C}$.

Similarly to the classical case, temporarily assuming that quantum dynamics is unitary has led to the following suggestion for characterising quantum common cause channels.

Definition 4.5. *A quantum channel $\rho_{BC|A}$ is compatible with A being the unitary common cause for B and C if and only if there is a unitary dilation U of the channel of the form*

$$\rho_{BC|A} = \text{Tr}_{F L_B L_B^* L_C L_C^*} \left(\rho_{BFC|L_B A L_C}^U \tau_{L_B}^{\text{id}} \tau_{L_C}^{\text{id}} \rho_{L_B} \rho_{L_C} \right) \quad (4.10)$$

where L_B has no causal influence on C and L_C has no causal influence on B in U .

Despite the specific way it was reached, this is a very natural quantum condition. It strongly suggests the following quantum analogue of Thm. 4.1.

Proposition 4.1. *The following are equivalent:*

¹To see this, consider the marginal output at B gained by preceding $\rho_{B\bar{B}|A\bar{A}}$ with a channel that discards the input at \bar{A} and replaces it with $\mathbb{1}/d_{\bar{A}}$. This must be the same as the marginal output from the original channel, so Eq. (4.5) gives the stated result.

1. $\rho_{BC|A}$ is compatible with A being the unitary complete common cause for B and C .
2. $\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$.

This proposition is proved as part of the more general Thm. 4.2 in Sec. 4.3.4. Note that condition (2) has no ordering ambiguity as Hermiticity guarantees that $[\rho_{B|A}, \rho_{C|A}] = 0$. It is not difficult to verify that both conditions of Prop. 4.1 reduce to their corresponding classical statements in any decohering limit (where all states and channels are diagonal in some choices of “classical” bases).

That Prop. 4.1 holds [Sec. 4.3.4] and is strongly analogous to Thm. 4.1 suggests the following definition for quantum conditional independence (at least where A is in the causal past of B and C).

Definition 4.6. *For systems A , B , and C related by quantum channel $\rho_{BC|A}$, the outputs are quantum conditionally independent given the input if and only if $\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$.*

Finally, with these definitions in hand, it is easy to state the quantitative part of quantum Reichenbach’s principle: in the case where B and C are correlated due exclusively to some complete common cause A , then B and C are quantum conditionally independent given A for the channel $\rho_{BC|A}$. Combined with the qualitative part discussed above, this completes the definition of quantum Reichenbach’s principle.

So in Prop. 4.1, the quantitative part of quantum Reichenbach’s principle is motivated by the (1) \Rightarrow (2) statement. The converse, (2) \Rightarrow (1), should be useful for causal inference, just as in Thm. 4.1

These definitions will form the motivation for most of the work in Chap. 5. For the remainder of this chapter, these definitions will be expanded and applied in various ways.

4.3.3 Alternative Expressions for Quantum Conditional Independence

There are many equivalent ways to define conditional independence of a channel $\mathbb{P}(Y, Z|X)$ classically. Two of these are given in Thm. 4.1: there exists a deterministic dilation that factorises between Y and Z ; and $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$. Two other useful definitions are as follows.

The first is that for every input $\mathbb{P}(X)$ to the channel, the joint distribution over input and output values $\mathbb{P}(Y = y, Z = z|X = x)\mathbb{P}(X = x)$ has vanishing conditional mutual information $I(Y : Z|X) = 0$. Equivalently, this need only hold when the input distribution is uniform $\mathbb{P}(X = x) = 1/|X|$, so that $I(Y : Z|X) = 0$ on $\hat{\mathbb{P}}(Y, Z, X) \stackrel{\text{def}}{=} \mathbb{P}(Y, Z|X)/|X|$, where $|X|$ is the cardinality of X .

The second definition is that $\mathbb{P}(Y, Z|X)$ is mathematically equivalent to the channel obtained by doing the following: (i) copy the input X , (ii) apply channel $\mathbb{P}(Y|X)$ to only one copy of X , (iii) apply channel $\mathbb{P}(Z|X)$ to only the other copy.

That these are both equivalent definitions to $\mathbb{P}(Y, Z|X) = \mathbb{P}(Y|X)\mathbb{P}(Z|X)$ are standard results and easily verified. Mathematically, there is no reason to consider one definition as “more fundamental” than the others (that depends on your philosophical bent). The definitions used in Thm. 4.1 were only given first because they are the closest to the motivational narrative being used. Another approach might have led to presenting other definitions first.

Like the conditions of Thm. 4.1, these two definitions also have natural quantum counterparts. They provide two new ways to define quantum conditional independence for a channel, equivalent to those in Prop. 4.1. Just as in the classical case, none of these definitions need be considered more fundamental than any other.

Proposition 4.2. *The following are equivalent to the conditions of Prop. 4.1 and to each other:*

3. $I(B : C|A) = 0$ for the quantum conditional mutual information evaluated on the (positive, trace-one) operator $\hat{\rho}_{BC|A}$.
4. The Hilbert space of A has a decomposition $\mathcal{H}_A = \bigoplus_i \mathcal{H}_{A_i^L} \otimes \mathcal{H}_{A_i^R}$ for which $\rho_{BC|A} = \sum_i \left(\rho_{B|A_i^L} \otimes \rho_{C|A_i^R} \right)$, where for each i , $\rho_{B|A_i^L}$ is a quantum channel from A_i^L to B and $\rho_{C|A_i^R}$ is a channel from A_i^R to C .

The proof of this, with Prop. 4.1, is in Sec. 4.3.4. It is easy to verify that condition (3) reduces to the classical case in a decohering limit when everything diagonalises in some choice of basis. Recall that $\hat{\rho}_{BC|A}$ is just $\rho_{BC|A}$ re-normalised to be trace-one.

Condition (4) deserves significantly more discussion, provided in Secs. 4.3.5, 5.4. For now, it suffices to note that any classical channel which copies the input and applies independent channels to each output will take the form of condition (4) when written as a quantum channel. It therefore generalises that alternative definition of conditional independence.

4.3.4 Proving Propositions 4.1 and 4.2

Propositions 4.1 and 4.2 are more naturally taken together as a single theorem. All four conditions given are equally valid ways to define quantum conditional independence for a channel and combining them into one theorem reflects this. Proving their equivalence together also requires fewer steps. The full theorem is restated here for convenience.

Theorem 4.2. *Given a quantum channel $\rho_{BC|A}$, the following are equivalent:*

1. $\rho_{BC|A}$ is compatible with A being the unitary complete common cause for B and C .
2. $\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$.
3. $I(B : C|A) = 0$ for the quantum conditional mutual information evaluated on the (positive, trace-one) operator $\hat{\rho}_{BC|A}$.
4. The Hilbert space of A has a decomposition $\mathcal{H}_A = \bigoplus_i \mathcal{H}_{A_i^L} \otimes \mathcal{H}_{A_i^R}$ for which $\rho_{BC|A} = \sum_i \left(\rho_{B|A_i^L} \otimes \rho_{C|A_i^R} \right)$, where for each i , $\rho_{B|A_i^L}$ is a quantum channel from A_i^L to B and $\rho_{C|A_i^R}$ is a channel from A_i^R to C .

Each of these conditions is an equivalent definition for when B and C are quantum conditionally independent given A in a channel $\rho_{BC|A}$.

This theorem will now be proved in parts. First, consider the following lemma proved in Ref. [HJPW04].

Lemma 4.1 [HJPW04, Thm. 6]. *For any tripartite quantum state ρ_{ABC} , the quantum conditional mutual information evaluated on that state vanishes $I(B : C|A) = 0$ if and only if the Hilbert space of the A system decomposes as $\mathcal{H}_A = \bigoplus_i \mathcal{H}_{A_i^L} \otimes \mathcal{H}_{A_i^R}$, such that*

$$\rho_{ABC} = \sum_i p_i \left(\rho_{BA_i^L} \otimes \rho_{CA_i^R} \right), \quad p_i \geq 0, \quad \sum_i p_i = 1, \quad (4.11)$$

where for each i , $\rho_{BA_i^L}$ is a quantum state on $\mathcal{H}_B \otimes \mathcal{H}_{A_i^L}$ and $\rho_{CA_i^R}$ is a quantum state on $\mathcal{H}_C \otimes \mathcal{H}_{A_i^R}$.

This provides the first part of the proof of Thm. 4.2.

Proof: (3) \Leftrightarrow (4). Applying Lem. 4.1 to $\hat{\rho}_{BC|A}$ immediately proves that (4) \Rightarrow (3) and gets most of the way to proving (3) \Rightarrow (4). To complete the proof, note that $\text{Tr}_{BC} \hat{\rho}_{BC|A} = \mathbb{1}_A/d_A$. It follows that each $\hat{\rho}_{B|A_i^L}$ given by Lem. 4.1 satisfies $\text{Tr}_B \hat{\rho}_{B|A_i^L} = \mathbb{1}_{A_i^L}/d_{A_i^L}$ and similarly for $\hat{\rho}_{C|A_i^R}$. Therefore each $\hat{\rho}_{B|A_i^L}$ and each $\hat{\rho}_{C|A_i^R}$ are appropriate quantum channels, completing the proof. \square

The next piece of the proof is perhaps the most involved as it demonstrates that the unitary dilation implied by condition (1) has the structure required by condition (4).

Proof: (1) \Rightarrow (4). Let $\rho_{BFC|L_B A L_C}^U$ be the Choi-Jamiołkowski operator for the unitary U as in the Sec. 4.3.2. Note that $\rho_{BC|A}^U \neq \rho_{BC|A}$ in general (the latter depends on particular choices for inputs at L_B and L_C —cf. Def. 4.5). The proof proceeds by demonstrating that several conditional mutual informations evaluated on $\hat{\rho}_{BFC|L_B A L_C}^U \stackrel{\text{def}}{=} \rho_{BFC|L_B A L_C}^U / (d_{L_B} d_A d_{L_C})$ vanish.

The first conditional mutual information of $\hat{\rho}_{BFC|L_B A L_C}^U$ that vanishes is

$$I(B : FC | L_B A L_C) = 0. \quad (4.12)$$

This follows by expanding in terms of von Neumann entropies

$$I(B : FC | L_B A L_C) = S(\hat{\rho}_{B|L_B A L_C}^U) + S(\hat{\rho}_{FC|L_B A L_C}^U) - S(\hat{\rho}_{BFC|L_B A L_C}^U) - S(\hat{\rho}_{\cdot|L_B A L_C}^U). \quad (4.13)$$

The third term is zero, since $\hat{\rho}_{BFC|L_B A L_C}^U$ is pure by unitarity. The final term is equal to $\log(d_{L_B} d_A d_{L_C})$, since $\hat{\rho}_{\cdot|L_B A L_C}^U = \mathbb{1}_{(L_B A L_C)}^* / (d_{L_B} d_A d_{L_C})$. Noting also that $\hat{\rho}_{BFC|\cdot}^U = \mathbb{1}_{BFC} / (d_B d_F d_C)$, and using the fact that the von Neumann entropy of the partial trace of a pure state is equal to the von Neumann entropy of the complementary partial trace, one finds that the first two terms equal $\log(d_F d_C)$ and $\log d_B$ respectively. Summing these, Eq. (4.12) follows.

Second

$$I(L_B : L_C | A) = 0, \quad (4.14)$$

which follows immediately from $\hat{\rho}_{\cdot|L_B A L_C}^U = \mathbb{1}_{(L_B A L_C)}^* / (d_{L_B} d_A d_{L_C})$.

Third,

$$I(B : L_C | L_B A) = 0. \quad (4.15)$$

As with Eq. (4.12), prove this by expanding

$$I(B : L_C | L_B A) = S(\hat{\rho}_{B|L_B A}^U) + S(\hat{\rho}_{\cdot|L_B A L_C}^U) - S(\hat{\rho}_{B|L_B A L_C}^U) - S(\hat{\rho}_{\cdot|L_B A}^U). \quad (4.16)$$

The second and fourth terms are entropies of maximally mixed states on their respective systems, which hence sum to $\log d_{L_C}$. Since there is no causal influence from L_C to B in U it follows that $\hat{\rho}_{B|L_B A L_C}^U = \hat{\rho}_{B|L_B A}^U \otimes \mathbb{1}_{L_C}^* / d_{L_C}$. Hence, the third term is $S(\hat{\rho}_{B|L_B A}^U) + \log d_{L_C}$, which gives Eq. (4.15).

Fourth and finally

$$I(C : L_B | A L_C) = 0, \quad (4.17)$$

which follows symmetrically to Eq. (4.15) by using the assumption that there is no influence from L_B to C in U .

Equations (4.12, 4.14, 4.15, 4.17) can be used to show that $\hat{\rho}_{BC|A}$ satisfies $I(B : C|A) = 0$. This follows from Ref. [LP08, Thm. 4.5], which states that quantum conditional mutual informations on partial traces of a multipartite quantum state satisfy the *semi-graphoid axioms* familiar from the classical formalism of causal networks [Pea09]. Therefore, the semi-graphoid axioms satisfied in this case are:

$$[I(X : Y|Z) = 0] \Rightarrow [I(Y : X|Z) = 0], \quad (4.18)$$

$$[I(X : YW|Z) = 0] \Rightarrow [I(X : Y|Z) = 0], \quad (4.19)$$

$$[I(X : YW|Z) = 0] \Rightarrow [I(X : Y|ZW) = 0], \quad (4.20)$$

$$[I(X : Y|Z) = 0] \wedge [I(X : W|YZ) = 0] \Rightarrow [I(X : YW|Z) = 0]. \quad (4.21)$$

By applying Eqs. (4.18–4.21) to Eqs. (4.12, 4.14, 4.15, 4.17) one finds

$$[I(B : FC|L_B A L_C) = 0] \Rightarrow [I(B : C|L_B A L_C) = 0] \quad (4.22)$$

$$[I(C : L_B|A L_C) = 0] \wedge [I(B : C|L_B A L_C) = 0] \Rightarrow [I(C : B L_B|A L_C) = 0] \quad (4.23)$$

$$[I(L_B : L_C|A) = 0] \wedge [I(L_C : B|L_B A) = 0] \Rightarrow [I(L_C : B L_B|A) = 0] \quad (4.24)$$

$$[I(B L_B : L_C|A) = 0] \wedge [I(B L_B : C|A L_C) = 0] \Rightarrow [I(B L_B : C L_C|A) = 0] \quad (4.25)$$

This shows that condition (1) implies $I(B L_B : C L_C|A) = 0$, calculated on $\hat{\rho}_{BC|L_B A L_C}^U$. Using Lem. 4.1 gives

$$\hat{\rho}_{BC|L_B A L_C}^U = \sum_i p_i \left(\hat{\rho}_{B|L_B A_i^L}^U \otimes \hat{\rho}_{C|A_i^R L_C}^U \right), \quad (4.26)$$

for some appropriate decomposition of \mathcal{H}_A^* and probability distribution $\{p_i\}_i$. This decomposition, with the fact that $\hat{\rho}_{L_B A L_C}^U = \mathbb{1}_{(L_B A L_C)^*} / (d_{L_B} d_A d_{L_C})$, implies

$$\rho_{BC|L_B A L_C}^U = \sum_i \left(\rho_{B|L_B A_i^L}^U \otimes \rho_{C|A_i^R L_C}^U \right), \quad (4.27)$$

where, for each i , the components satisfy $\text{Tr}_B \rho_{B|L_B A_i^L}^U = \mathbb{1}_{(L_B A_i^L)^*}$ and $\text{Tr}_C \rho_{C|A_i^R L_C}^U = I_{(A_i^R L_C)^*}$. Definition 4.5 shows that the operator $\rho_{BC|A}$ is obtained by acting with this channel on the input states ρ_{L_B} and ρ_{L_C} respectively for L_B and L_C . Finally, therefore

$$\rho_{BC|A} = \sum_i \left(\rho_{B|A_i^L} \otimes \rho_{C|A_i^R} \right), \quad (4.28)$$

where $\text{Tr}_B \rho_{B|A_i^L} = \mathbb{1}_{(A_i^L)^*}$ and $\text{Tr}_C \rho_{C|A_i^R} = \mathbb{1}_{(A_i^R)^*}$, as required. \square

To prove the converse, that condition (4) implies (1), one can show how to construct an appropriate unitary dilation.

Proof: (4) \Rightarrow (1). Each channel $\rho_{B|A_i^L}$ can be dilated to a unitary transformation V_i , with ancilla input L_B in a fixed state ρ_{L_B} , such that V_i acts on $\mathcal{H}_{L_B} \otimes \mathcal{H}_{A_i^L}$. Similarly, $\rho_{C|A_i^R}$ can be dilated to a unitary transformation W_i , with ancilla L_C in a fixed state ρ_{L_C} , acting on $\mathcal{H}_{A_i^R} \otimes \mathcal{H}_{L_C}$. By choosing the dimension of L_B large enough, a single system L_B and state ρ_{L_B} can be used for each value of i and similarly for L_C .

For each i , let V'_i be the unitary that acts as $V_i \otimes \mathbb{1}_{A_i^R} \otimes \mathbb{1}_{L_C}$ on the subspace $\mathcal{H}_{L_B} \otimes \mathcal{H}_{A_i} \otimes \mathcal{H}_{L_C}$ and as zero on each subspace $\mathcal{H}_{L_B} \otimes \mathcal{H}_{A_j} \otimes \mathcal{H}_{L_C}$ where $j \neq i$. Similarly, for each i let W'_i be the unitary that acts as $\mathbb{1}_{L_B} \otimes \mathbb{1}_{A_i^L} \otimes W_i$ on the subspace $\mathcal{H}_{L_B} \otimes \mathcal{H}_{A_i} \otimes \mathcal{H}_{L_C}$, and as zero on each subspace $\mathcal{H}_{L_B} \otimes \mathcal{H}_{A_j} \otimes \mathcal{H}_{L_C}$ for every $j \neq i$. Using these, define unitaries

$$V \stackrel{\text{def}}{=} \sum_i V'_i, \quad (4.29)$$

$$W \stackrel{\text{def}}{=} \sum_i W'_i. \quad (4.30)$$

Note that: they commute $[V, W] = 0$; V acts as the identity on L_C ; and W acts as the identity on L_B .

It can easily be checked that the unitary $WV = VW$ is a unitary dilation for the channel $\rho_{BC|A}$ in the form of condition (4), with ancillae L_B and L_C . From the form of V and W it can also be seen that there can be no causal influence from L_C to B in WV (the output at B already exists, by the action of V , before W acts on L_C at all). Similarly, since $WV = VW$, there can be no causal influence from L_B to C . The unitary WV is therefore the unitary dilation required by condition (1). \square

These final two pieces complete the proof of Thm. 4.2.

Proof: (2) \Rightarrow (3). Assuming a channel $\rho_{BC|A} = \rho_{B|A}\rho_{C|A}$ it follows that:

$$\rho_{BC|A} = \exp(\log \rho_{B|A} + \log \rho_{C|A}) \quad (4.31)$$

$$\log \rho_{BC|A} = \log \rho_{B|A} + \log \rho_{C|A} \quad (4.32)$$

$$\log \rho_{BC|A} + \log \rho_{\cdot|A} = \log \rho_{B|A} + \log \rho_{C|A} \quad (4.33)$$

$$\log(d_A^{-1} \rho_{BC|A}) + \log(d_A^{-1} \rho_{\cdot|A}) = \log(d_A^{-1} \rho_{B|A}) + \log(d_A^{-1} \rho_{C|A}). \quad (4.34)$$

The first line follows because $[\rho_{B|A}, \rho_{C|A}] = 0$ (as noted in Sec. 4.3.2); the third because $\rho_{\cdot|A} = \mathbb{1}_{A^*}$ and therefore $\log \rho_{\cdot|A}$ is the zero matrix; and the final line by adding $2 \log d_A^{-1}$ to both sides. It is proved in Ref. [Rus02] that if ρ_{XYZ} is any trace-one density operator, then $\log \rho_{XYZ} + \log \rho_Z = \log \rho_{XZ} + \log \rho_{YZ}$ is equivalent to $I(X : Y|Z) = 0$. Therefore, the conditional mutual information of $\hat{\rho}_{BC|A}$ vanishes. \square

Proof: (4) \Rightarrow (2). Condition (4) implies

$$\rho_{B|A} = \sum_i \left(\rho_{B|A_i^L} \otimes \mathbb{1}_{(A_i^R)^*} \right), \quad (4.35)$$

$$\rho_{C|A} = \sum_j \left(\mathbb{1}_{(A_j^L)^*} \otimes \rho_{C|A_j^R} \right). \quad (4.36)$$

Taking the product, terms where $i \neq j$ have support on orthogonal subspaces and so vanish, therefore

$$\begin{aligned} \rho_{B|A} \rho_{C|A} &= \sum_{i,j} \left(\rho_{B|A_i^L} \otimes \mathbb{1}_{(A_i^R)^*} \right) \left(\mathbb{1}_{(A_j^L)^*} \otimes \rho_{C|A_j^R} \right) = \sum_i \rho_{B|A_i^L} \otimes \rho_{C|A_i^R} \quad (4.37) \\ &= \rho_{BC|A}. \quad (4.38) \end{aligned}$$

□

Combining these proof steps, any of the conditions of Thm. 4.2 can be seen to imply any other. Moreover, some of the contents of the proofs are somewhat revealing. For instance, in the proof for (4) \Rightarrow (1) a specific possible structure for the unitary dilation is revealed, one that makes the no-causal-influence properties explicit. Some of the theorems used from other papers also point to places in the literature concerned with similar ideas and which may be fruitful places to apply these results.

4.3.5 Circuits for Common Cause Channels

The definitions for quantum conditional independence given in Thm. 4.2 can be considered in the language of quantum circuits. This is useful both to elucidate their meanings and for comparison with the corresponding classical definitions.

Figure 4.5 illustrates several equivalent ways to view a classical conditionally-independent channel $\mathbb{P}(Y, Z|X)$ (*viz.* one where X could be a complete common cause). Equality (1) of that figure simply repeats that any classical channel can be dilated to a deterministic channel by a function f and random variable Λ that absorbs all stochasticity. The other three equalities are more interesting and may be understood as forming a commuting diagram, showing two equivalent ways to get from the top left to bottom right. That is, two different ways of understanding that $\mathbb{P}(Y, Z|X)$ is compatible with X being the deterministic common cause for Y and Z .

Equality (2) of Fig. 4.5 states that if Y and Z are conditionally independent given X , then the channel $\mathbb{P}(Y, Z|X)$ may be achieved by first copying X and then applying separate channels $\mathbb{P}(Y|X)$ and $\mathbb{P}(Z|X)$ to each copy. The symbol \blacktriangleright is used to represent




Figure 4.5: Circuit representations of a classical channel $\mathbb{P}(Y, Z|X)$ where Y and Z are conditionally independent given X (equivalently, where X could be a complete common cause for Y and Z). The symbol \blacktriangleright represents the classical copy operation. Note that equalities (1) and (4) do not depend on conditional independence when viewed in isolation, but they do in the context of the other equalities in the diagram.

the copying operation. Equality (4) mirrors equality (1): the channels $\mathbb{P}(Y|X)$ and $\mathbb{P}(Z|X)$ can be individually dilated to functions f_Y and f_Z respectively.

Finally, equality (3) follows as X is a *complete* common cause for Y and Z . Together, equalities (1) and (3) illustrate the justification for classical Reichenbach's principle given in Sec. 4.2.2.

Mirroring this, Fig. 4.6 illustrates the analogous conditions for quantum conditional independence of channel $\rho_{BC|A}$. In order to do this some non-standard notation in quantum circuits has been introduced in order to capture condition (4) of Thm. 4.2. This figure may similarly be read as a commuting diagram.

Equality (1) of Fig. 4.6 reaffirms that $\rho_{BC|A}$ has a unitary dilation U . The system being traced-out (the symbol \ddagger) is F from Sec. 4.3.2.

Equality (2) expresses condition (4) of Thm. 4.2. In order to do this the symbol \oplus has been introduced. \oplus may be read as re-expressing the Hilbert space of input system A as a direct sum of factorising spaces $\mathcal{H}_A = \bigoplus_i \mathcal{H}_{A_i^L} \otimes \mathcal{H}_{A_i^R}$. The output wire $\{A_i^L\}$ then carries the left factor of each subspace, which is then acted upon by the relevant channel from $\{\rho_{B|A_i^L}\}$ (similarly for the other output wire). This is non-standard for quantum circuits, where wires normally correspond to a single Hilbert space and adjacent wires represent the tensor product of Hilbert spaces. Here, wires labelled as $\{A_i^L\}$ and $\{A_i^R\}$ only represent a whole Hilbert space when considered together as a pair. This has the




Figure 4.6: Circuit representations of a quantum channel $\rho_{BC|A}$ where B and C are conditionally independent given A (equivalently, where A could be a complete common cause for B and C). The meaning of the symbol \textcircled{i} is discussed in the text. \ddagger denotes taking the partial trace (discarding) of the indicated subsystem.

advantage of allowing factorised channels $\{\rho_{B|A_i^L}\}$ and $\{\rho_{C|A_i^R}\}$ to act on separate wires. The interpretation of these symbols will be discussed further below.

Equality (4) then applies unitary dilations to each $\rho_{B|A_i^L}$ and $\rho_{C|A_i^R}$, such that all dilations (for each i) use a common ancillary input L_B and L_C respectively.

There are two ways to understand equality (3). First, simply as the conjunction of equalities (1), (2), and (4). Second, as expressing that the unitary U has the no-causal-influence properties demanded by Thm. 4.2. There is no direct way to express the lack of causal influence in a unitary diagrammatically. However, in the proof of Thm. 4.2 the existence of a unitary dilation of the form shown here is seen to be equivalent to the existence of a unitary with appropriate no-causal-influence properties. So the equalities (1) and (3) taken together might be viewed as illustrating the justification of quantum Reichenbach of Sec. 4.3.2 (as in the classical case) but the correspondence is not immediate.

The introduction of the new symbol \textcircled{i} to quantum circuits should, by rights, be accompanied by a thorough description of its meaning. However, in this case it is useful to leave the meaning somewhat ambiguous. For the purposes of this thesis \textcircled{i} need only have meaning in the diagrams of Fig. 4.6. There are multiple ways to interpret the symbol such that those diagrams are well-defined and it would be premature to pick one above the others before a more general diagrammatic use of \textcircled{i} has been found.

One interpretation of \textcircled{i} is the *passive* interpretation: it simply re-interprets the in-

put Hilbert space as a direct sum of factorising spaces indexed by i , analogously to reinterpreting a single Hilbert space as a tensor product of two factors. Indeed, when i only takes one value \oplus is just simple factorisation. In this view, \oplus is entirely reversible and no physical operation occurs until the output wires are acted on. As noted above, the output wires only form a complete Hilbert space when taken together, but allow appropriate sets of channels to act independently on each factor.

Another valid interpretation of \oplus is the *active* interpretation: it represents a von Neumann measurement on A defined by the linear subspaces labelled by i , followed by a factorisation of the output system depending on the outcome i . This is clearly a non-reversible physical operation.

A conditionally independent channel satisfying Thm. 4.2 is decohering across the linear subspaces labelled by i . Under the passive interpretation, this decoherence doesn't take hold until the sets of channels $\{\rho_{B|A_i^L}\}$ and $\{\rho_{C|A_i^R}\}$ are applied. Under the active interpretation on the other hand, the decoherence occurs at \oplus due to the measurement. This allows the subsequent channels, $\rho_{B|A_i^L}$ etc., to be ordinary quantum channels, where the appropriate ones are selected depending on the outcome i .

Both interpretations are compatible with the definitions of Thm. 4.2. The passive interpretation might be preferred as a description of channels that satisfy quantum conditional independence, while the active interpretation gives a concrete way to achieve quantum conditional independence operationally. Of course, there are other possible interpretations too.

If \oplus is only used in Fig. 4.6, why should the interpretation matter? Comparing Figs. 4.5 and 4.6 suggests that \oplus is, in some way, analogous to the classical copy \Uparrow . It represents a general way for two agents to independently act on a single system A , just as classically two agents can independently use input X by first copying it. There is certainly potential for fleshing out the relationship between \oplus and \Uparrow more generally.

The symbol \oplus is most closely related to condition (4) of Thm. 4.2 and both active and passive interpretations can be applied to that definition. Given the above discussion, condition (4) seems to require channels where Bob and Clare can work independently on a single input. Importantly, it shows that there are more general ways of achieving this than simply factorising the quantum input. Condition (4) therefore appears to hint at a general characterisation of when agents can share a single input while working independently.

4.3.6 Examples

Definitions 4.5 and 4.6 respectively defined common causes and conditional independence for quantum channels, which were then used to identify a quantum Reichenbach's principle.




Figure 4.7: (a) A generic unitary transformation from AD to BC . (b) The corresponding causal structure, making no assumptions about the properties of U .

It is now time to consider some examples to see how these apply in practice, as well as to double check that the definitions given are reasonable.

4.3.6.1 Generic Unitary Transformations

Consider the circuit of Fig. 4.7(a) where systems A and D unitarily evolve to B and C . Figure 4.7(b) shows the causal structure one would expect for this basic situation. It would be very troubling indeed if the definition proposed in this chapter forbade AD from being a complete common cause for B and C (that is, if B and C were not quantum conditionally independent given AD).

Fortunately, this is not the case. The channel $\rho_{BC|AD}$ in this example is its own unitary dilation which trivially satisfies Def. 4.5. Indeed, it is easy to verify directly that the other conditions of Thm. 4.2 are also satisfied. For example, $\hat{\rho}_{BC|AD}$ is a pure state where AD is maximally entangled with BC , so $I(B : C|A) = 0$. Therefore Def. 4.6 and Thm. 4.2 pass the most basic of sanity-checks.

4.3.6.2 Coherent and Incoherent Copies

Perhaps the simplest non-trivial example of a classical common cause is copying a bit. That is, the classical channel for binary random variables $X, Y, Z \in \{0, 1\}$

$$\mathbb{P}(Y, Z|X = x) = \delta(Y, x)\delta(Z, x) \quad (4.39)$$

which simply sets both outputs Y and Z to match the input X . Clearly this channel satisfies classical conditional independence and, perhaps most importantly, it is intuitively obvious that this should be the case: Y and Z simply take the value of X , so X must explain all correlations between them.

One quantum generalisation for this classical channel is the *incoherent copy* channel for qubits $d_A = d_B = d_C = 2$

$$\alpha|0\rangle_A + \beta|1\rangle_A \rightarrow |\alpha|^2|00\rangle_{BC}\langle 00| + |\beta|^2|11\rangle_{BC}\langle 11|. \quad (4.40)$$

This channel reduces to the classical bit-copy in the case where inputs are diagonal in the $\{|0\rangle, |1\rangle\}$ basis. The Choi-Jamiołkowski state for this channel is

$$\rho_{BC|A}^{\text{inc}} \stackrel{\text{def}}{=} |000\rangle_{BCA^*} \langle 000| + |111\rangle_{BCA^*} \langle 111|. \quad (4.41)$$

It is easy to verify that this satisfies all of the conditions of Thm. 4.2 and therefore satisfies quantum conditional independence according to Def. 4.6. This should be unsurprising, the incoherent copy is essentially a classical channel which removes any quantum coherence.

There is, however, another easy quantum generalisation for the classical bit-copy. That is, the *coherent copy* channel

$$\alpha|0\rangle_A + \beta|1\rangle_A \rightarrow \alpha|00\rangle_{BC} + \beta|11\rangle_{BC} \quad (4.42)$$

defined by Choi-Jamiołkowski operator

$$\rho_{BC|A}^{\text{coh}} \stackrel{\text{def}}{=} (|000\rangle + |111\rangle)_{BCA^*} (\langle 000| + \langle 111|). \quad (4.43)$$

This channel also reduces to the classical bit-copy in the same diagonal cases. However, this channel fails to satisfy the conditions of Thm. 4.2 and therefore B and C are not conditionally independent given A . In particular, it is easy to verify that $I(B : C|A) = 1$ on $\hat{\rho}_{BC|A}^{\text{coh}}$. A cannot be considered a complete common cause for B and C in the coherent copy.

This should seem troubling. Just as in the classical case, B and C take the value of A in Eq. (4.42), so how can A fail to be a complete common cause? There are at least three ways to understand this apparent discrepancy and see that, far from being a failure of quantum conditional independence, it is exactly what should hold.

The first is to recall the discussion of Sec. 4.3.5, where it was noted that conditionally independent channels capture the idea that two agents can “act independently” on a single input. Classically, this is easy: simply copy the input and give one to each agent. In quantum theory, this is achieved by using $\textcircled{\otimes}$ in place of the classical copy. The coherent copy creates arbitrary entanglement between two qubits from a single qubit. From this perspective, if Eq. (4.42) satisfied quantum conditional independence then agents would be able to create entanglement by acting independently and without communication on a single qubit. This would certainly be surprising, perhaps even perverse (for one, it would smell of remote entanglement preparation [BHL⁺05]), and therefore it is encouraging that Thm. 4.2 rules it out.

The second is to consider the deterministic/unitary dilations of the channels Eqs. (4.39, 4.42), shown in Figs. 4.8 and 4.9 respectively. Classically, the causal influence in a CNOT




Figure 4.8: (a) Classical deterministic dilation of the bit-copy from X to Y and Z . The ancilla Λ is prepared with the value 0 then a CNOT gate is applied, controlled on X . (b) The corresponding causal structure of this dilation, note that there is no causal influence from Λ to Y .




Figure 4.9: (a) Quantum unitary dilation of the coherent copy from A to BC . (b) The corresponding causal structure of this dilation, note the explicit back-action from L to B .

gate is unidirectional, as shown in Fig. 4.8(b). The value at Λ has no causal influence on Y . However, this is not the case for the quantum CNOT gate, which has a back-action (see also Ref. [SW12]). This is reflected in Fig. 4.9, where L has a causal influence on both B and C and therefore is an additional common cause (preventing A from being the complete common cause).

These diagrams show just one example of a dilation for each channel (though arguably the most natural), but similar remarks hold for other possible dilations. In particular, it is not possible to find a unitary dilation of Eq. (4.42) where the ancilla does not act as an additional common cause for B and C .

One might worry that this back-action should not matter in the quantum CNOT, since the ancilla L always takes the state $|0\rangle_L$. But this does not change the fact that there is no way to achieve Eq. (4.42) unitarily without introducing an additional common cause. Moreover, for those who take pure quantum states to represent “maximal but incomplete” information about the system [CFS02b, Fuc02, Spe07, Lei06, Spe16] a pure quantum state is not analogous to a classical point distribution².

Finally, the third way to understand why the coherent copy fails to satisfy conditional independence is via the Bayesian updating procedure discussed in Sec. 5.3.5.2. The details of this example are discussed in that section, for now it suffices to say that if A is a complete common cause for B and C , then any knowledge about C gained at B should be expressible

²Indeed, in Spekkens’ Toy Model [Spe07] (a subset of quantum theory which explicitly models quantum states as maximal but incomplete information and the CNOT as a deterministic map, briefly discussed in Sec. 2.1.3) classical conditional independence fails for the coherent copy channel.

in terms of knowledge about A . Classically, this is always the case for a complete common cause and Sec. 5.3.5.2 demonstrates the same for quantum causal models. However, this is not generally possible for the coherent copy: B can (in general) locally learn more about C than can be expressed via A . This is a key property of a common cause that fails due to the entanglement between B and C .

So the quantum incoherent copy can be the result of a common cause, just as the classical copy channel, while the coherent copy cannot. Reviewing the reasons given above, the core of this discrepancy seems to come from the strength of entanglement that the coherent copy can create and how another qubit is necessary for its creation.

4.3.6.3 Bell Experiments

In Sec. 4.1.2, Bell experiments were given as one of the main motivations for quantum conditional independence. Simply put, Reichenbach's principle using classical conditional independence fails in these experiments. It is therefore fitting to check how the notion of quantum conditional independence introduced in Sec. 4.3.2 deals with Bell experiments.

A standard Bell experiment is of the following form. Dave prepares a pair of qubits in the entangled singlet state $|\Psi^-\rangle \stackrel{\text{def}}{=} (|0\rangle|1\rangle - |1\rangle|0\rangle) / \sqrt{2}$ and then gives the whole system to Alice. Alice then noiselessly distributes the qubits, one to Bob and one to Clare. Bob and Clare are each prevented from any means of causal influence on the other for the whole experiment (this is normally done by invoking relativistic locality and ensuring they are spacelike separated). Bob and Clare each independently decide on a choice of measurement (from a pre-determined set of possibilities) on their qubit which they perform, obtaining outcomes B_r and C_r respectively. This scenario is illustrated in Fig. 4.10. Note that the slightly unusual step has been taken here of separating the preparation of $|\Psi^-\rangle$ (by Dave) from the distribution of the qubits (by Alice).

The key result from Ref. [WS15] is that in general there is no classical causal model (and certainly no classical common cause) that explains the correlations between B_r and C_r without fine-tuning. A satisfying result would be that B_r and C_r are appropriately quantum conditionally independent, since the experimental set-up seems to demand an explanation in terms of a common cause from Alice and Dave.

While B_r and C_r are experimental outcomes, it is necessary to consider them as encoded in quantum systems to apply quantum conditional independence as defined here. That is, the outcomes B_r and C_r are just represented as diagonal density operators of quantum systems in some obvious way.

It immediately follows that B and C are quantum conditionally independent given A , since the channel $\rho_{BC|A}$ that describes Alice distributing the qubits to Bob and Clare

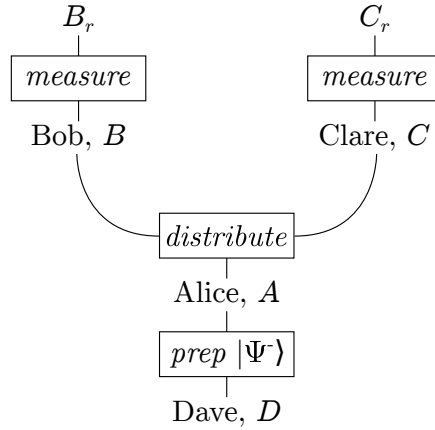


Figure 4.10: Schematic representation of a Bell experiment as described in the text. The measurements of Bob and Clare are chosen by them independently from a set of predefined choices.

is simply the identity channel. So conditional independence follows from Sec. 4.3.6.1. Since the measurement procedures of Alice and Bob are entirely factorised, it follows that B_r and C_r must also be quantum conditionally independent given A . It is simple to verify from Thm. 4.2 that composing a factorised channel (such as these measurements) after a quantum conditionally independent channel must result in another conditionally independent channel.

However, if one instead asks whether B and C (or B_r and C_r) are quantum conditionally independent given D , then the answer depends what system Dave starts with to prepare $|\Psi^-\rangle$. If, for example, Dave starts with a single qubit, then the composite channel $\rho_{BC|D}$ is a variant of the coherent copy and is therefore not conditionally independent. However, if Dave starts with a pair of qubits and unitarily evolves them to $|\Psi^-\rangle$, then the composite channel $\rho_{BC|D}$ is unitary and the channel satisfies conditional independence as above.

Should this be troubling? Not at all. This analysis shows that statistics from Bell experiments are quantum conditionally independent because all factorised quantum measurements are quantum conditionally independent, even when made on entangled states. This is as it should be, since conditional independence is a property of the channel and not of the channel-with-input. On the other hand, the cases where B and C (and therefore B_r and C_r) are not quantum conditionally independent given D are those where Dave's starting system is too simple to underwrite the correlations. But this also occurs classically: condition on too simple an ancestor and variables that were conditionally independent are rendered dependent again. One only expects conditional independence when conditioning on a complex enough ancestor.

4.3.7 Generalisation to k Outputs

Theorem 4.2 only considers quantum channels from a single input A to a pair of outputs BC . It is natural to ask whether a similar set of equivalences hold for channels from a single input A to an arbitrary finite set $\{B_l\}_{l=1}^k$ of k outputs. This is more than idle curiosity; it is a first step towards generalising to full quantum causal models, defined in Chap. 5.

Consider such a channel $\rho_{B_1\dots B_k|A}$, and let \bar{B}_l denote the whole joint output system excepting only B_l . Rather than starting from scratch, consider the following natural generalisation of Def. 4.5.

Definition 4.7. *A quantum channel $\rho_{B_1\dots B_k|A}$ is compatible with A being the unitary complete common cause for B_1, \dots, B_k if and only if there is*

- (a) *a unitary dilation U of the channel in terms of ancillae L_1, \dots, L_k with factorised initial state $\rho_{L_1} \otimes \dots \otimes \rho_{L_k}$ such that*
- (b) *for each l , L_l has no causal influence on \bar{B}_l .*

Beyond being perhaps the easiest generalisation of Def. 4.5, there is not yet any reason to suppose that this is a good definition. Its strength comes from the following theorem, which generalises Thm. 4.2.

Theorem 4.3. *Given a quantum channel $\rho_{B_1\dots B_k|A}$, the following are equivalent:*

1. *$\rho_{B_1\dots B_k|A}$ is compatible with A being a complete common cause of B_1, \dots, B_k .*
2. *$\rho_{B_1\dots B_k|A} = \rho_{B_1|A} \dots \rho_{B_k|A}$, where $[\rho_{B_l|A}, \rho_{B_m|A}] = 0$ for all l and m .*
3. *$I(B_l : \bar{B}_l|A) = 0$ for all l when evaluated on the (positive, trace-one) operator $\hat{\rho}_{B_1\dots B_k|A}$.*
4. *The Hilbert space of A has a decomposition $\mathcal{H}_A = \bigoplus_i \mathcal{H}_{A_i^1} \otimes \dots \otimes \mathcal{H}_{A_i^k}$ for which $\rho_{B_1\dots B_k|A} = \sum_i \left(\rho_{B_1|A_i^1} \otimes \dots \otimes \rho_{B_k|A_i^k} \right)$, where for each i and l , $\rho_{B_l|A_i^l}$ is a quantum channel from A_i^l to B_l .*

Each of these conditions is an equivalent definition for when B_1, \dots, B_k are quantum conditionally independent given A in a channel $\rho_{B_1\dots B_k|A}$.

Defining quantum conditional independence with these conditions is done by analogy with the classical and $k = 2$ cases. Again, each condition is an equally valid definition and there is no particular *a priori* reason to privilege one above any other. The whole theorem reduces to Thm. 4.2 if $k = 2$. The proof follows very similarly to that of Thm. 4.2 and is presented here in parts.

Proof: (3) \Rightarrow (2). The proof proceeds by induction and repeated use of Thm. 4.2. To this end, suppose the result holds for $\rho_{B_1 \dots B_n | A}$ for some $n < k$ (the inductive hypothesis) and that $I(B_{n+1} : \bar{B}_{n+1} | A) = 0$ as in condition (3).

By Thm. 4.2, this implies that

$$\rho_{B_1 \dots B_{n+1} | A} = \rho_{B_1 \dots B_n | A} \rho_{B_{n+1} | A} \quad (4.44)$$

where $[\rho_{B_1 \dots B_n}, \rho_{B_{n+1}}] = 0$. The inductive hypothesis implies

$$\rho_{B_1 \dots B_n | A} = \rho_{B_1 | A} \cdots \rho_{B_n | A} \quad (4.45)$$

where the ρ_{B_j} factors pairwise commute. Together these imply the factorisation of condition (2) for $\rho_{B_1 \dots B_{n+1} | A}$. To show that $\rho_{B_{n+1}}$ commutes with each ρ_{B_l} , simply note that $\rho_{B_1 \dots B_n} \rho_{B_{n+1}} = \rho_{B_{n+1}} \rho_{B_1 \dots B_n}$ and trace over all $B_{m \neq l}$. Therefore, using Thm. 4.2 as the base case the result holds for all $k \geq 2$ by induction. \square

Proof: (2) \Rightarrow (3). This follows immediately from Thm. 4.2 by letting $B = B_l$ and $C = \bar{B}_l$ for each l . \square

Proof: (3) \Rightarrow (4). Assume that (3) holds, so that $I(B_l : \bar{B}_l | A) = 0$ for all l . Since quantum conditional mutual information cannot increase if systems are discarded and is non-negative [NC00] it follows that

$$0 = I(B_2 : B_1, B_3, \dots, B_k | A) = I(B_2 : B_3, \dots, B_k | A). \quad (4.46)$$

By Thm. 4.2, $I(B_1 : \bar{B}_1 | A) = 0$ implies a decomposition $\mathcal{H}_A = \bigoplus_i \mathcal{H}_{A_i^L} \otimes \mathcal{H}_{A_i^R}$ such that

$$\rho_{B_1 \dots B_k | A} = \sum_i \rho_{B_1 | A_i^L} \otimes \rho_{B_2 \dots B_k | A_i^R} \quad (4.47)$$

$$\Rightarrow \hat{\rho}_{B_2 \dots B_k | A} = \sum_i \frac{p_i}{d_{A_i^L}} \mathbb{1}_{A_i^L} \otimes \hat{\rho}_{B_2 \dots B_k | A_i^R} \quad (4.48)$$

where $p_i = d_{A_i^L} d_{A_i^R} / d_A$ form a probability distribution. The terms in $\hat{\rho}_{B_2 \dots B_k | A}$ only have support on orthogonal subspaces, so [NC00]

$$S(\hat{\rho}_{B_2 \dots B_k | A}) = H(\{p_i\}) + \sum_i p_i \log d_{A_i^L} + \sum_i p_i S(\hat{\rho}_{B_2 \dots B_k | A_i^R}), \quad (4.49)$$

$$S(\hat{\rho}_{B_2 | A}) = H(\{p_i\}) + \sum_i p_i \log d_{A_i^L} + \sum_i p_i S(\hat{\rho}_{B_2 | A_i^R}), \quad (4.50)$$

$$S(\hat{\rho}_{B_3 \dots B_k | A}) = H(\{p_i\}) + \sum_i p_i \log d_{A_i^L} + \sum_i p_i S(\hat{\rho}_{B_3 \dots B_k | A_i^R}), \quad (4.51)$$

$$S(\hat{\rho}_{\cdot | A}) = H(\{p_i\}) + \sum_i p_i \log d_{A_i^L} + \sum_i p_i S(\hat{\rho}_{\cdot | A_i^R}). \quad (4.52)$$

Substituting into

$$I(B_2 : B_3, \dots, B_k | A) \stackrel{\text{def}}{=} S(\hat{\rho}_{B_2|A}) + S(\hat{\rho}_{B_3 \dots B_k|A}) - S(\hat{\rho}_{B_2 \dots B_k|A}) - S(\hat{\rho}_{\cdot|A}) \quad (4.53)$$

the $H(\{p_i\})$ terms and the $\sum_i p_i \log d_{A_i^L}$ terms cancel, leaving

$$I(B_2 : B_3, \dots, B_k | A) = \sum_i p_i I(B_2 : B_3, \dots, B_k | A_i^R) = 0. \quad (4.54)$$

Therefore $I(B_2 : B_3 \dots B_k | A_i^R) = 0$ for each i by non-negativity and Thm. 4.2 can be applied to $\rho_{B_2 \dots B_k | A_i^R}$. Iterating this procedure yields the required decomposition. \square

Proof: (4) \Rightarrow (3). This also follows immediately from Thm. 4.2 by letting $B = B_l$ and $C = \bar{B}_l$ for each l . \square

Proof: (1) \Rightarrow (4). Assume that condition (1) holds. Note from Def. 4.4 that as there is no causal influence from L_l to \bar{B}_l for every l , there is also no causal influence from \bar{L}_l to B_l for every l . Therefore, by partitioning $B_1 \dots B_k$ into B_l and \bar{B}_l and partitioning $L_1 \dots L_k$ into L_l and \bar{L}_l , Thm. 4.2 implies that $I(B_l : \bar{B}_l | A) = 0$. Thus, since l is arbitrary, this implies (3) which further implies (4). \square

Proof: (4) \Rightarrow (1). Follows by a straightforward extension of the corresponding proof in Sec. 4.3.4. \square

Together, these steps can be used to show that any condition of Thm. 4.3 implies any other, thus completing the proof.

As well as being interesting in its own right, Thm. 4.3 also places further confidence in the quantum conditional independence defined here being correct. The fact that the simplest multipartite generalisation of Thm. 4.2 also holds is a testament to the naturalness of Thm. 4.2 and Def. 4.6.

4.4 Summary

This chapter started to consider the nature of ontological causal influences in quantum theory. A natural starting point is Reichenbach's principle, a cornerstone of causal reasoning but one that fails in certain quantum experiments. In particular, classical causal models fail to adequately explain Bell experiments [Bel87, WS15]. Since classical causal models generalise Reichenbach's principle, it is sensible to first find a natural quantum version of the principle before attempting to define a full framework of quantum causal models.

Finding a quantum Reichenbach's principle becomes easier if the original is split into two parts: qualitative and quantitative [Sec. 4.2.1]. The qualitative part can be commuted

to quantum theory with almost no change. The quantitative part applies specifically to the case of a complete common cause.

One way to derive the classical quantitative part is by assuming fundamentally deterministic dynamics [Sec. 4.2.2]. An analogous procedure was followed to justify a quantum quantitative part by assuming fundamental unitarity. This gave rise to Def. 4.6 of quantum conditional independence of B and C given A for a channel $\rho_{BC|A}$. Such quantum conditional independence gave a natural quantum quantitative part and therefore a full quantum Reichenbach’s principle [Sec. 4.3.2].

Classical conditional independence can be defined in many equivalent ways, each lending itself to different intuitions. It was then proposed that the same can be said for quantum conditional independence of a channel (Props. 4.1, 4.2). In Thm. 4.2, the four proposed definitions were shown to be equivalent. Each of these is a generalisation of a definition for classical conditional independence and each reduces to the corresponding classical case when states and channels appropriately decohere into some choice of “classical” bases.

One of these definitions, condition (4) of Thm. 4.2, is particularly interesting. It identifies a particular structure that conditionally independent channels must follow. In Sec. 4.3.5 it was shown that this structure has an operational interpretation and plays a similar role to the classical copying map. It appears that this structure generally allows two agents to “act independently” on a single input.

Theorem 4.3 generalised quantum conditional independence to the case of a channel with $k > 2$ outputs. It showed that the easy generalisations of the definitions in Thm. 4.2 hold in this more general scenario, giving extra confidence that those definitions are correct.

Some examples were considered in Sec. 4.3.6. After verifying that the proposed definitions act appropriately in some basic cases, particular attention was given to Bell experiments since they were a key motivation for developing quantum Reichenbach in the first place. It was shown that the outcomes of Bell experiments can be considered to arise from a quantum common cause, even though Bell’s theorem shows they cannot arise from a classical common cause [Bel87]. The conditional independence does depend on what is chosen to be conditioned on, but this is also the case for any classical conditionally independent channel.

Having motivated, defined, and tested quantum Reichenbach’s principle, the next step is to follow the classical example and generalise it to a full framework of quantum causal models. This will follow in the next chapter. A full discussion of these results, with open questions for further investigation, is deferred until Sec. 5.4 once the full framework of quantum causal models has been defined.

5

Quantum Causal Models

5.1 The Need for Quantum Causal Models

The previous chapter began to consider causal influences as an ontological issue in quantum theory. By looking at the simple scenario of a complete common cause it was discovered that a significant revision of the classical ways of describing causation was required. Such a revision was provided for that scenario in the form of a quantum Reichenbach's principle, utilising a notion of quantum conditional independence. In this chapter, these discussions are extended to more general causal scenarios. Much of the work presented in this chapter has been published in Ref. [ABH⁺17].

5.1.1 The Reality of Causal Influences

Reichenbach's principle is a powerful link between the probabilistic/statistical notion of correlation and the notion of a causal influence between two events. It supports the idea that causal influences are "real," even for those (such as probabilistic Bayesians) who prefer to think of probabilities as purely subjective.

But Reichenbach's principle only applies to some very simple causal scenarios, with two events of interest. A natural generalisation of the principle is found in the framework of classical causal models [Pea09, SGS01], which describes arbitrary causal structures and their relationship with possible probability distributions. Just as Reichenbach's principle can support the reality of causal influences in its limited realm of applicability, so do

classical causal models more generally¹.

A classical causal model, defined fully in Sec. 5.2, comes in two parts. First, the *causal structure* simply specifies what causal influences exist between events (or *nodes*). Simple causal structures have already been encountered in the previous chapter, e.g. Figs. 4.2, 4.9(b). Second, the probability distribution over these events compatible with the causal structure. The framework of classical causal models specifies which causal structures are permissible and which distributions are compatible with any given causal structure. Note the similarity to Reichenbach’s principle. The qualitative part simply demands the causal structure take some particular form, while the quantitative part ensures the distribution is compatible. This is no accident: Reichenbach’s principle becomes a corollary of the more general framework of causal models.

5.1.2 Causation in General Quantum Experiments

Given the discussion of the previous chapter, one should expect classical causal models to be insufficient for describing general quantum experiments. It has already been noted in Sec. 4.1.2 that no classical causal model can describe a general Bell experiment without fine-tuning. But Bell experiments are specifically designed to show that a simple classical common cause explanation fails. Are there other types of quantum experiments that also resist explanation by a classical causal model?

Several such experiments have been considered in the literature [GHZ89, BGP10, HWB11, BRGP12, Fri12, TSCA14]. It should perhaps be unsurprising that many are variations or generalisations of Bell experiment set-ups. In any case, the general point is clear: there are many conceivable quantum experiments that resist natural explanation by classical causal models².

Just as in the case of Bell experiments and Reichenbach’s principle, these experiments demand a revision of classical ideas of causality (each in their own way). Classical causal models are simply insufficient to adequately explain many quantum experiments. A satisfactory revision would ideally describe all such experiments without fine tuning, superluminal influences, or other similar undesirable features.

Beyond this, there is another lesson to take from the search for novel Bell-like experiments. Several techniques have been developed for deriving bounds on the observable

¹Another important result in this regard is the de Finetti theorem [Fin93, DF80, Fin75] (which also has quantum counterparts [BH13, BL09, CT09, CKMR07, CFS02c]). Amongst other things, this theorem provides subjective Bayesians justification for behaving *as if* the results of ancestrally independent trials are due to a single unknown probability distribution.

²Here “natural” might mean “in the absence of fine-tuning” or “in the absence of superluminal causal influences” or similar reasonable restrictions depending on the exact experiment considered.

statistics from particular classical causal models [LS15, WSF16, Cha16, RBB⁺16]. It has been noted that, given an analogous framework of quantum causal models, these techniques might be adapted to derive similar bounds for achievable quantum statistics from causal models with the same causal structures [CLG14, CMG15, CKBG15]. Then, by seeing where the quantum bounds differ from the corresponding classical bounds, one could identify experiments where quantum statistics violate classical causal models. A framework of quantum causal models would hopefully enable this to be done systematically.

Such a framework of quantum causal models will be presented in Sec. 5.3. This is found by generalising quantum Reichenbach as given in the previous chapter, just as classical causal models generalise Reichenbach’s principle.

The quantum causal models presented here are entirely natively quantum. They are not classical causal models with quantum additions. They do, however, reduce to the classical case in appropriate decohering limits [Sec. 5.3.4]. The framework builds on the foundation of quantum Reichenbach from Chap. 4 and contains it as a special case. These features strengthen the claim that this is an appropriate definition of quantum causal models and speak to a certain degree of robustness.

Use of the framework will be illustrated in Sec. 5.3.5 by examples. These include the case of a confounding common cause and simple Bayesian updating across a common cause. Finally, Sec. 5.4 will summarise the chapter and discuss the implications from both this chapter and the previous.

Just as in Chap. 4, all random variables and graphs will be assumed to be finite. Similarly, all quantum systems will be assumed to be finite-dimensional. This is done for clarity and no major conceptual changes needed to extend the results to infinite cases are anticipated.

5.1.3 Previous Approaches to Quantum Causal Models

This is certainly not the first time that moves in the direction of quantum causal models have been made³. This section will review some of this previous and related work to put the results of this chapter into their proper context.

An early approach to causality in quantum theory started small, with a pair of quantum systems B, C acted on by a single quantum operation. This led to defining the properties “(semi)-causal” and “(semi)-localisable” for such operations, as well as establishing relationships between these properties [BGNP01, ESW02, PHHH06]. This was extended in Ref. [SW05] to the tripartite case BAC with very strong results on the structure of

³That is probably the work of Refs. [Tuc95, Tuc12], where transition amplitudes are taken to replace conditional probabilities from classical causal models.

“local” unitaries. These results are related to the channels with “no causal influence” used in Secs. 4.3.2, 4.3.4.

References [LP08, LS13] made strides towards quantum causal models by viewing density operator quantum theory as a non-commutative generalisation of probability theory. There, density operators took the role of probability distributions and Choi-Jamiołkowski operators took the role of conditional distributions. This is very similar to the way quantum Reichenbach was approached in the previous chapter. Formally, there is a certain amount of overlap, with some results from Ref. [LP08] being used in proofs in Sec. 4.3.4. However, Refs. [LP08, LS13] primarily aimed to develop “causally neutral” quantum causal models that could be used for Bayesian inference, like classical causal models. This is rather different from traditional approaches to quantum theory, where spacelike and timelike relationships between systems are fundamentally different. This focus on quantum-theory-as-probability-theory led to causal models of a different character and the programme has been hampered by difficulties in defining quantum-state-like objects for timelike separated systems [HHP⁺16].

The programme of deriving Bell-like bounds in more general causal structures was mentioned in Sec. 5.1.2. This inspired independent formulations of quantum causal models in Ref. [HLP14] and Ref. [Fri16] with the aim of providing frameworks for deriving such bounds. Both approaches are based on some underlying operational theory: the arrows of the causal structure are taken to represent systems of that theory while the nodes (events) represent transformations. The operational framework of *general probabilistic theories* (GPTs, [Har01, Har11, CDP10, CDP11]) is the basis for Ref. [HLP14], while Ref. [Fri16] applies to theories based on symmetric monoidal categories. Both of these are more general operational theories that include classical and quantum as special cases. Reference [PB15] constructs specifically quantum causal models compatible with Ref. [Fri16].

Both frameworks provide some unification due to their generality. They also prove some novel results (notably, analogues to d-separation theorems in Refs. [HLP14, PB15]). However, neither approach is natively quantum. They only define conditional independence between observed classical outcomes, rather than between quantum systems as in Thm. 4.2. In particular, one cannot condition on a quantum node in those frameworks. The framework defined in this chapter, by contrast, will generalise quantum Reichenbach and will apply directly to quantum systems just as that principle did.

The approaches of Refs. [HLP14, Fri16, PB15] demonstrate that there are close links between quantum causal models and operational formulations of quantum theory. This is especially true of those operational formulations that consider relativistic causal structure, including the causaloid framework [Har09], the multi-time formalism [AV07, APTV09,

APT14, SGB⁺14], quantum combs [CDP09, Chi12, CDPV13], categorical quantum theory [CK17, CL13], and process matrices [OCB12, ACB14]. In particular, many of these frameworks [APTV09, CDP09, OCB12] use a pair of isomorphic Hilbert spaces to represent a system at a point in spacetime: an “input” system which is received then transformed to an “output” system. This has been particularly useful for describing local *interventions*, where the transformation from input to output may be freely chosen. This approach to interventions will be used in the nodes of quantum causal models later.

A similar interventional approach to nodes in causal models has also been considered classically [RR13], as noted in Ref. [CS16]. This was done to describe possible counterfactual interventions and similar properties make the quantum interventional approach useful here.

The previous work that most closely resembles the quantum causal models defined here is Ref. [CS16]. There, a definition for quantum causal models is given based on the process matrix formalism [OCB12, ACB14]. As in this chapter, the dual-Hilbert-space interventional method noted above is used for nodes in the causal structure. Nonetheless, there are important differences between quantum causal models as defined here and those of Ref. [CS16] which will be discussed in Sec. 5.4.

Perhaps the most important way that the quantum causal models defined in this chapter differ from all previous attempts is that they are generalised from a quantum Reichenbach’s principle. As demonstrated in the previous chapter, quantum Reichenbach is robust, general, and natively quantum. If this is the correct way to generalise Reichenbach’s principle to quantum systems, then one would expect that its natural generalisation to quantum causal models would be similarly correct. The philosophy of this chapter and the previous is also deliberately neutral and sticks to structures from vanilla quantum theory as far as possible. This has been done to allow the results to be as widely applicable as possible.

5.2 Classical Causal Models Generalise Reichenbach’s Principle

5.2.1 Classical Causal Models in Two Parts

Like Reichenbach’s principle, the framework of classical causal models [Pea09, SGS01] neatly divides into two parts. First, *causal structures* between random variables represented by directed acyclic graphs (DAGs). Second, the *Markov condition* which defines when a joint distribution over variables is compatible with any given causal structure.

Causal structures are simply a formal way to express causal relationships between events represented by random variables. A DAG consists of a set of nodes $\{X_i\}_i$ joined by arrows (a directed graph), such that it is impossible to travel from any node back to itself by following arrows (acyclic). It is convenient to use standard genealogical terminology to describe relationships between these nodes. For any node X_i , $\text{Pa}(X_i) = \text{Pa}(i)$ is the set of parent nodes for X_i . Similarly, $\text{Ch}(X_i) = \text{Ch}(i)$ is the set of children nodes of X_i and $\text{De}(X_i) = \text{De}(i)$ is the set of descendants of X_i (which conventionally includes X_i itself). Finally, $\text{Nd}(X_i) = \text{Nd}(i)$ is the set of “non-descendants” of X_i , viz. the complement of $\text{De}(i)$.

The Markov condition specifies which joint distributions over these variables are compatible with a given causal structure representing what actually occurred. A joint distribution $\mathbb{P}(\{X_i\}_i)$ is said to be *Markov for a given DAG* if and only if the nodes are those same random variables and the joint distribution can be written in the form

$$\mathbb{P}(\{X_i\}_i) = \prod_i \mathbb{P}(X_i | \text{Pa}(i)) \quad (5.1)$$

(recalling that each $\mathbb{P}(X_i | \text{Pa}(i))$ can be calculated from $\mathbb{P}(\{X_i\}_i)$ once the DAG is known).

This defines the *framework* of causal models. A *specific* causal model is given by: a set of random variables $\{X_i\}_i$ representing events, a DAG with nodes of those random variables, and a set of conditional probability distributions $\mathbb{P}(X_i | \text{Pa}(i))$. One could equivalently just specify a joint distribution that is Markov for the graph, but it is often more practical⁴ to give each conditional distribution which guarantees a Markov joint distribution can be constructed.

This framework generalises Reichenbach’s principle. If X and Y are ancestrally independent, then the Markov condition requires that $\mathbb{P}(Y, Z) = \mathbb{P}(Y)\mathbb{P}(Z)$, which is the qualitative part of Reichenbach’s principle. If there is a complete common cause X for Y and Z in the causal structure (e.g. Fig. 4.1), then the Markov condition guarantees that $\mathbb{P}(Y, Z | X) = \mathbb{P}(Y | X)\mathbb{P}(Z | X)$, which is the quantitative part. Moreover, a loose generalisation of Reichenbach’s principle for causal models may be stated as: correlations between variables should have causal explanations in the causal structure.

⁴More practical in two ways. First, when constructing a causal model for a given situation the individual conditional probabilities are often simply easier to discover or estimate. Second, it is more economical, as union of the spaces of different conditional distributions is normally much smaller than the space of joint distributions (you always get a Markov joint distribution by specifying conditional distributions, but not every joint distribution is Markov).

5.2.2 Justifying the Markov Condition

In Sec. 4.2.2, one possible justification for the quantitative part of Reichenbach’s principle from the qualitative part was presented. This was done by temporarily assuming fundamental determinism for illustrative purposes. A very similar argument can be used to justify the Markov condition [Eq. (5.1)] from the qualitative part of Reichenbach’s principle. This is strictly stronger than the argument in Sec. 4.2.2 since, as seen above, the Markov condition implies Reichenbach’s principle.

To this end suppose once again that classical dynamics is fundamentally deterministic. The task is to prove that if a situation has a causal structure given by some DAG, then the qualitative part of Reichenbach’s principle requires that the distribution over variables is Markov for that DAG. This argument closely mirrors that of Sec. 4.2.2 and so will be covered quickly.

Just as in Sec. 4.2.2, a determinist will always view the classical maps between nodes in the causal structure to be the result of *deterministic dilations*. That is, for every node X there must, in reality, also be some latent node Λ that is a cause for X in a more fundamental causal structure. This Λ ensures the classical channel that outputs X is the result of a deterministic dilation with input $\text{Pa}(X) \times \Lambda$. Simply put, any stochasticity in the original causal model is explained by functions including hidden latent variables in a more fundamental causal model.

Moreover, each latent variable must be unique to each node and have no parents. If this were not the case then adding the latent variables would invalidate the original causal structure by introducing new common cause links. In particular therefore, the latent variables are ancestrally independent and thus, by the qualitative part of Reichenbach’s principle, their distributions factorise.

Putting this together, one can say the following.

Definition 5.1. *A joint distribution $\mathbb{P}(\{X_i\}_i)$ is deterministically compatible with a causal structure given by DAG G if and only if:*

- (a) *the nodes of G are the variables $\{X_i\}$ and*
- (b) *there exists causal structure G' obtained from G by adding a node Λ_i and a single arrow $\Lambda_i \rightarrow X_i$ for each X_i such that*
- (c) *there exist distributions $\mathbb{P}(\Lambda_i)$ and functions⁵ $f_i : \text{Pa}(i) \times \Lambda_i \rightarrow X_i$ that form deterministic dilations for each channel $\mathbb{P}(X_i | \text{Pa}(X_i))$ derived from the joint distribution.*

⁵Here $\text{Pa}(i)$ is being used to refer to the parents of X_i in G . Of course, in G' each Λ_i is a parent for the corresponding X_i .

Just as in Sec. 4.2.2, this mathematics is independent from the temporary assumption of determinism used to reach it. By simply entertaining the logical possibility of a deterministic dilations, the following theorem from Ref. [Pea09] completes the justification of the Markov condition.

Theorem 5.1 (Ref. [Pea09]). *Given a joint distribution $\mathbb{P}(\{X_i\}_i)$ and a causal structure the following are equivalent:*

1. $\mathbb{P}(\{X_i\}_i)$ is *deterministically compatible with the causal structure.*
2. $\mathbb{P}(\{X_i\}_i)$ is *Markov for the causal structure, satisfying Eq. (5.1).*

Theorem 5.1 is clearly analogous to Thm. 4.1. It also admits similar readings: both as the justification presented above and as proving that Eq. (5.1) and Def. 5.1 are equivalent definitions of the Markov condition. As in Sec. 4.2.2, this depends on whether condition (1) is taken as a causal or probabilistic statement respectively.

Once again, this argument is not meant as a complete first-principles derivation of classical causal models. It is instead simply illustrative and helps to motivate the approach to quantum causal models that follows.

5.3 Quantum Causal Models Generalise Quantum Reichenbach

In the previous section, it was seen that the framework of classical causal models can be seen as a generalisation of Reichenbach’s principle to more general causal structures. Moreover, one way to justify the crucial Markov condition of classical causal models was shown by assuming fundamental determinism (mirroring the justification of the quantitative part of Reichenbach’s principle in Sec. 4.2.2).

This will now be used to motivate a definition for quantum causal models. In particular, quantum causal models should generalise both quantum Reichenbach and classical causal models. The definition of quantum causal models that follows naturally achieves both of these things, putting it in a strong position.

5.3.1 Defining Quantum Causal Models

In Secs. 4.2.2, 4.3.2, and 5.2.2 justifications were given for Reichenbach’s principle, quantum Reichenbach, and classical causal models respectively. Each of these proceeded by temporarily assuming either fundamental determinism (in the classical cases) or fundamental unitarity (in the quantum case).

It would be neat, therefore, to continue that pattern here. That is, to justify a definition for quantum causal models by assuming fundamental unitarity. Unfortunately, this is difficult to do *a priori* since it is unclear what sort of mathematical object a quantum analogue for the Markov condition should apply to. For Reichenbach’s principle, the quantitative part applies to the classical channel. Correspondingly, for quantum Reichenbach the quantitative part applies to the quantum channel. For classical causal models, the Markov condition applies to the joint probability distribution. Here lies the problem. As already noted, there is no convenient or accepted quantum analogue for a joint probability distribution for causally-related systems [HHP⁺16].

The approach taken here, therefore, is to propose a definition that most simply generalises quantum Reichenbach to general causal structures. The resulting definition can then be checked to satisfy other natural criteria.

As with the qualitative part of Reichenbach’s principle, the causal structures from classical causal models can be commuted to quantum theory with almost no change. This is because a causal structure is, mathematically, just a DAG. All that remains is to specify what the nodes should correspond to.

Mathematically, the nodes of a classical causal structure are taken to be random variables. Physically, they can be thought of as events/systems in local regions of spacetime which can take several values/physical states. In quantum theory, the most general way to describe the same thing is with a quantum instrument applied to a system. This covers both the interpretation as an event with outcome value and as a system with a state.

Therefore, each node A_i of a quantum causal structure will correspond to two Hilbert spaces: an *input space* $\mathcal{H}_{A_i}^{\text{in}} = \mathcal{H}_i^{\text{in}}$ and an *output space* $\mathcal{H}_{A_i}^{\text{out}} \stackrel{\text{def}}{=} \mathcal{H}_i^{\text{out}} \stackrel{\text{def}}{=} (\mathcal{H}_i^{\text{in}})^*$ which is its dual. The fact that the output Hilbert space is the dual to the input Hilbert space is simply a convenient mathematical convention that will be useful later, morally they can be thought of as “the same” spaces. This allows the interpretation that each node A contains the local laboratory of some agent, who performs some quantum instrument with input from $\mathcal{H}_i^{\text{in}}$ and output to $\mathcal{H}_i^{\text{out}}$. It also allows for less general interpretations, such as each node simply representing the state of some system when the quantum instrument is just an identity map.

As noted in Sec. 5.1.3, this interventional approach to quantum nodes has been used many times before in the literature. It is well-known as a general and convenient way to allow for both agent interventions and measurement outcomes where they might be needed.

The first part of a quantum causal model is therefore to specify a causal structure as a DAG and interpret the nodes as pairs of input/output Hilbert spaces as above. To complete

the causal model, by analogy with a classical causal model, channels between the nodes are required.

The simplest way to generalise quantum Reichenbach to these causal structures is to specify, for each node A_i , a quantum channel $\rho_{A_i|\text{Pa}(i)}$ from $\bigotimes_{A_j \in \text{Pa}(i)} \mathcal{H}_j^{\text{out}}$ to $\mathcal{H}_i^{\text{in}}$ (using the notational convention from Sec. 4.3.1). Then, to ensure that these channels are compatible, simply require that they commute pairwise. That is, for all nodes A_i and A_j , require $[\rho_{A_i|\text{Pa}(i)}, \rho_{A_j|\text{Pa}(j)}] = 0$. As such, a complete channel from common parents to their children can be taken to be the product of each of these channels $\rho_{A_i|\text{Pa}(i)} \rho_{A_j|\text{Pa}(j)} \cdots$. The fact that each channel commutes with every other ensures that such a product still defines a valid quantum channel via the Choi-Jamiołkowski isomorphism.

It is convenient to wrap all of these channels into a single object. This can easily be done by taking their product. The result is the *model state*

$$\sigma \stackrel{\text{def}}{=} \prod_i \rho_{A_i|\text{Pa}(i)} \quad (5.2)$$

which is an operator on $\bigotimes_i (\mathcal{H}_i^{\text{in}} \otimes \mathcal{H}_i^{\text{out}})$ where, recall, each factor commutes with every other.

This completes the definition of a quantum causal model. It consists of a causal structure with local laboratories for the nodes $\{A_i\}_i$ and a set of channels $\rho_{A_i|\text{Pa}(i)}$ which pairwise commute. These channels together define a model state by Eq. (5.2).

Equivalently, a quantum causal model could be taken as a causal structure together with a model state σ over $\bigotimes_i (\mathcal{H}_i^{\text{in}} \otimes \mathcal{H}_i^{\text{out}})$ which is required to satisfy Eq. (5.2). This way, Eq. (5.2) is analogous to the Markov condition and is therefore called the *quantum Markov condition*.

Crucially, this definition generalises quantum Reichenbach. If the causal network is taken to be a complete common cause from A to BC [Fig. 4.3], then the quantum Markov condition requires that the channel factorises as $\rho_{BC|A} = \rho_{B|A} \rho_{C|A}$. This is the reason for claiming that this definition is the simplest to generalise quantum Reichenbach. Moreover, it also generalises classical causal models, as shall be shown explicitly in Sec. 5.3.4

More work is required to flesh out the meaning of these quantum causal models and how they should be used, but first some remarks should be made about the relationship between this approach and traditional classical causal models.

5.3.2 Bayesian vs. do-conditionals

Strictly speaking, an important distinction should be made between two types of conditional in classical causal models, discussed at length in Ref. [Pea09]. The *Bayesian conditional*

is derived from a joint probability distribution by Bayes' rule $\mathbb{P}(Y|X) \stackrel{\text{def}}{=} \mathbb{P}(Y, X)/\mathbb{P}(X)$. This represents the set of probability distributions that should be assigned to Y given that X is found to have any given value. The *do-conditional* $\mathbb{P}(Y|\text{do } X)$, on the other hand, gives the set of probability distributions for Y given that an agent has intervened and set X to any definite value. Rather than following Bayes' rule, this is found by modifying the causal model to remove any arrows entering X and setting $\mathbb{P}(X)$ to the appropriate point distribution for the specifically set value.

A classical channel from X to Y should strictly, therefore, be written $\mathbb{P}(Y|\text{do } X)$ and similarly quantum channels $\rho_{B|A}$ are more analogous to do- than Bayesian conditionals. This relates to the problem with joint states over causal structure in quantum theory discussed above [HHP⁺16].

Since Bayesian conditionals are derived from a joint distribution, finding a close quantum analogue for the Bayesian conditional seems difficult at best (finding a joint state over the network is not straightforward). For this reason, no such analogue is given here. Rather, the quantum Markov condition has been given in terms of quantum channels $\rho_{A|\text{Pa}(A)}$, which do not require such a joint state but are more analogous to do-conditionals.

Fortunately, this need not break the analogy between quantum and classical causal models. If a probability distribution $\mathbb{P}(\{X_i\}_i)$ is Markov for a causal network, then for each node $\mathbb{P}(X_i|\text{Pa}(i)) = \mathbb{P}(X_i|\text{do } \text{Pa}(i))$ [Pea09]. That is, the Bayesian conditionals match the form of the classical channels in the network. Therefore one can equivalently define classical causal models as causal structures supplemented with do-conditionals $\mathbb{P}(X_i|\text{do } \text{Pa}(i))$ for each node. It is this definition to which quantum causal models are more closely analogous.

These distinctions have not and will not be used in the bulk of this thesis for brevity. Reference [CS16] discusses similar issues in the context of its related approach to causal models.

Of course the other main difference between how quantum causal models and classical causal models have been defined here is that quantum causal models use an interventionist approach to nodes, as noted in Sec. 5.1.3.

5.3.3 Making Predictions

A quantum causal model would not be much use if it did not give predictions for measurements. Fortunately, once a quantum instrument has been specified at each node joint probabilities for their outcomes can be found.

For a quantum causal model with nodes $\{A_i\}_i$, let $\{\mathcal{E}_i^{k_i}\}_{k_i}$ be the quantum instrument at node A_i , expressed as a set of completely positive maps from $\mathcal{H}_i^{\text{out}}$ to itself. Index k_i

labels the classical outcomes of the intervention in that local laboratory. An equivalent way to define these instruments is with their Choi-Jamiołkowski isomorphic states on $\mathcal{H}_i^{\text{out}} \otimes \mathcal{H}_i^{\text{in}}$

$$\tau_i^{k_i} = \tau_{A_i}^{k_i} \stackrel{\text{def}}{=} \sum_{p,q} \mathcal{E}_i^{k_i} (|p\rangle_{\text{out}}\langle q|) \otimes |p\rangle_{\text{in}}\langle q| \quad (5.3)$$

where $\{|p\rangle_{\text{in}}\}_p$ is an arbitrary orthonormal basis on $\mathcal{H}_i^{\text{in}}$ and $\{|p\rangle_{\text{out}}\}_p$ its dual basis in $\mathcal{H}_i^{\text{out}}$.

In the special case where the instrument is the identity (*i.e.* there is no intervention) then this state is the *linking operator* used in Sec. 4.3.1

$$\tau_i^{\text{id}} = \tau_{A_i}^{\text{id}} \stackrel{\text{def}}{=} \sum_{p,q} |p\rangle_{\text{out}}\langle q| \otimes |p\rangle_{\text{in}}\langle q|. \quad (5.4)$$

With this notation, the joint probability for any set of outcomes $\{k_i\}_i$ given choices of quantum instruments is given by

$$\mathbb{P}(\{k_i\}_i) = \text{Tr} \left(\sigma \left(\tau_1^{k_1} \otimes \tau_2^{k_2} \otimes \dots \right) \right). \quad (5.5)$$

It is easy to verify that this produces the same probabilities as standard quantum theory.

The linking operator allows one to remove or ignore nodes from a causal network. This process is called *linking out* and corresponds to the classical marginalisation process. For example, suppose node A is ignored. This can be represented by linking it out from the model state

$$\text{Ln}_A \sigma \stackrel{\text{def}}{=} \text{Tr}_{A^{\text{in}} A^{\text{out}}} \left(\sigma \tau_A^{\text{id}} \right). \quad (5.6)$$

The linking operation Ln may be thought of as a modified partial trace, which includes insertion of the linking operator. Doing this produces a new model state which can be used to obtain outcome probabilities from the other nodes as above.

There is no general reason to expect $\text{Ln}_A \sigma$ to satisfy the quantum Markov condition for an appropriate causal structure even when σ did for the original causal structure. However, in the special case where A has no children, then it is easy to verify that $\text{Ln}_A \sigma$ does satisfy the quantum Markov condition for the causal structure obtained by removing node A .

5.3.4 Classical Limits

As noted above, it is crucial that any reasonable definition of quantum causal networks be a generalisation of classical causal networks. It shall now be shown that this is the case for quantum causal models as defined in Sec. 5.3.1.

Quantum theory describes a classical situation in a decohering limit. That is, when there is a choice of orthonormal basis for each system such that all states and channels are

diagonal with respect to those bases. In such a case, the elements of those bases become the classical states, the quantum states probability distributions over them, and the channels become classical channels.

Consider any decohering limit for a quantum causal network and let all of the quantum instruments be the identity. By taking the product of the model state σ with the linking operators and tracing out the input spaces of each node, a state diagonal in these bases is obtained

$$\varsigma \stackrel{\text{def}}{=} \text{Tr}_{A_1^{\text{in}} A_2^{\text{in}} \dots} \left(\sigma \left(\tau_1^{\text{id}} \otimes \tau_2^{\text{id}} \otimes \dots \right) \right). \quad (5.7)$$

This is a positive trace-one operator over $\mathcal{H}_{A_1}^{\text{out}} \otimes \mathcal{H}_{A_2}^{\text{out}} \otimes \dots$.

The diagonal entries of ς (in the decohering bases) form a probability distribution over the decohered classical states of each $\mathcal{H}_{A_i}^{\text{out}}$. The claim is that if σ satisfies the quantum Markov condition [Eq. (5.2)] in such a classical limit, then the distributions encoded in ς satisfy the classical Markov condition for the same causal structure. This will now be shown.

Suppose σ satisfies Eq. (5.2). By writing σ as a product where each factor $\rho_{A|\text{Pa}(A)}$ appears before the corresponding factor for any of its children, then each operator τ_A^{id} can be commuted through to appear to the immediate right of the corresponding $\rho_{A|\text{Pa}(A)}$. The result is a product with one factor

$$\text{Tr}_{A^{\text{in}}} \left(\rho_{A|\text{Pa}(A)} \tau_A^{\text{id}} \right) \quad (5.8)$$

for each node A , in any order where parents appear before children.

Suppose the linking operator is expanded in terms of the decohering basis $\{|a\rangle\}_a$ for A , so that Eq. (5.8) takes the form

$$\sum_{a,b} A^{\text{in}} \langle b | \rho_{A|\text{Pa}(A)} | a \rangle_{A^{\text{in}}} \otimes | a \rangle_{A^{\text{out}}} \langle b |. \quad (5.9)$$

Recalling that $\rho_{A|\text{Pa}(A)}$ is diagonal in the decohering bases, all terms vanish except where $b = a$. Letting $\{|\bar{p}\rangle\}_p$ be the product basis of the decohering bases for each node in $\text{Pa}(A)$, the operator $\rho_{A|\text{Pa}(A)}$ can be expanded to find

$$\sum_{a,p} \left\{ \langle a | \mathcal{E}_{A|\text{Pa}(A)} \left(|\bar{p}\rangle_{\text{Pa}(A)^{\text{in}}} \langle \bar{p}| \right) | a \rangle \right\} | a \rangle_{A^{\text{out}}} \langle a | \otimes |\bar{p}\rangle_{\text{Pa}(A)^{\text{out}}} \langle \bar{p}|. \quad (5.10)$$

Here the fact that $\mathcal{E}_{A|\text{Pa}(A)}$ only produces non-zero output when the input is diagonal in the $\{|\bar{p}\rangle\}_p$ basis (due to decoherence in the classical bases) has been used. Finally, note that the factor in braces is simply the probability that the classical values of $\text{Pa}(A)$ represented by $|\bar{p}\rangle$ are mapped by $\mathcal{E}_{A|\text{Pa}(A)}$ to the classical value $|a\rangle$ of A . That is, $\langle a | \mathcal{E}_{A|\text{Pa}(A)} (|\bar{p}\rangle \langle \bar{p}|) | a \rangle = \mathbb{P}(A = a | \text{Pa}(A) = p)$.




Figure 5.1: Causal structures where B causes C , but they also have additional common causes. (a) A is the only common cause (though not a complete common cause due to the arrow from B to C). (b) D is an additional common cause.

This shows that ς is a product of factors, each diagonal in the product basis of the decohering bases

$$\varsigma = \prod_A \left\{ \sum_{a,p} \mathbb{P}(A = a | \text{Pa}(A) = p) |a\rangle_{A^{\text{out}}} \langle a| \otimes |\bar{p}\rangle_{\text{Pa}(A)^{\text{out}}} \langle \bar{p}| \right\}. \quad (5.11)$$

Each diagonal entry in this basis is a product of conditional probability distributions $\mathbb{P}(A = a | \text{Pa}(A) = p)$, which forms a joint distribution for the variables to take values a, p, \dots etc. Thus, the entries are joint distributions over the graph which are Markov for the graph, as claimed.

In this way, quantum causal models naturally reduce to classical causal models in the appropriate limit with no interventions.

5.3.5 Examples

Quantum causal models have now been defined, shown to predict measurement outcomes, and shown to reproduce classical causal models in the relevant limits. Since quantum causal models generalise quantum Reichenbach, the examples of Sec. 4.3.6 are also examples for quantum causal models. What follows are further examples to illustrate the properties of quantum causal models.

5.3.5.1 Confounding Common Cause

Consider a quantum causal model with causal structure shown in Fig. 5.1(a). Such a model requires channels $\rho_{C|AB}$, $\rho_{B|A}$ and ρ_A to be specified, which must commute pairwise. This produces a model state $\sigma = \rho_{C|AB} \rho_{B|A} \rho_A$ on $\mathcal{H}_C^{\text{out}} \otimes \mathcal{H}_C^{\text{in}} \otimes \mathcal{H}_B^{\text{out}} \otimes \mathcal{H}_B^{\text{in}} \otimes \mathcal{H}_A^{\text{out}} \otimes \mathcal{H}_A^{\text{in}}$.

First, note some properties of these states. Since A has no parents, its “channel” is from the trivial system, which is simply a quantum state ρ_A on $\mathcal{H}_A^{\text{in}}$. Some of the commutation

relations, such as $[\rho_{B|A}, \rho_A] = 0$, follow immediately, since they only act non-trivially on different Hilbert spaces. However, the requirement that $\rho_{C|BA}$ and $\rho_{B|A}$ commute is a significant restriction, as they both act on $\mathcal{H}_A^{\text{out}}$. By Thm. 4.2, this requires that there is a decomposition $\mathcal{H}_A^{\text{out}} = \bigoplus_i \mathcal{H}_{A_i^L}^{\text{out}} \otimes \mathcal{H}_{A_i^R}^{\text{out}}$ where $\rho_{C|BA}$ acts only on the right-hand factor in each subspace and $\rho_{B|A}$ only on the left hand factors.

This is very different from *classical* causal models of the same causal structure. It is easy to check that any classical distribution $\mathbb{P}(A, B, C)$ is Markov for Fig. 5.1(a) interpreted as a classical causal structure. There is no analogue of the strong constraint that $[\rho_{C|BA}, \rho_{B|A}] = 0$ puts on the channels possible in a quantum causal network. Classically, both B and C may depend arbitrarily on A , whereas in a quantum causal network these dependencies must satisfy a compatibility condition required by commutation.

Figure 5.1(b) shows a variation of this causal network, where D is an additional common cause. This might represent, for example, a system interacting with an environment. A , B , and C are the system at different times and D is the initial state of the environment. The causal arrow from A to C is necessary since information can flow into the environment from A and back to the system at C (nodes for the environment at later times have been omitted from the diagram).

Clearly, if nodes A and D are considered together then this is the exact same form as Fig. 5.1(a). Therefore, the Hilbert space $\mathcal{H}_A^{\text{out}} \otimes \mathcal{H}_D^{\text{out}}$ must factorise in linear subspaces as above, with the left hand factor going to B and the right hand factor to C . This *does not* imply that each individual space $\mathcal{H}_A^{\text{out}}$ and $\mathcal{H}_D^{\text{out}}$ must decompose in this way, only that the tensor product of the two does. Suppose that ρ_D is fixed in some pure state $|0\rangle\langle 0|$ and consider removing D from the structure (returning to Fig. 5.1(a)). There is no reason to expect that the resulting model state, found by Eq. (5.6), satisfies the quantum Markov condition for Fig. 5.1(a). This is because the channels out of $\mathcal{H}_A^{\text{out}}$ will not generally have the required decomposition.

Once again, this is in marked contrast with the case of Fig. 5.1(b) interpreted as a *classical* causal structure. Again, there is no restriction on the form of the classical channels out of A or D (or both together). Moreover, if D is fixed to have some definite value and then is marginalised out of the causal model, the resulting distribution will always be Markov for Fig. 5.1(a). Of course, it must be, since every distribution over A , B , and C is Markov for Fig. 5.1(a).

One way to understand this to consider pure quantum states to contain some irreducible stochasticity. For example, one may prefer to think of them as “maximal but incomplete information” about a system (see also Sec. 4.3.6.2). In this view, the idea that a pure state of D could still underwrite correlations between B and C is a natural one.




Figure 5.2: Causal structures for simple complete common causes in (a) the classical case and (b) the quantum case.

These examples demonstrate a curious thing: that quantum causal models appear to place *more* restrictions on the same causal structure than their classical counterparts. Of course, quantum theory and quantum causal models contain classical theory and classical causal models respectively, so there is no contradiction. However, this shows that causal structure tempers some of the extra power of quantum theory in a way that it does not classically.

5.3.5.2 Simple Bayesian Updating with a Complete Common Cause

Consider again the case of X being a complete classical common cause for Y and Z , illustrated in Fig. 5.2(a). Classically, an important property of a complete common cause is that any knowledge about Z gained at Y is expressible via X . That is, if an agent learns something about Y , then the derived knowledge they gain about Z is entirely the result of derived knowledge they have learned about X . The new information “follows the arrows” of the causal model (both backwards then forwards). Indeed, properly formalised, this can be viewed as another definition for classical conditional independence.

Specifically, the causal model specifies channels $\mathbb{P}(Y|X)$ and $\mathbb{P}(Z|X)$ together with distribution $\mathbb{P}(X)$. Suppose that an agent learns that $Y = y$. This allows them to replace the initial distribution $\mathbb{P}(X)$ with $\tilde{\mathbb{P}}(X) \stackrel{\text{def}}{=} \mathbb{P}(X|Y = y)$. Then passing this new distribution through the channel $\mathbb{P}(Z|X)$ they find

$$\mathbb{P}(Z|Y = y) = \sum_{x \in X} \mathbb{P}(Z|X = x) \tilde{\mathbb{P}}(X = x). \quad (5.12)$$

The result is exactly the same as if the agent had simply calculated $\mathbb{P}(Z|Y)$ from $\mathbb{P}(Y, Z, X)$ using Bayes’ rule. For this procedure to work generally, it is essential that Y and Z are conditionally independent given X .

This is a natural and intuitive property of a complete common cause: it underwrites all correlations and therefore can be used to mediate information gained. It is therefore desirable, perhaps even essential, for the same procedure to work in quantum causal models.

For the corresponding quantum causal model of Fig. 5.2(b), suppose that an agent makes some measurement at B , obtaining outcome k_B . This corresponds to the operator $\tau_B^{k_B}$ from a quantum instrument, as in Eq. (5.3). It is well known from standard quantum theory that the agent can consider the system at C to be in an updated state $\sigma_{C|k_B}$ based on this measurement outcome due to “collapse” upon measurement. So long as A is a complete quantum common cause for B and C , this can be expressed as the result of a modified input state $\tilde{\rho}_A$ being passed into the channel $\rho_{BC|A}$ and then marginalising over B . This mirrors exactly the classical updating procedure above.

The updated state at A required to do this is

$$\tilde{\rho}_A \stackrel{\text{def}}{=} \frac{\rho_A \text{Tr}_{B^{\text{in}} B^{\text{out}}} \left(\tau_B^{k_B} \rho_{B|A} \right)}{\text{Tr} \left(\rho_A \tau_B^{k_B} \rho_{B|A} \right)}. \quad (5.13)$$

To verify that this state behaves as required, simply apply the channel $\rho_{BC|A} = \rho_{B|A} \rho_{C|A}$ and trace out the output at B^{in} . The result is equal to (up to normalisation)

$$\sigma_{C|k_B} = \text{Tr}_{B^{\text{out}}} \left((\mathcal{E}_B^{k_B} \otimes \text{id}_C)(\mathcal{E}_{BC|A}(\rho_A)) \right) \quad (5.14)$$

which is exactly the marginal state on C after the agent finds the outcome k_B as predicted by standard quantum theory. In other words, rather than taking the joint state on BC and measuring B , the agent can find the marginal at C after measurement by inputting $\tilde{\rho}_A$ into the channel $\rho_{BC|A}$. Just as in the classical case, to make this work conditional independence of the channel $\rho_{BC|A}$ was necessary. Also similarly to the classical case, all that is needed to find $\tilde{\rho}_A$ is the original state on A , the channel $\rho_{B|A}$ and the outcome k_B . The agent can express $\tilde{\rho}_A$ in complete ignorance of C or the channel to it.

Consider, for example, the coherent copy channel discussed in Sec. 4.3.6.2. There it was seen that this channel, summarised in Eq. (4.42), does not satisfy conditional independence, which initially seems surprising. However, this updating procedure provides another way of understanding why it should not.

Suppose the input to the coherent copy [Eq. (4.42)] is $|+\rangle_A = (|0\rangle + |1\rangle)/\sqrt{2}$. The output is therefore the Bell state $|\Phi^+\rangle_{BC} = (|00\rangle + |11\rangle)/\sqrt{2}$. If a measurement were made at B and outcome $|+\rangle_B$ obtained then the resulting collapsed marginal at C would become $|+\rangle_C$. However, there is no state which can be input to that channel to produce $|+\rangle_C$ as the marginal output at C . Therefore, $\tilde{\rho}_A$ doesn’t exist for this channel and measurement. This is in violation of the discussion above, where it was seen that every common cause channel and subsequent measurement can be understood in this way.

This whole procedure of knowledge “following the arrows” is reminiscent of Bayesian updating. It is, however, much more limited in scope. This discussion only directly applies

to complete common causes. Recall also the disanalogy between quantum causal models and Bayesian conditionals noted in Sec. 5.3.2.

5.4 Summary and Discussion

This chapter completes the investigation of ontological causal influences in quantum theory in this thesis. Following the classical example, this has been done by generalising quantum Reichenbach of Chap. 4 to a full framework of quantum causal models. Since quantum Reichenbach is well motivated, the resulting framework is placed on a solid foundation.

In order to respect the great success of the theory of classical causal models, any framework of quantum causal models should also be a generalisation of that theory. To this end, Sec. 5.2 briefly outlined the theory of classical causal models and demonstrated one justification for them, mirroring the justification for Reichenbach’s principle in Sec. 4.2.2.

Just as with Reichenbach’s principle, finding a quantum generalisation of classical causal models becomes easier when the original is bisected. The causal part, causal structures, can be used in quantum theory with little change. All that is needed is a specification of what the nodes correspond to. The real challenge is finding a generalisation of the probabilistic/statistical part—the Markov condition Eq. (5.1).

To this end, the definition of quantum causal models in Sec. 5.3.1 sought to most simply generalise quantum Reichenbach to general causal networks. Each node in the causal structure was modelled as a local laboratory, with input system, output system, and the possibility of some agent who intervenes there. This allowed the nodes to be thought of as events, observations, or system states as needed. It also facilitated a quantum Markov condition that applies to a global *model state* [Eq. (5.2)] that is largely disconnected, avoiding the problems with defining true joint quantum states over time [HHP⁺16].

To further justify this somewhat speculative definition, Sec. 5.3.3 demonstrated how it can be used to obtain predictions for experiments in line with standard quantum theory. Further, Sec. 5.3.4 explicitly showed how these quantum causal models generalise classical causal models. This was followed by some examples in Sec. 5.3.5, which touch on the power of quantum causal models for explaining phenomena and highlight some key differences with classical causal models.

As noted in Sec. 5.1.3, this is not the first time that a framework of quantum causal models has been defined. The framework given here, however, has two key advantages over previous attempts. First, it is a natively quantum framework. All classical systems described by the network are quantum systems in some classical decohering limit. This is a great strength over frameworks which treat quantum systems as add-ons to classical

causal models. Here, quantum systems are the only first-class citizens. Second, it rests on a thoroughly motivated quantum Reichenbach’s principle. So, at least in complete common cause scenarios, the causal models inherit this motivation and have many close analogies to the classical case.

In particular, the important differences between quantum causal models given here and those of Ref. [CS16] should be noted. Both use a similar set-up, with input and output spaces at each node with possible interventions between them. The important difference is in how these nodes can be linked. In Ref. [CS16], the output space from one node must factorise as a tensor product of spaces for each outgoing causal arrow. This is much stronger than the condition here, which requires that output Hilbert spaces must factorise *in linear subspaces* (which may be one-dimensional) *when they are a complete common cause*. Compare, for example, the causal models of Sec. 5.3.5.1 where this thesis only requires the combined AD system to factorise in this way, while Ref. [CS16] would require A and D to individually factorise as simple tensor products.

The definition of quantum Reichenbach provided in Sec. 4.3 rests on a concept of quantum conditional independence of outputs B and C given input A , which has four equivalent definitions given in Thm. 4.2. A corresponding set of definitions of conditional independence of $k > 2$ outputs given one input are given in Thm. 4.3. Consider again each of these defining conditions.

Condition (1) fits perfectly with the idea that all quantum dynamics is fundamentally unitary, a popular position known as the “church of the larger Hilbert space”. It defines exactly which unitaries can be considered to act as complete common cause channels (satisfying conditional independence of outputs given input). In such a view, therefore, quantum conditional independence provides a strong characterisation of the structure of ontological transformations. Moreover, this definition of conditional independence can be found just by assuming a larger Hilbert space position, as in Sec. 4.3.2.

Condition (2) is perhaps the most pleasingly simple, but is not directly terribly useful. Not only is the Choi-Jamiołkowski state $\rho_{BC|A}$ a somewhat roundabout description of the channel, but the product of two such states has no direct physical meaning in general. The great exception, of course, is in a decohering classical limit, when it becomes the familiar condition for classical conditional independence.

Condition (3) is a very direct generalisation of the corresponding classical condition. To obtain the state $\hat{\rho}_{BC|A}$ as a simple joint quantum state, one can prepare a maximally entangled input $|\Phi\rangle = \sum_i |i\rangle_A |i\rangle_A$ for any basis $\{|i\rangle\}_i$ of \mathcal{H}_A and input half of it into the channel $\rho_{BC|A}$. The joint output state, including the extra A system, will be mathematically equal to $\hat{\rho}_{BC|A}$. This closely mirrors how one can obtain the corresponding classical

distribution $\hat{\mathbb{P}}(Y, Z, X)$ as discussed in Sec. 4.3.3. Condition (3) also provides an interesting link to the study of approximate quantum Markov chains [ILW08, FR15, Wil15, STH16, LW14, JRS⁺15, BLW15, SFR16]. Typically, (approximate) quantum Markov chains are defined for a tripartite joint quantum state with (approximately) vanishing conditional mutual information. In the approximate case, questions of how characterise quantum states with small-but-non-zero conditional mutual information are of great importance. Such results should also be very important for the study of quantum causal models, especially if one wants to reason about causal structures from imperfect experimental data.

Condition (4) is perhaps the most curious, but is also very persuasive. As discussed in Sec. 4.3.5, it seems to directly capture the idea that such channels are those allowing two agents to independently act on a single input. A general procedure for this is to perform a von Neumann measurement on the input and factorise the resulting state in a manner that can depend on the measurement result. Equivalently, the channel factorises into two, potentially differently in different linear subspaces. The factorisation in linear subspaces is analogous to the classical copy operation, in so far as it allows two agents to act independently on a single input. This is an interesting new analogue for the copy operation. In Sec. 4.3.5, these operations were represented in circuit diagrams with the new symbol $\textcircled{\otimes}$.

The use of $\textcircled{\otimes}$ in circuit diagrams immediately suggests extending circuit formulations of quantum theory to natively support such structures. In particular, it would be interesting to see how $\textcircled{\otimes}$ could be incorporated into the diagrammatic reformulation of quantum theory of Refs. [Coe10, CK17] based on category-theoretical underpinnings.

Perhaps, if an operational formalisation of the concept of two agents “independently” acting on a single input were found, one could derive condition (4) as a consequence. This would provide another strong justification for quantum Reichenbach as defined here, as well as being an interesting result in its own right.

Classical causal networks have become invaluable in the study of statistics. In particular, the field of causal inference [Pea09, SGS01, LS15, WSF16, CLM⁺14], which aims to deduce properties of the causal structure from uncontrolled statistical data. This has significant applications in many areas of science and beyond. The quantum causal models presented here should provide the appropriate framework for achieving similar utility for quantum experiments.

As noted in Sec. 5.1.2, classical causal models have provided techniques for deriving Bell-like inequalities for more general causal structures [LS15, WSF16, Cha16, RBB⁺16]. The quantum causal models presented here should therefore prove useful in adapting these

techniques to finding bounds on observables in quantum experiments with the same causal structures [CLG14, CMG15, CKBG15].

There has been much recent interest in the idea of “indefinite causal structures” in quantum theory. That is, considering the possibility of a coherent superposition of different causal structures [Chi12, OCB12, MRSR16, FB16]. This may be significant for the project of unifying quantum and general relativistic theories [Har09]. One might expect the framework of quantum causal models defined here, as well as the understanding of common causes, to be able to significantly contribute to this conversation.

Finally, there is a notable way in which the development of the quantum causal models presented in Sec. 5.3.1 could be improved. Section 4.2.2 presented a justification of Reichenbach’s principle from fundamental determinism and Sec. 5.2.2 did the same for classical causal models. By way of analogy, Sec. 4.3.2 presented a justification for quantum Reichenbach from fundamental unitarity, giving the first definition of quantum conditional independence. What is missing is the analogous justification of quantum causal models from fundamental unitarity. That is, one might reasonably expect a framework of quantum causal models which generalises quantum Reichenbach to be derivable by assuming that the mappings between nodes are all unitary, supplemented by latent nodes which satisfy the qualitative part of Reichenbach’s principle. This has not yet been achieved. The interconnected nature of general causal structures makes this more difficult, since a global unitary evolution normally assumes a sense of global time. However, it is suspected that the task is not insurmountable and the result would place quantum causal models on even firmer footing.

6

Causal Loops: Time Travel to the Past

6.1 Closed Timelike Curves and the Ontology of Time Travel

This thesis is primarily concerned with the restrictions that quantum theory places on ontology and causality. These topics meet head on when considering the possibility of time travel to the past in a quantum universe. Quantum theory is famously linear but as soon as time travel is included non-linearities creep in. This destroys the equivalence between many ontological interpretations possible in standard quantum theory. Moreover, the lack of consensus towards quantum ontology means a range of possible models for time travel should be considered. Most of the work in this chapter has been published in Ref. [All14].

6.1.1 Why Study Time Travel?

There are few areas of physics in which one confronts the idea of time travel to the past. Indeed, from the causal perspective, true causal loops are normally ruled out by hypothesis (this assumption being built into the use of directed *acyclic* graphs in causal models, Chap. 5). When it is discussed, time travel is often associated with particular ways of thinking about quantum theory and quantum field theory. One sometimes hears the view that in quantum teleportation, for example, the teleported information travels back in time to the point at which entanglement was created, before proceeding forward in time to the recipient [Pen98, Jos98, Joz04, Tim06]. Another example might be remarks about antiparticles in quantum field theory being akin to particles “travelling backwards in time”.

The time travel considered in this chapter is rather different. Here, it will be assumed that time travel into the past is a possible physical process and the question is asked: how could or should quantum theory be modified to account for time travel?

There are a few ways to motivate this. Perhaps the most common is in relation to general relativity and *closed timelike curves* (CTCs), discussed in Sec. 6.1.2. More pertinent to this thesis are the insights it may give to the relationship between causality and ontology in quantum theory. In a classical and ostensibly deterministic universe our familiar ideas of causality and physical ontology are strong and considering time travel to the past is easy enough to be the subject of many of popular stories. For quantum theory, on the other hand, there are a wide variety of possible ontological interpretations, none of which are uncontroversial. As shall be seen, this gives rise to many reasonable models for quantum time travel. Time travel to the past forces us to consider quantum features such as indeterminism and inseparability of states in a new light.

These conversations are valuable even if time travel to the past is not possible in our universe. Treated as a thought experiment, time travel can draw particular attention to issues that may be overlooked in other situations. In the words of Deutsch in the original paper on quantum time travel as approached in this thesis: “It is curious that the analysis of a physical situation which might well not occur should yield so many insights into quantum theory” [Deu91].

Taking the opposite perspective, theoretical studies into quantum time travel may yield *a priori* restrictions on whether and how such effects are logically possible. It is well known that classical treatments of time travel are prone to paradoxes, but a common feature of most quantum approaches to time travel is that certain paradoxes can be avoided altogether. This may be taken as evidence that time travel to the past is not as nonsensical as our classically-influenced brains might assume. On the other hand, if a very well-motivated argument is found stating that no sensible account of quantum time travel can be given this may be taken as evidence that time travel to the past is impossible.

There are also computational motivations for studying time travel. One of the main tools in the complexity theorist’s box is analysing the abilities of different computational paradigms when given access to powerful additional resources, often specified as “oracles” [AB09]. This allows the derivation of so-called “relativised” separation results, which establish relationships between complexity classes given access to certain oracles. The situation with time travel to the past need be no different, allowing comparison of different models of computation when given access to these anachronistic resources. This is of particular interest in quantum computation. Most confirmed computational speed-ups of quantum computation over classical computation are in tasks which involve an oracle’s

resources [NC00]. It is therefore interesting to compare the capabilities of quantum and classical computers with access to different models of time travel.

Finally, by incorporating time travel to the past, one obtains a non-linear extension to quantum theory [Haw95, And95, FW95, Har94]. Standard quantum theory is characteristically linear and yet the possibility of non-linear extensions to it remain a subject of interest, especially with regard to any hypothetical post-quantum physics [GRW86, Wei89, Gol08, Ken05, Pen98, AW09, Aar05, CMP12]. In this view, quantum time travel provides an interesting source of reasonable non-linear extensions which can serve as either examples or counter-examples for how non-linear quantum theory can or must behave.

6.1.2 Closed Timelike Curves and General Relativity

By far the most common motivation given for studying time travel to the past is that general relativity allows for closed timelike curves (CTCs). CTCs are features of exotic spacetimes which allow massive particles to travel to their own past, while heading apparently forwards in time at all points. As such they give a concrete possible mechanism for time travel to the past. It has long been known that solutions containing CTCs can be found to the Einstein field equations [Göd49, Bon80, Got91]. Unsurprisingly, there are doubts that any such solutions could be found in nature [Haw92, MTY88] but it has not been possible to absolutely preclude them. It is therefore prudent to treat them as real possibilities and, since our universe is quantum, the problem of quantum theory along CTCs (that is, quantum time travel to the past) must be addressed.

The most obvious approach to analysing quantum theory with CTCs is to use quantum field theory and general relativity in curved spacetimes. This is far from easy, however. Spacetimes containing CTCs are globally non-hyperbolic and typically contain no Cauchy surfaces. This makes them incompatible with standard relativistic quantum field theory and without a well-defined initial value problem [BD82, Wal10].

Nonetheless, there is a tradition of attempts to model quantum time travel using the path-integral approach to quantum theory [EKT91, BD82, Har94, Pol94]. The approach taken in this chapter will, however, be somewhat more abstract. Rather than dealing directly with curved spacetimes, one can use the quantum circuit approach [Deu91]. The result is an abstraction away from the precise time travel mechanism, retaining only the essential feature that some system is sent into its own past [All14]. This move has turned out to be fruitful, giving rise to the two most popular current models of time travel in quantum theory: *D-CTCs* (or “Deutschian” CTCs) and *P-CTCs* (or “post-selection” CTCs).

This general approach to modelling quantum time travel will be detailed in Sec. 6.2 together with a short discussion of the types of paradoxes that time-travel can cause in the classical case. Section 6.3 will then briefly outline the models of D- and P-CTCs and discuss the ontological issues raised, in particular with respect to non-linearity. The main results of the chapter will be in Sec. 6.4, where alternatives to the D- and P-CTC models are found. One of these alternative models, dubbed T-CTCs, will be fully fleshed out and compared with the established models. Finally, Sec. 6.5 will summarise these results and discuss their significance for the ontology of quantum theory, time travel, and non-linear extensions to quantum theory more generally.

6.2 Modelling Time Travel with Quantum Circuits

6.2.1 The Standard Form of Time Travel Circuits

The quantum circuit approach to time travel is based on a particular form of circuit introduced in Ref. [Deu91]. This convenient building block, from which all other circuits involving time travel can be built, will be called the *standard form circuit*. It allows different circuits and models to be easily and concisely specified.

The quantum circuit model is an abstraction from the precise physical mechanisms that neatly separates quantum evolution from spatial motion. Quantum interactions described by unitary gates are assumed to only occur in small, freely falling, non-rotating regions of spacetime so that they obey non-relativistic quantum theory. To include time travel to the past in such a quantum circuit model is therefore equivalent to saying that the classical paths quantum systems take between gates are allowed to go back in time. Within this approach, different models are then defined by their behaviour when sending systems back in time. For clarity, this chapter will only deal with finite-dimensional Hilbert spaces.

The standard form circuit, illustrated in Fig. 6.1, contains a single time travel event in a localised spacetime region and a single unitary quantum interaction U . To simplify the discussion, it will be assumed that there is some system in the circuit which does not go back in time¹. The *chronology violating* (CV) system arrives from its own future in the state τ_i and after the interaction is said to be in the state τ_f . Meanwhile, the *chronology respecting* (CR) system arrives from the unambiguous past in the state ρ_i and emerges into the unambiguous future in the state ρ_f . Therefore, a standard form circuit is completely

¹This assumption could be dropped, but at the expense of a more fiddly discussion. It also seems reasonable to expect any model containing non-localised time travel events to be extendible to a model where all are localised.




Figure 6.1: Schematic diagram for the standard form circuit described in the text. The chronology respecting (CR) and chronology violating (CV) systems are shown entering the gate labelled by unitary U . Time, when it is unambiguous at least, increases up the diagram. The double bars represent the time travel event and may be thought of as two depictions of the same spacelike hypersurface forming a CTC. The dashed lines represent spacelike boundaries of the region in which time travel takes place; CV system is restricted to that region.

specified by these two systems and U , while a model for quantum theory with time travel is a general specification of ρ_f given U and ρ_i .

In cases where the CV system is initially known to be in a pure state it may be written as $|\phi\rangle$ so that $\tau_i = |\phi\rangle\langle\phi|$. Similarly, in cases where the CR system input or output states are known to be pure, they may be written $|\psi_{i,f}\rangle$ so that $\rho_{i,f} = |\psi_{i,f}\rangle\langle\psi_{i,f}|$. In Sec. 6.3.3 it will be seen that ρ_i can always be purified to some $|\psi_i\rangle$.

A standard assumption used in all circuit models of quantum time travel is that the CR and CV systems are not initially entangled. That is, before the action of U they are in the product state $\rho_i \otimes \tau_i$. [DFI10]. It has been argued that this is an unreasonable assumption and that prior entanglement should be considered as the CV system contains ρ_i in its past [Pol94]. That possibility will not be considered here because alternative assignment methods all have undesirable features. The problem is one of finding an assignment procedure that gives a joint state ω on $\mathcal{H}_{\text{CR}} \otimes \mathcal{H}_{\text{CV}}$ given only a state ρ_i on \mathcal{H}_{CR} . The procedure assumed here—where $\omega = \rho_i \otimes \tau_i$ —is the only one for which: ω is always positive, $\text{Tr}_{\text{CV}} \omega = \rho_i$, and mixtures are preserved [Ali95].

6.2.2 Time Travel Paradoxes and the Classical Model

In order to define the possible types of paradox in time travel it is useful to leave quantum theory to one side briefly and concentrate on classical time travel. Consider a classical version of the standard form circuit, with classical states $\tilde{\rho}_{i,f}$ and $\tilde{\tau}_{i,f}$ and some classical

dynamical evolution \tilde{U} replacing their quantum counterparts. Since classically it can be assumed that there is no fundamental stochasticity, assume that $\tilde{\rho}_{i,f}$ and $\tilde{\tau}_{i,f}$ are both ontic.

The standard way of introducing time travel into classical theories is to impose a consistency condition on states that go back in time; that is, $\tilde{\tau} \stackrel{\text{def}}{=} \tilde{\tau}_i = \tilde{\tau}_f$. In other words, the ontic state that emerges into the past is required to be the same one that left from the future. With knowledge of a $\tilde{\rho}_i$ and \tilde{U} , a consistent $\tilde{\tau}$ can be deduced, from which $\tilde{\rho}_f$ may be calculated.

Classical time travel to the past may give rise to both *dynamical consistency* paradoxes and *information* paradoxes. This chapter takes a slightly unconventional view on paradoxes whereby a computation, which algorithmically produces unambiguous output from an input, is never paradoxical. So a theory that predicts “absurdly” powerful communication or computational abilities will not be called paradoxical because of them, even if those abilities make the theory appear unreasonable or hard to accept. Some “absurd” conclusions should probably be expected when “absurdity” of time travel has been assumed.

A dynamical consistency paradox is a situation in which a consistent history of events is not possible [Deu91, LMGP⁺11b, LMGP⁺11a]. The usual example is the “grandfather paradox” where a grandchild travels back to kill their infant grandfather. Dynamical consistency paradoxes occur when a model fails to specify any valid final state from some initial state and evolution. Classically, this is exactly because the consistency condition cannot be satisfied². The only way to avoid these paradoxes in general is to disallow the interactions that lead to them from the model [All14].

An information paradox is a situation with consistent dynamics but information that appears from nowhere; *viz.* the information has not been computed. The prototypical example is the “unproved theorem” paradox: a mathematician reads the proof of a theorem from a book only to travel back in time to author that same book. The proof has no ultimate source.

Under these definitions, information paradoxes arise if and only if a theory contains a *uniqueness ambiguity*: the model specifies more than one final state given some initial state and evolution, but fails to give probabilities for each possibility. Any other dynamically consistent evolution is counted as a computation and information paradoxes have been defined as exactly those where an uncomputed output is produced. For example, in the classical unproved theorem paradox there are many time travelling states $\tilde{\tau}$ compatible with the consistency condition, each producing a different $\tilde{\rho}_f$ —one $\tilde{\tau}$ produces a theorem

²It is worth noting that, although these paradoxes appear to be possible in classical models of physics, such situations are quite difficult, if not impossible, to construct in classical models with continuous state spaces. Reference [AM13] is a useful introduction to such issues.

answering $P \stackrel{?}{=} NP$, another answering $BPP \stackrel{?}{=} BQP$, many others where the “theorem” is nonsense, *etc.* It should be noted that the equivalence of the uniqueness ambiguity and information paradoxes is not universally used, but it follows from the definitions used here as described above. Readers preferring different definitions may replace each further instance of “information paradox” with “uniqueness ambiguity” without affecting the meaning.

For example, some authors prefer a wider definition of paradox in which time travel circuits are counted as having information paradoxes when the only consistent evolution reveals a fixed point of some given function [Deu91, LMGP⁺11c]. Such circuits uniquely produce solutions to problems that are hard if $P \neq NP$ very rapidly [GJ79]. As such, they are counted as very powerful computations rather than paradoxes as defined here. This is not to claim that such processes are necessarily reasonable, but rather to reflect that they are qualitatively different from the uniqueness ambiguities equivalent to information paradoxes as defined here. If there is no uniqueness ambiguity then any information that appears as a result of the time travel circuit is uniquely specified by the structure of, and input to, the circuit [LMGP⁺11c]. It is therefore reasonable to say that this circuit is algorithmically computing as instructed.

6.3 Quantum Time Travel Models, Non-linearity, and Ontology

The model of D-CTCs, introduced by Deutsch, was the first to use the quantum circuit model to analyse quantum time travel [Deu91]. P-CTCs are a more recent development which make use of post-selection and ideas from quantum teleportation to construct a very different model. In this section each will be briefly introduced and discussed, focussing on what they suggest about ontology and non-linearity in quantum theory.

6.3.1 Overview of D-CTCs

While not the original line of reasoning, the D-CTC model can be rapidly constructed by assuming that reduced density operators are ontic states, following Ref. [WB12]. Just as in the classical model, this imposes a consistency condition $\tau \stackrel{\text{def}}{=} \tau_i = \tau_f$ on the ontic time travelling states. For a standard form circuit, this implies

$$\tau = \mathcal{G}(\tau) \stackrel{\text{def}}{=} \text{Tr}_{\text{CR}} (U(\rho_i \otimes \tau)U^\dagger). \quad (6.1)$$

Noting that the CR and CV systems have been separated using a partial trace in Eq. (6.1), this suggests that the final state should be given by

$$\rho_f = \text{Tr}_{\text{CV}} (U(\rho_i \otimes \tau)U^\dagger). \quad (6.2)$$

The key difference between this and ordinary unitary quantum theory is the implied map

$$U(\rho_i \otimes \tau)U^\dagger \rightarrow \rho_f \otimes \tau, \quad (6.3)$$

that replaces the quantum state after U with the product of its reduced density operators. This defines the action of the time travel event in D-CTCs.

Between them, Eqs. (6.1, 6.2) define the D-CTC model of time travel. Given ρ_i and U , one can solve Eq. (6.1) to get τ and then calculate ρ_f . Equation (6.2) takes the role of an equation of motion and is clearly both non-linear and non-unitary in general.

The first thing to note is that a solution τ for Eq. (6.1) always exists. This follows from Schauder's fixed point theorem which guarantees that every trace-preserving quantum channel, such as \mathcal{G} , has at least one fixed point [Sch30, Tyc35, Zei85]. Therefore dynamical consistency paradoxes cannot arise in the D-CTC model³. On the other hand, solutions for τ are not always unique. Therefore, D-CTCs have the same uniqueness ambiguity present in classical time travel and hence information paradoxes.

To avoid uniqueness ambiguities the *maximum entropy rule* has been suggested [Deu91], stating that one should choose the unique τ with maximum von Neumann entropy. However, this is not universally accepted as part of the D-CTC model and alternative principles do exist [Pol94, DFI10]. Interestingly, these ambiguities vanish in the presence of arbitrary non-zero noise. Suppose one incorporates a noise channel \mathcal{N} applied to the CV system so that Eq. (6.1) becomes $\tau = \mathcal{N}(\mathcal{G}(\tau))$. As a noisy channel, $\mathcal{N}(\mathcal{G}(\cdot))$ should be *strictly contractive* (that is, trace distance should always decrease under its action) [Rag02] and therefore have a unique fixed point⁴ [NC00].

Thus, D-CTCs are free from dynamical consistency paradoxes but can have information paradoxes (which vanish under arbitrarily small noise). The non-linearity of the model enables them to have the following abilities, beyond those of ordinary quantum circuits.

They can solve any problem in PSPACE in polynomial time and are therefore likely to be vastly more powerful even than quantum computers [AW09]. D-CTCs are also capable of producing discontinuous mappings from ρ_i to ρ_f , which means that (for all practical purposes) the model loses predictive power near these discontinuities.

³A direct proof of this can be found in Ref. [Deu91]

⁴This argument resolves a conjecture from Ref. [All14].

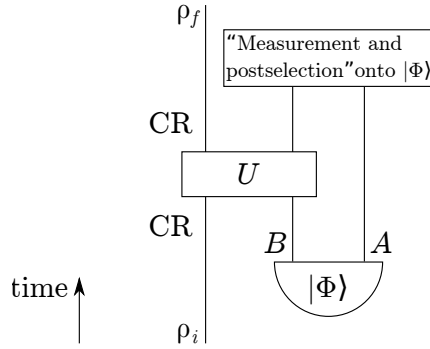


Figure 6.2: Schematic illustration of the protocol defining the action of P-CTCs as described in the text. In standard teleportation, Alice and Bob share entangled systems A and B and Alice can teleport a state to Bob by performing a joint measurement with A and obtaining the outcome $|\Phi\rangle$. This situation differs from that of standard teleportation in two ways. First, the system that Alice teleports is the same as B , just at a later time. Second, that Alice can get the outcome $|\Phi\rangle$ by postselection with certainty and so no classical communication to Bob is required to complete the teleportation.

Given any finite set of pure states (which are not necessarily orthogonal), there is a D-CTC circuit that can render them distinguishable with a single measurement [BW12]. In this way, D-CTCs can violate the Holevo bound [BHW09, Hol73]. Given this distinguishing ability, it is perhaps unsurprising that D-CTC circuits are also capable of cloning arbitrary quantum states [AMRM12, BWW13] (though, of course, entanglement is not cloned as this is forbidden by monogamy [CKW00]).

6.3.2 Overview of P-CTCs

The P-CTC model is due to Svetlichny [Sve09] (inspired by diagrammatic approaches to quantum theory [CK17]) and Lloyd *et al.* [LMGP⁺11b, LMGP⁺11a] (based on the unpublished work of Bennett and Schumacher [BS05] and inspired by Ref. [HM04]). Reference [BW12] contains an accessible introduction. It is defined by ignoring the precise mechanism behind the time travel and postulating only that the effect is mathematically equivalent to teleportation into the past, achieved by the following unphysical operational protocol schematically illustrated in Fig. 6.2.

Prepare two copies of the CV system, A and B , in the maximally entangled state $|\Phi\rangle \propto \sum_i |i\rangle_B |i\rangle_A$, where $\{|i\rangle\}_i$ is any orthonormal basis of the CV system. Let B interact with the CR system as in the standard form circuit and then perform a joint measurement on B and A which contains outcome $|\Phi\rangle$. The unphysical step is to postselect this measurement on the outcome $|\Phi\rangle$. This is equivalent to simply projecting the tripartite system of CR, B , and A onto $\langle\Phi|$ and then renormalising the resulting state. Comparing this to the standard

quantum teleportation protocol, the effect is to “teleport” the final state of B back onto the B system in the past. This protocol may be simulated in the laboratory by manual postselection of measurement outcomes [Sve09, LMGP⁺11b].

If the CR system is initially in the pure state $|\psi_i\rangle$ then the effect of this protocol is (up to renormalisation)

$$\begin{aligned} {}_{BA}\langle\Phi|U_{\text{CR},B}|\psi_i\rangle|\Phi\rangle_{BA} &\propto \sum_i {}_B\langle i|U_{\text{CR},B}|i\rangle_B|\psi_i\rangle \\ &\propto \text{Tr}_{\text{CV}}(U)|\psi_i\rangle. \end{aligned} \quad (6.4)$$

Generalising this result to mixed input ρ_i and including the renormalisation, the general action of a P-CTC becomes

$$\rho_f = \frac{P\rho_iP^\dagger}{\text{Tr}(P\rho_iP^\dagger)}, \quad P \stackrel{\text{def}}{=} \text{Tr}_{\text{CV}}(U). \quad (6.5)$$

The map of Eq. (6.5) completely specifies the action of a P-CTC standard form circuit. So having obtained this result, one can take Eq. (6.5) as the definition of the P-CTC model and take the above unphysical protocol simply as an argument showing that the model is equivalent to teleporting the final state of the CV system into the past. This equation of motion is both non-linear and non-unitary in general.

Since Eq. (6.5) maps each ρ_i onto a specific ρ_f without ambiguity, P-CTCs cannot suffer uniqueness ambiguities and so are not vulnerable to information paradoxes. On the other hand, the P operators defined in Eq. (6.5) can act as $P\rho_iP^\dagger = 0$ on some input states ρ_i . These cases are dynamical consistency paradoxes as no consistent outcome can be obtained from these inputs and interactions. However, such dynamical consistency paradoxes vanish under arbitrary non-zero noise [All14]. Note how P-CTCs generically suffer from dynamical consistency paradoxes, while D-CTCs do not, but D-CTCs suffer from information paradoxes, while P-CTCs do not, and in both cases the paradoxes vanish under the action of any finite noise.

As with the D-CTC model, the non-linearity of this model permits various super-quantum abilities. P-CTC circuits can solve any problem in PP in polynomial time [LMGP⁺11a], making them likely much more powerful than quantum computers but likely much less powerful than D-CTCs. This follows because P-CTCs and quantum theory with postselection are computationally equivalent [BW12, Aar05]. Unlike D-CTCs, P-CTCs can never produce discontinuous evolutions because of the more mild form of non-linearity, as shown in Sec. 6.3.3.

There are P-CTC circuits that can render any finite set of linearly independent pure states distinguishable with a single measurement [BHW09]. Note that, while still able to

distinguish non-orthogonal quantum states with certainty, this is much less powerful than distinguishing with D-CTCs, as linear independence is required (in particular, the Holevo bound cannot be violated in this way). P-CTCs are also capable of generically deleting arbitrary quantum states, something impossible in standard linear quantum theory [PB00].

6.3.3 The Problems of Non-linearity

Even from this brief overview, it should be clear that the non-linearity of D-CTCs and P-CTCs causes some striking departures from standard linear quantum theory. Central features, including no-cloning and indistinguishability of non-orthogonal states, are broken due to the non-linearity and non-unitarity of Eqs. (6.2, 6.5). However, the abilities of both models are still bounded and it is sensible to consider how and why these two models differ.

There are different types of non-linear maps one can consider applying to quantum states. Generally, non-linearity occurs due to input state appearing at quadratic or higher orders in an equation of motion, as in D-CTCs. A special case of non-linearity is *renormalisation non-linearity*, where the equation of motion is linear except for a scalar factor which simply normalises the final state. P-CTCs are renormalisation non-linear since Eq. (6.5) would be entirely linear were it not for the renormalising denominator. A non-linear equation may be called *polynomial non-linear* if it is not renormalisation non-linear.

This distinction facilitates a more general discussion of different non-linear theories and their features. For example, renormalisation non-linear equations cannot lead to the discontinuous state evolutions possible with D-CTCs. It is simple to verify that if $\rho_f(\rho_i) = L(\rho_i) / \text{Tr}(L(\rho_i))$ is a renormalisation non-linear equation of motion (that is, if L is linear in this equation) then

$$\lim_{\epsilon \rightarrow 0} \rho_f(\rho_i + \epsilon\sigma) = \rho_f(\rho_i) \quad (6.6)$$

for any σ so the mapping is always continuous.

There are also different types of mixed quantum states that can be identified. Typically, one encounters *proper mixtures*—epistemic mixtures due to the observer’s ignorance of the actual quantum state—and *improper mixtures*—the way to describe only part of an entangled system. Both of these are described using density operators and in standard quantum theory two mixtures described by the same density operator behave identically regardless of the type.

The operational equivalence of different types of mixture directly facilitates the operational equivalence of many interpretations and ontologies for quantum theory. For

example, Everettian ontologies have only a single global quantum state that is hyper-entangled and evolves deterministically, implying an absence of proper mixtures in favour of improper mixtures⁵. A GRW-style collapse theory [GRW86] would, on the other hand, allow both types of mixture and specifically have many fewer improper mixtures.

In any extension to quantum theory with non-linearity, however, it is not valid to describe proper mixtures using density operators. This follows because, for a non-linear evolution \mathcal{E} acting on some ensemble of states and corresponding probabilities $\{(\rho_j, p_j)\}_j$, applying \mathcal{E} to the initial density operator is not generally the same as the density operator obtained by applying \mathcal{E} to each individual state in the ensemble:

$$\mathcal{E} \left(\sum_j p_j \rho_j \right) \neq \sum_j p_j \mathcal{E}(\rho_j). \quad (6.7)$$

On the other hand, density operators are the correct way to describe improper mixtures under non-linear evolution. By examining the derivation of reduced density operators, as given in [NC00] for example, it is easily seen that linearity of operations is not assumed at any point.

A third distinct type of mixed state becomes particularly evident when discussing time travel models [Deu91, BW12]. These are *true mixtures*: mixtures that are not entangled to any reference system and yet would still be described as mixed by an observer with maximum knowledge. These arise naturally in the model of D-CTCs, where it is possible for a mixed output to be produced from a pure state. Moreover, the CV state τ in D-CTCs will generally be mixed and cannot be purified, so must also be a true mixture [PCA11].

However, both improper and true mixtures may still be validly described by density operators in a non-linear theory (unlike proper mixtures). It therefore follows that non-linear theories do not treat true and improper mixtures differently. So whilst models of time travel may introduce true mixtures conceptually, they do not affect the way in which calculations are performed. As such the purification theorem still holds for ρ_i and one may always assume that $\rho_i = |\psi_i\rangle\langle\psi_i|$ by simply extending U to act on the purification ancilla as the identity.

⁵Of course, any Everettian observer could still choose to be ignorant, but this does not give rise to mixed states in the same way. To be clear, consider how proper mixtures arise in objective collapse models. An observer O sets up a measurement M on a system S , but chooses (by not looking, or what have you) to remain ignorant of the outcome of M . In an objective collapse model, O now knows (assuming sufficient understanding of S , M , and quantum theory) that the state of the universe has now collapsed into one of multiple possible states with corresponding probabilities; this ensemble forms a proper mixture which O uses to describe S after M . In an Everettian model, however, O now knows that the universal state has evolved into some macroscopic superposition, the only uncertainty is about which branch of this O is in. For more on this see, for example, Ref. [Alb10, especially §3.2].

The different ontologies implied by proper, improper, and true mixtures must be carefully borne in mind when dealing with non-linear models, such as those of time travel. Insufficient clarity on this point caused significant controversy over the capabilities of D-CTCs, the so-called “linearity trap” [BLSS09, RM10, CM10, CMP12, BW12, Pie13]. The solution is to not only specify the density operator of a mixture in a non-linear model, but also specify the ontology of the mixture so that the correct treatment can be used [All14].

Moreover, since all non-linear evolutions treat proper and improper mixtures differently in general then if the difference is observable the result is an *entanglement detector*: a device capable of telling whether or not a system is entangled with another. In any interpretation that involves instantaneous disentanglement by measurement then an entanglement detector necessarily facilitates superluminal signalling [Pie13, §3.1.3]. D-CTCs and P-CTCs therefore both lead to superluminal signalling with such an interpretation.

Instantaneous disentanglement is not a necessary feature of quantum measurement, however. There are many ontologies that do not require instantaneous disentanglement. Moreover, one can construct alternative models of quantum measurement that prevent superluminal signalling in non-linear theories by construction without subscribing to a particular interpretation [CD02, Ken05]. So whilst non-linearity does not necessarily lead to signalling, it may do depending on the ontology [Pie13, §3].

6.3.4 The Role of Ontology

Any non-linear extension to quantum theory (and, in particular, models of time travel) raises ontological problems compared to standard quantum theory. The clearest examples are in the differences that non-linearity introduces between proper and improper mixtures. As discussed in Sec. 6.3.3, this essentially means that density operators cannot be used to describe proper mixtures in a non-linear model.

Traditionally, the only differences between the types of mixture in quantum theory have been ones of interpretation and preferred ontology. Non-linearity breaks this long-standing equivalence of different interpretations. For example, with non-linearity Everettian quantum theory need not produce the same predictions as Bohmian quantum theory simply because they differ ontologically. Therefore, if a method of time-travel to the past were discovered (a real CTC, for example) then in principle it could be used to experimentally distinguish between certain ontological interpretations.

There is a wrinkle in this argument, however, that points towards the other role of ontology in models of time travel. That is, there are at least two models of time travel in quantum theory each of which is somewhat well-motivated: D-CTCs and P-CTCs. They are only “somewhat” well-motivated as neither comes with a solid first-principles argument

deriving the model, they are instead developed using plausibility arguments. A clear idea of a quantum ontology will play a key role in selecting the correct model for time travel.

Compare this to the classical case. An uncontroversial model for classical time travel to the past was briefly outlined in Sec. 6.2.2. Why was this model used and why is it not controversial? Because, as far as physics is concerned, there is broad consensus on the essential features of classical ontology. Classical physics suggests that systems have a local ontology that describes their entire state at any point, so it is easy to simply transport this local ontology back in time when required.

As an aside, it is interesting to note that in Ref. [AW09] a curiously different model of classical time travel was used. In that paper, it was claimed that the computational power of D-CTCs is equivalent to a classical computer equipped with time travel abilities (both being PSPACE). However, this is only true if one takes the unusual position of requiring that classical time travel imposes a consistency condition on *probability distributions* rather than ontic states, as in Sec. 6.2.2. In this way, Ref. [AW09] avoids dynamical consistency paradoxes normally associated with classical time travel⁶.

To illustrate how ontology influences the plausibility of different time travel models, consider D-CTCs. When first introduced, the suggested interpretation of D-CTCs was in an Everettian ontology where the CV system heading back in time enters a different branch or “world” of the quantum state [Deu91]. This is not an ontological assumption from which the model is derived, but rather a suggested interpretation that makes a certain amount of sense of the model and was likely influenced by Deutsch’s pre-existing preference for that interpretation⁷. On the other hand, D-CTCs can be analysed from an epistemic realist perspective on quantum state ontology—see Sec. 2.1.2—and found to be inconsistent [WB12]. The conclusion is clear: someone, like Deutsch, who prefers Everettian ontology can consistently choose the D-CTC model (though they are not necessarily forced to) while an epistemic realist cannot. Moreover, as in Sec. 6.3.1, one can try to justify D-CTCs by assuming that reduced mixtures are ontic states and therefore suitable candidates for a classical-like consistency condition, but taking this ontological route seriously requires a careful justification of the very unusual quantum ontology that results.

Interestingly, P-CTCs were introduced without any suggested ontological basis at all. In fact, one of the papers that introduced the model preferred to call it “effective

⁶Another way to look at this oddity is to note that, since classical theory is a subset of quantum theory, P-CTCs should be able to model classical time travel but P-CTCs only have the computational power of PP. If classical time travel has the power of PSPACE this would seem to suggest $PP = PSPACE$ which would be a hugely surprising result (of course, this is not implied mathematically, but merely illustrates the odd model of classical time travel assumed).

⁷In Ref. [Deu91] he repeatedly refers to Everettian quantum theory as “unmodified” quantum theory.

quantum time travel” that one might simulate in a lab using manual postselection [Sve09]. The primary motivation is that quantum teleportation creates a quantum communication channel and that the (unphysical) protocol described in Sec. 6.3.2 modifies this channel to communicate with the past. Despite this, there are still ontological arguments to be made for and against P-CTCs. For example, P-CTCs have been shown to be compatible with one model for quantum time travel that takes the path-integral approach, rather than the circuit approach used here [Pol94]. On the other hand, the P-CTC model fails to specify a state for the system that travels back in time analogous to τ in D-CTCs. It is therefore unclear how to answer questions like “what is the state of the CR and CV system just after applying U ?” with P-CTCs, for example. This may not be necessary for some mechanisms of time travel, but it would certainly be bizarre if there were not a well-defined state for a system traversing a CTC.

It is a strength of the abstract quantum circuit-based approach that it can easily isolate these ontological concerns. In particular, it facilitates direct comparisons between the classical and quantum cases which then suggests two core features of quantum theory that contribute to the ontological ambiguity found in quantum time travel. First, the stochasticity found in quantum theory is normally thought to be irreducible, while classical physics is normally taken to be fundamentally deterministic. Classically, this allows one to unproblematically consider the state of the CV system to be in a definite ontic state, even if the precise state is not known. Second, quantum states of the CV and CR systems are generally non-separable, whereas classically there is always a well-defined concept of the state of the CV system separate from the CR system. Classically, this allows the CV system to be extracted from the future and transplanted to the past easily, while quantum mechanically some non-trivial operation is required to separate the two. With D-CTCs this is a partial trace, while with P-CTCs it is a projection. Without these two features, defining an uncontroversial quantum model for time travel would likely be as easy as it is classically.

On the other hand, the circuit approach also introduces some further ambiguities as a result of being somewhat divorced from the precise physical mechanisms. In particular, one prominent ambiguity might be called the *dynamical ambiguity* [All14]. In a D-CTC circuit, for example, Eq. (6.3) specifies a non-trivial dynamical change but it is not clear exactly when this should occur. There is a similar ambiguity in P-CTCs: there is a physical change that has no well-defined location. These ambiguities are related to how, in quantum circuits, one can slide gates along wires freely without changing the overall effect of the circuit and is probably therefore a fundamental ambiguity for this circuit approach. While a little inelegant, this ambiguity need not be of particular concern, however, because it is

entirely unobservable in both models (the predictions are the same regardless of where Eq. (6.3) is placed, for example).

So ontology can and should inform one's approach to quantum time travel. Likewise, theorising about quantum ontology can highlight certain ontological issues such as non-separability in quantum theory and whether it can make sense to discuss a fundamentally different sort of evolution on one part of a system (CV) from the other (CR). There is also one further link between the studies of quantum time travel and ontology: models of quantum time travel might directly influence the development of ontologies for quantum theory. Suppose, for example, that a physical CTC were discovered. Experiments on it could reveal a particular model of quantum time travel to be correct which should then, in turn, suggest some quantum ontologies as more plausible than others (*e.g.* D-CTCs may suggest Everettian ontology and rule out many epistemic realist ontologies). It may even be possible for this to happen entirely theoretically: if an exceptionally well-motivated and natural model for quantum time travel were developed, then this would probably also provide clues towards natural ontologies for quantum theory.

6.4 Alternative Time Travel Models

The discussion of D-CTC and P-CTC models in the previous section makes three things clear. First, there can be more than one reasonable way to extend quantum theory to include time travel to the past. Second, there are several reasons to dislike either of these models, depending on one's philosophical bent. Third, the ontological and interpretational foundations of each are far from certain but, if one thing is known, they definitely differ.

These points raise some questions. What other reasonable models of quantum theory with time travel might exist? How might they compare to these existing examples? What ontological implications would they suggest? These questions will be tackled in this section. Before developing some new theories, it will be useful to first review some background on integrating over quantum states.

6.4.1 Integrating over Quantum States

Consider any d -dimensional quantum system with Hilbert space \mathcal{H} . For any scalar function $\mathcal{J} : \mathcal{H} \rightarrow \mathbb{C}$ one can consider the integral over the pure states

$$J = \int_{\mathcal{P}(\mathcal{H})} d[\phi] \mathcal{J}(\phi), \quad (6.8)$$

where the integration measure $d[\phi]$ is yet to be defined. Conveniently, there exists a unique natural measure over $\mathcal{P}(\mathcal{H})$ that is invariant under unitary transformations given by taking

a random unitary matrix distributed according to the Haar measure on the group $U(d)$ [ZS01]. One way to write this is in the Hurwitz parametrisation [BZ06, ZS01, Hur97] defined with respect to some orthonormal basis $\{|\alpha\rangle\}_{\alpha=0}^{d-1}$ of \mathcal{H} , such that any pure state $|\phi\rangle$ takes the form

$$|\phi\rangle = \prod_{\beta=d-1}^1 \sin \theta_\beta |0\rangle + \sum_{\alpha=1}^{d-2} e^{i\varphi_\alpha} \cos \theta_\alpha \prod_{\beta=d-1}^{\alpha+1} \sin \theta_\beta |\alpha\rangle + e^{i\varphi_{d-1}} \cos \theta_{d-1} |d-1\rangle, \quad (6.9)$$

for some parameters $\theta_\alpha \in [0, \pi/2]$ and $\varphi_\alpha \in [0, 2\pi)$. In this parametrisation, the integration measure takes the form

$$d[\phi(\theta_\alpha, \varphi_\alpha)] = \prod_{\alpha=1}^{d-1} \cos \theta_\alpha (\sin \theta_\alpha)^{2\alpha-1} d\theta_\alpha d\varphi_\alpha. \quad (6.10)$$

This is then a natural measure for integrating over pure quantum states that is unique up to a multiplicative constant.

This natural measure is invariant under unitary operations, so that under $|\phi\rangle \rightarrow U|\phi\rangle$ the measure transforms as $d[\phi] \rightarrow d[U\phi] = d[\phi]$. It may be useful to observe that the Hurwitz parametrisation is a generalisation of the Bloch sphere parametrisation often used for qubits $\mathcal{H} = \mathbb{C}^2$ where the measure is, up to a scalar, the rotationally-invariant area measure on a sphere:

$$|\phi\rangle = \sin \theta |0\rangle + e^{i\varphi} \cos \theta |1\rangle, \quad (6.11)$$

$$d[\phi] \propto \sin(2\theta) d(2\theta) d\varphi. \quad (6.12)$$

Integrals over mixed quantum states can be considered similarly. For some scalar function $\mathcal{J} : \mathcal{D}(\mathcal{H}) \rightarrow \mathbb{C}$, consider the integral over the density operators on \mathcal{H} ,

$$J = \int_{\mathcal{D}(\mathcal{H})} d[\tau] \mathcal{J}(\tau), \quad (6.13)$$

for some integration measure $d[\tau]$. Unlike $\mathcal{P}(\mathcal{H})$, there is no unique natural measure on $\mathcal{D}(\mathcal{H})$ [BZ06] and so one must be chosen, along with a useful way to parametrise τ . The result is that there is no unique natural way to define J ; it will depend on the choice of measure used.

6.4.2 Desiderata

When considering how new models of quantum theory with time travel might be developed, it is useful to first consider how it might be desirable for such a theory to behave. A list of possible desirable features follows. Of course, all desiderata are linked to various philosophical prejudices, but there is still utility in considering them.

1. The model should have sound physical motivation and an ontological interpretation.
2. The model should be compatible with standard quantum theory, at least in so far as current experiments are concerned. In the case of CTCs it should reproduce quantum theory locally along the CTC, as well as in spacetime regions far from the CTC. It is also expected to be locally approximately consistent with special relativity and relativistic causal structure. Specifically, it should not allow superluminal signalling.
3. The model should not allow dynamical inconsistencies. In other words, it should not have disallowed evolutions that lead to dynamical consistency paradoxes for any U or ρ_i .
4. The model should specify ρ_f uniquely given U and ρ_i . If multiple possible output states are considered, then probabilities for each of these should be specified. In other words, it should not have uniqueness ambiguities that lead to information paradoxes (Sec. 6.2.2).
5. The model should specify a state τ that travels back in time; this should either be uniquely specified or an ensemble of possibilities with corresponding probabilities should be uniquely specified.
6. Given a pure ρ_i , prejudice might require either or both of ρ_f and τ to also be pure.
7. The model should not be able to distinguish non-orthogonal states in a single measurement, neither should it be able to clone arbitrary quantum states.

Feature (1) is the most subtle and subjective of these but arguably the most important. It was discussed to some extent in Sec. 6.3.4 for D- and P-CTCs.

D-CTCs have feature (2) so long as ontological assumptions regarding collapse are made that rule out superluminal signalling. P-CTCs only have feature (2) in the presence of finite noise, and even then similar assumptions about collapse are required to rule out signalling. However, as noted in Sec. 6.3.3, adding any non-linear evolution to quantum theory opens up the possibility of signalling in this way.

Feature (3) is not in the D-CTC model but is in the P-CTC model, while feature (4) is definitely in P-CTCs but is only in D-CTCs by adding an extra postulate or in the presence of noise, as noted in Sec. 6.3.1.

Neither the D-CTC nor P-CTC models have feature (5). P-CTCs do not specify any τ , while D-CTCs specify τ but not necessarily uniquely. Feature (6) is perhaps the least

compelling feature listed and is one that neither D-CTCs nor P-CTCs have. Feature (7) is also not one respected by either P-CTCs or D-CTCs.

Of course, it would be ambitious to ask for a model that satisfies all these desiderata. Notably, the standard way of introducing time travel into classical mechanics does not satisfy features (2), (3), (4), or (5). However, they do form a helpful guide for where one might start looking for new time travel models.

6.4.3 Some Alternative Models

The above desiderata can be used to guide the construction of new models of time travel in quantum theory. In this section two overlapping classes of new models will be considered: *weighted D-CTCs* and *transition probability models*. An example of the latter, dubbed *T-CTCs*, will be fully fleshed out in Sec. 6.4.5.

Weighted D-CTCs represent an extension of the model of D-CTCs. These are described by parametrising the convex subset of density operators τ_α allowed by the consistency condition Eq. (6.2) with α and then assigning a weight $w_\alpha \geq 0$ to each of these valid CV states. This weighted mixture is then used to calculate $\tau = \int d\alpha w_\alpha \tau_\alpha / \int d\alpha w_\alpha$. The D-CTC protocol can then be used with this uniquely determined choice of τ .

Clearly, this describes a whole class of models based on how the parametrisation is done and which integration measure is chosen. For example, one could weight all possibilities equally $w_\alpha = 1$, giving a *uniform* weighted D-CTC model. In terms of the desiderata, this theory would gain features (4) and (5) at least, possibly at the expense of feature (1) depending on the details and motivation of the model. Such a model is essentially that of D-CTCs, with an alternative to the maximum entropy rule.

Transition probability models make use of some useful intuition from standard quantum theory. It is common to say that the probability of an initial state $|I\rangle$ to transition into a final state $|F\rangle$ under the unitary transformation V is given by the transition probability $|\langle F|V|I\rangle|^2$. More precisely, what is meant is that $|\langle F|V|I\rangle|^2$ is the Born rule probability of finding the system in state $|F\rangle$ if one were to measure the system to see if it is in state $|F\rangle$ after the transformation. As an example of this useful way of thinking, consider starting with a bipartite system, initially in state $|\psi_i\rangle|\phi\rangle$, and act upon it with the unitary U ; the “probability of finding the second system in $|\phi\rangle$ ” after the transformation is $p(\phi) = \|\langle\phi|U|\psi_i\rangle|\phi\rangle\|^2$.

The transition probability can be generalised to apply to mixed states ρ and τ . One way to do this is to use the Hilbert-Schmidt inner product $\text{Tr}(\rho\tau)$, which is the probability for ρ to be found in an eigenstate of the proper mixture τ and for τ to be a realisation of that same eigenstate, averaged over all eigenstates of τ . Another option is the square of

the fidelity, the interpretation of which involves considering ρ as an improper mixture on which a measurement is performed by projective measurement of the larger purified system [Uhl11]. Both of these options reduce to the transition probability in the case of both states being pure. The use of any mixed state transition probability must be motivated by its operational meaning in context. For the remainder of this chapter it will be assumed that the appropriate generalisation of transition probability to mixed states is the Hilbert-Schmidt inner product, so that the probability for a bipartite system initially in the state $\rho_i \otimes \tau$ to have the second system found in the state τ after some unitary transformation U is given by⁸ $p(\tau) = \text{Tr}(\tau U(\rho_i \otimes \tau)U^\dagger)$.

Transition probability models are obtained by applying these ideas to time travel. The choices that need to be made to define a specific model include: whether pure or mixed CV states are used, which $|\phi\rangle$ or τ are to be considered, and how ρ_f should be separated from the CV system.

One example of a transition probability model is also a weighted D-CTC model. This is obtained by choosing the weights of a weighted D-CTC model to be the transition probabilities $w_\alpha = p(\tau_\alpha) = \text{Tr}(\tau_\alpha^2)$, giving the equation of motion

$$\rho_f = \frac{\int d\alpha \text{Tr}(\tau_\alpha^2) \text{Tr}_{\text{CV}}(U(\rho_i \otimes \tau_\alpha)U^\dagger)}{\int d\alpha \text{Tr}(\tau_\alpha^2)}. \quad (6.14)$$

Another collection of transition probability models is found by integrating over all possible initial CV states, weighted by the corresponding transition probability, and so

$$\tau_i = Z^{-1} \int_{\mathcal{P}(\mathcal{H}_{\text{CV}})} d[\phi] p(\phi) |\phi\rangle\langle\phi| \quad (6.15)$$

is used if pure CV states are considered, or similarly with an integral over $\mathcal{D}(\mathcal{H}_{\text{CV}})$ if mixed CV states are considered. $Z > 0$ normalises the state τ_i and $p(\phi)$ is the transition probability as described above. There are two options for then isolating ρ_f : either take the partial trace as with D-CTCs or use the same partial projection used to calculate the transition probability. The first of these, using pure CV states, gives a model with the equation of motion

$$\rho_f = Z^{-1} \int d[\phi] p(\phi) \text{Tr}_{\text{CV}}(U(|\psi_i\rangle\langle\psi_i| \otimes |\phi\rangle\langle\phi|)U^\dagger), \quad (6.16)$$

which is one of various models on this theme.

⁸One curiosity of using this generalisation of transition probability is that the probability for ρ to transition to ρ under unitary $\mathbb{1}$ is strictly less than unity for mixed ρ . This is simply a reflection of the fact that mixed states can be viewed as epistemic states over the pure states.

Another such transition probability model is the *model of T-CTCs*, which is the model found by integrating over all pure CV states, weighted by transition probability, but using the partial projection to find ρ_f . T-CTCs will be developed fully in Sec. 6.4.5, complete with a discussion of their physical motivation and ontological implications. Some of the other theories that are variations on this theme will be briefly considered in Sec. 6.4.5.6.

6.4.4 The Uniqueness Ambiguity and Epistemic Reasoning

Before proceeding to detail the model of T-CTCs, some remarks are in order about the uniqueness ambiguity. In Sec. 6.2.2 this ambiguity was introduced as a necessary and sufficient condition for a model to suffer information paradoxes, as defined in that section. The argument hinges on the idea that if the final state is uniquely determined by the initial state and dynamics, then what has occurred can be regarded as a (possibly very powerful) computation and is therefore not paradoxical according to that definition.

The argument still holds if the unique final state is an epistemic state (that is, a probability distribution over ontic states), so long as the probabilities in the epistemic state are determined by the physics rather than purely epistemic principles. If the probabilities are physically determined then any new information obtained can be viewed as being due to a probabilistic computation. For any particular final state to be likely, the physics must not only establish that final state as a possibility, but also that the corresponding probability is sufficiently high. Models for probabilistic computation are well-established and certainly not paradoxical.

Compare this to D-CTCs without either noise or the maximum entropy rule. In Sec. 6.3.1 it was claimed that D-CTCs suffer from the uniqueness ambiguity and therefore information paradoxes. This is different from a probabilistic computation since D-CTCs assign no probabilities to the possible final states; they are merely left as possibilities.

Therefore, as far as the definitions in Sec. 6.2.2 go, information paradoxes are still avoided if uniqueness ambiguity is avoided; *viz.*, when a unique physically determined epistemic state is specified. It is for this reason that requirements (4) and (5) of Sec. 6.4.2 allow for uniquely specified epistemic states.

For example, suppose a time travel circuit is designed to produce previously unknown theorems. If this circuit produces unique theorems with certainty from the input, then, as discussed in Sec. 6.2.2, this is a type of computation: a novel automated theorem prover. Similarly, if the circuit produces one of a selection of possible theorems from the input, each with a given probability, then the circuit is performing a (possibly novel) probabilistic computation. On the other hand, if a model allows for a circuit that could produce one of

a range of possible theorems but has no way of giving a probability for each, then that is an information paradox.

6.4.5 T-CTCs

The model of T-CTCs may be motivated as follows. Consider a CR observer watching a standard form time travel circuit evolve and suppose that the primitive states of quantum theory are pure states. Because of this, and the purification theorem, take $\rho_i = |\psi_i\rangle\langle\psi_i|$ to be pure.

This observer watches a CV system emerge from the future in some unknown pure state $|\phi\rangle$. This is then observed to interact with a CR system, initially in state $|\psi_i\rangle$, via unitary U . The CV system then proceeds to head back in time. At this point, the observer may judge whether any given $|\phi\rangle$ is a consistent initial CV state. If someone were to have measured the CV system immediately before it travelled back in time, then the probability of their finding that any given $|\phi\rangle$ is a consistent initial state would be $p(\phi) = \|\langle\phi|U|\psi_i\rangle\|^2$. So for any pair $|\phi_1\rangle$ and $|\phi_2\rangle$, the former would be found to be consistent $p(\phi_1)/p(\phi_2)$ times more often than the latter. It therefore seems reasonable to conclude that $|\phi_1\rangle$ is $p(\phi_1)/p(\phi_2)$ times more likely to have been the initial state than $|\phi_2\rangle$. The observer therefore considers τ_i as a proper mixture over all $|\phi\rangle \in \mathcal{P}(\mathcal{H}_{CV})$, each weighted by $p(\phi)$, Eq. (6.15). On the other hand, consistency demands that if $|\phi\rangle$ was the initial CV state, then on heading back in time the CV system must be found to be in the same state again. So the observer can describe the final state of the CR system in each case by the partial projection $\langle\phi|U|\psi_i\rangle\langle\psi_i|U|\phi\rangle / \|\langle\phi|U|\psi_i\rangle\|^2$ consistent with this being the case. The resulting final state for the CR system is, therefore,

$$\rho_f = Z^{-1} \int d[\phi] U_\phi |\psi_i\rangle\langle\psi_i| U_\phi^\dagger, \quad (6.17)$$

$$U_\phi \stackrel{\text{def}}{=} \langle\phi|U|\phi\rangle, \quad (6.18)$$

$$Z \stackrel{\text{def}}{=} \int d[\phi] \langle\psi_i|U_\phi^\dagger U_\phi|\psi_i\rangle, \quad (6.19)$$

where the operator U_ϕ acts only on \mathcal{H}_{CR} and the constant $Z > 0$ is defined to normalise ρ_f .

Equations (6.17–6.19) define the behaviour of the model of T-CTCs. The CR input state $|\psi_i\rangle$ was assumed to be pure, but these equations can equally be applied to mixed input states by simply replacing instances of $|\psi_i\rangle\langle\psi_i|$ with ρ_i .

This is not intended as a derivation, but a motivational explanation for T-CTCs akin to those given for D-CTCs and P-CTCs. As with those models, T-CTCs are defined by the Eqs. (6.17–6.19) rather than by any particular interpretation. One might even use a similar

argument to motivate other models, including some of the other transition probability models mentioned in Sec. 6.4.3.

Several features of T-CTC immediately follow from Eqs. (6.17–6.19). First, it is a non-unitary and non-linear model. Second, it is only renormalisation non-linear and as such it only gives rise to continuous evolutions, as discussed in Sec. 6.3.3. Third, there is no ambiguity in the equation of motion (6.17), so there is no uniqueness ambiguity and no information paradoxes. Before proceeding to consider what other features T-CTCs may have, it will first be useful to re-write Eq. (6.17) in a simpler form.

6.4.5.1 Simplification of the T-CTC Equation of Motion

The equation of motion (6.17) for T-CTCs in its current form is rather opaque. In order to more easily calculate with the model it is useful to perform the integration in generality and thereby simplify this equation.

Let $\{|\alpha\rangle\}_{i=0}^{d-1}$ be an orthonormal basis for the d -dimensional CV system and expand the unitary U in the Kronecker product form in this basis $U = \sum_{\alpha,\beta} A_{\alpha\beta} \otimes |\alpha\rangle\langle\beta|$, where $A_{\alpha\beta}$ are operators on the CR system. In this form, the equation of motion is

$$\rho_f = Z^{-1} \sum_{\alpha,\beta,\gamma,\delta} I_{\alpha\beta,\gamma\delta} A_{\alpha\beta} |\psi_i\rangle\langle\psi_i| A_{\gamma\delta}^\dagger, \quad (6.20)$$

having defined the integrals

$$I_{\alpha\beta,\gamma\delta} = \int d[\phi] \langle\phi|\alpha\rangle\langle\beta|\phi\rangle\langle\phi|\delta\rangle\langle\gamma|\phi\rangle. \quad (6.21)$$

Now consider expanding both $d[\phi]$ and $|\phi\rangle$ in the Hurwitz parametrisation [Eqs. (6.9, 6.10)] with respect to the same basis. Since, for each α , $d\varphi_\alpha$ factorises out of the measure, any integrand in which the only φ_α -dependence is an integer power of $e^{i\varphi_\alpha}$ will integrate to zero. Considering the integrals in Eq. (6.21), every integrand will have such a phase factor unless at least one of the two following conditions is met: $\alpha = \beta$ and $\gamma = \delta$, or $\alpha = \gamma$ and $\beta = \delta$. In these cases, all phase factors will cancel out and the phase integrals will not come to zero. Discarding these zero integrals in Eq. (6.20) it is, therefore, found that

$$\rho_f = Z^{-1} \left(\sum_{\alpha \neq \beta} I_{\alpha\beta,\alpha\beta} A_{\alpha\beta} |\psi_i\rangle\langle\psi_i| A_{\alpha\beta}^\dagger + \sum_{\alpha \neq \beta} I_{\alpha\alpha,\beta\beta} A_{\alpha\alpha} |\psi_i\rangle\langle\psi_i| A_{\beta\beta}^\dagger + \sum_{\alpha} I_{\alpha\alpha,\alpha\alpha} A_{\alpha\alpha} |\psi_i\rangle\langle\psi_i| A_{\alpha\alpha}^\dagger \right). \quad (6.22)$$

By unitary invariance of the integration measure one may rotate $|\phi\rangle$ in each of these integrals so that only the $|d-1\rangle$ and $|d-2\rangle$ components contribute. Therefore, for $\alpha \neq \beta$, $I_{\alpha\beta,\alpha\beta}$ and $I_{\alpha\alpha,\beta\beta}$ are both equal to

$$\begin{aligned} & \int d[\phi] |\langle\phi|d-1\rangle|^2 |\langle\phi|d-2\rangle|^2 \\ &= (2\pi)^{d-1} \left(\int \prod_{\gamma=1}^{d-3} d\theta_\gamma \cos \theta_\gamma (\sin \theta_\gamma)^{2\gamma-1} \right) \\ & \times \int d\theta_{d-1} d\theta_{d-2} \cos^3 \theta_{d-1} \cos^3 \theta_{d-2} \sin^{2d-1} \theta_{d-1} \sin^{2d-5} \theta_{d-2}, \quad (6.23) \end{aligned}$$

where in the final line the integrand has been expanded out in the Hurwitz parametrisation. Similarly

$$\begin{aligned} I_{\alpha\alpha,\alpha\alpha} &= \int d[\phi] |\langle\phi|d-1\rangle|^4 \\ &= (2\pi)^{d-1} \left(\int \prod_{\gamma=1}^{d-3} d\theta_\gamma \cos \theta_\gamma (\sin \theta_\gamma)^{2\gamma-1} \right) \\ & \times \int d\theta_{d-1} d\theta_{d-2} \cos^5 \theta_{d-1} \cos \theta_{d-2} \sin^{2d-3} \theta_{d-1} \sin^{2d-5} \theta_{d-2}. \quad (6.24) \end{aligned}$$

By evaluating the final lines of Eqs. (6.23, 6.24) it is seen that, for $\alpha \neq \beta$, the ratio $I_{\alpha\alpha,\alpha\alpha}/I_{\alpha\beta,\alpha\beta} = 2$. From Eq. (6.22) one therefore finds

$$\rho_f \propto \sum_{\alpha,\beta} \left(A_{\alpha\beta} |\psi_i\rangle \langle\psi_i| A_{\alpha\beta}^\dagger + A_{\alpha\alpha} |\psi_i\rangle \langle\psi_i| A_{\beta\beta}^\dagger \right). \quad (6.25)$$

Finally, note the following identities, which may readily be verified by expanding the traces: $P \stackrel{\text{def}}{=} \text{Tr}_{\text{CV}} U = \sum_\alpha A_{\alpha\alpha}$ and $\sum_{\alpha,\beta} A_{\alpha\beta} |\psi_i\rangle \langle\psi_i| A_{\alpha\beta}^\dagger = \text{Tr}_{\text{CV}} \left(U (|\psi_i\rangle \langle\psi_i| \otimes \mathbb{1}) U^\dagger \right)$. Using these, and introducing a normalising scalar $z > 0$ (which is generally different from Z used before), the final form of the equation of motion becomes

$$\rho_f = z^{-1} \left(P |\psi_i\rangle \langle\psi_i| P^\dagger + d \text{Tr}_{\text{CV}} \left(U \left(|\psi_i\rangle \langle\psi_i| \otimes \frac{\mathbb{1}}{d} \right) U^\dagger \right) \right). \quad (6.26)$$

Equation (6.26) is in a much more revealing form than Eq. (6.17). It shows that the T-CTC equation of motion is a weighted mixture of the corresponding P-CTC equation of motion (6.5) with an ordinary quantum channel. This gives the impression that T-CTCs are akin to noisy P-CTCs. Like Eq. (6.17), this simplified equation of motion can equally be applied to mixed CR input states ρ_i .

6.4.5.2 The P Operator

The operator P in Eq. (6.26) is the same as used in P-CTCs in Eq. (6.5). It is the partial trace of a unitary operator and is therefore not generally unitary itself. For instance, for a general P there can be states $|\phi\rangle$ in $\mathcal{P}(\mathcal{H})$ for which $P|\phi\rangle = 0$, which would not be possible if P were unitary.

So P does not generally preserve state norms, but the effect of P on state norms is still bounded. Consider P acting on a vector $|\psi\rangle$ and let $\{|\alpha\rangle\}_{\alpha=0}^{d-1}$ be any orthonormal basis on the d -dimensional CV system. Using the triangle inequality and unitarity one can bound $\|P|\psi\rangle\|$ as follows

$$\|P|\psi\rangle\| \leq \sum_{\alpha} \|\langle\alpha|U|\psi\rangle|\alpha\rangle\| \leq \sum_{\alpha} \|U|\psi\rangle|\alpha\rangle\| = d \|\psi\|. \quad (6.27)$$

Moreover, this bound can be achieved as demonstrated in Ref. [BW17] and is therefore the tightest possible general bound.

6.4.5.3 Paradoxes

It has already been noted that T-CTCs have no uniqueness ambiguities and therefore no information paradoxes. With the simplified equation of motion (6.26) to hand it is now easy to see that T-CTCs are also always dynamically consistent. Even though it is possible for $P|\psi_i\rangle = 0$, the second term of Eq. (6.26) will always give a non-zero density operator. Therefore, T-CTCs contain neither type of paradox identified in Sec. 6.2.2. Unlike P-CTCs and D-CTCs, no noise or extra rule is required to avoid these paradoxes.

An example is instructive. Consider the following toy model of an unproved theorem paradox as a standard form time travel circuit, introduced in Ref. [LMGP⁺11b] and illustrated in Fig. 6.3. The CR system is a pair of qubits, M and B representing the mathematician and book respectively, initially in state $|0\rangle_B|0\rangle_M$. The CV system is a single qubit representing the mathematician heading back in time. The unitary U consists of a pair of CNOT gates representing the writing and reading of the book with a swap for when the mathematician swaps places with their time-travelling self, as illustrated. Clearly, by extending this toy model to use N qubits for each of M , B , and CV it would allow a theorem to be encoded in an N -bit string.

For this circuit the P operator is

$$P = |0\rangle_B\langle 0| \otimes |0\rangle_M\langle 0| + |0\rangle_B\langle 1| \otimes |1\rangle_M\langle 1| + |1\rangle_B\langle 1| \otimes |0\rangle_M\langle 1| + |1\rangle_B\langle 0| \otimes |1\rangle_M\langle 0| \quad (6.28)$$

and therefore $P|00\rangle_{BM} = |00\rangle_{BM} + |11\rangle_{BM}$. To get the second term of Eq. (6.26), note that $\text{Tr}_{CV}(U(|00\rangle_{BM}\langle 00| \otimes \mathbb{1})U^\dagger) = |00\rangle_{BM}\langle 00| + |11\rangle_{BM}\langle 11|$. The output the unproved




Figure 6.3: Toy model of an unproved theorem paradox as a standard form circuit from Ref. [LMGP⁺11b]. The two CR qubits are labelled B for the book and M for the mathematician. The unitary gates illustrated represent, from bottom to top, the book being written with a CNOT, the book being read with a CNOT, and the mathematician swapping places with their future self.

theorem T-CTC circuit is therefore

$$\rho_f = \frac{1}{2}|00\rangle_{BM}\langle 00| + \frac{1}{4}|00\rangle_{BM}\langle 11| + \frac{1}{4}|11\rangle_{BM}\langle 00| + \frac{1}{2}|11\rangle_{BM}\langle 11|. \quad (6.29)$$

This result may also be verified, with rather more effort, directly from the integral expression Eq. (6.17).

Compare this to the cases of D-CTCs and P-CTCs. For D-CTCs, any state diagonal in the computational basis will satisfy Eq. (6.1) for this circuit. The result is a one-parameter continuous family of consistent evolutions representing any probabilistic mixture of possible “theorems” along the CV system. Notably, since the D-CTC description is diagonal it is entirely classical, except the consistency condition requires consistency on probability distributions over the theorems rather than the theorems themselves⁹. For P-CTCs, Eq. (6.28) gives $|\psi_f\rangle = \frac{1}{\sqrt{2}}(|0\rangle_B|0\rangle_M + |1\rangle_B|1\rangle_M)$ as the circuit’s output, so one obtains an equal superposition over all possible “theorems”.

6.4.5.4 Computation

Some basic facts about the computational abilities of T-CTCs can be read straight from Eq. (6.26). The second term could be realised in ordinary quantum theory and so is limited to the power of BQP, while the first term is the P-CTC equation of motion (6.5). It follows immediately that P-CTCs can therefore trivially simulate T-CTCs, so clearly T-CTCs cannot efficiently solve any problems that are not contained within PP. This shows that P-CTCs are at least as computationally powerful as T-CTCs.

⁹*C.f.* the unusual model for classical time travel used in Ref. [AW09], discussed in Sec. 6.3.4

Moreover, the form of Eq. (6.26) suggests that T-CTCs may be less powerful than P-CTCs. This is because, for a T-CTC, ρ_f is only a pure state if either $P|\psi_i\rangle = 0$ or if the two terms in Eq. (6.26) are equal. So every T-CTC algorithm that outputs a pure state is achievable on an ordinary quantum computer in exactly the same way. Any potential algorithm for a T-CTC-equipped computer that makes computational use of the extra power of first term in Eq. (6.26) would therefore have to output mixed states. This observation also prevents T-CTCs from being able to perform an arbitrary postselected quantum measurement, since many postselected measurement outcomes are pure states. Therefore, one could not prove that T-CTCs have the power of PP using the same method used for P-CTCs [LMGP⁺11a].

Despite this, it has recently been shown that T-CTCs can efficiently simulate any P-CTC circuit to arbitrary precision [BW17] (that is, the undesired “error term” can be made exponentially small with a linear number of CV qubits). This does not invalidate the above comments, which demonstrate that T-CTCs cannot *perfectly* simulate P-CTCs. This approximate simulation is, however, enough to show that T-CTCs do indeed have the computational power of PP.

6.4.5.5 Mixed States and Non-linearity

In Sec. 6.3.3 it was noted that both improper and true mixtures are validly described with density operators in non-linear extensions of quantum theory and that the purification still holds so $\rho_i = |\psi_i\rangle\langle\psi_i|$ may always be assumed. This remains true in the model of T-CTCs. It also remains true that proper mixtures are not validly described by density operators and that non-linearity of T-CTCs allows for the possibility of creating an entanglement detector and thereby signalling, exactly as with D-CTCs and P-CTCs.

Another consequence of non-linearity is that both D-CTCs and P-CTCs are capable of distinguishing non-orthogonal states in a single measurement. However, this is not the case with T-CTCs as shall now be shown.

Consider the problem of distinguishing between two states ρ and σ . The probability of success when using a single optimal measurement is given by $\frac{1}{2}(1 + D(\rho, \sigma))$, where $D(\rho, \sigma) \stackrel{\text{def}}{=} \frac{1}{2} \text{Tr} |\rho - \sigma|$ is the trace distance between the states [FG99]. Therefore, ρ and σ are perfectly distinguishable in a single measurement if and only if $D(\rho, \sigma) = 1$.

Another measure of distinguishability of quantum states is the the fidelity between ρ and σ , defined as [NC00]

$$F(\rho, \sigma) \stackrel{\text{def}}{=} \text{Tr} \sqrt{\rho^{1/2} \sigma \rho^{1/2}}. \quad (6.30)$$

In the case of pure states $|a\rangle$ and $|b\rangle$, the Fidelity takes on the particularly simple form $F(|a\rangle, |b\rangle) = |\langle a|b\rangle|$.

Suppose one wishes to distinguish quantum states using a T-CTC. Only pure state inputs need be considered, so what is required is a bound on the distinguishability of the output states of some T-CTC circuit, ρ_f^a and ρ_f^b , for which the input states were $|a\rangle$ and $|b\rangle$, respectively. Note first that trace distance is bounded by fidelity [NC00]

$$D(\rho_f^a, \rho_f^b) \leq \sqrt{1 - F(\rho_f^a, \rho_f^b)}. \quad (6.31)$$

It is then useful to separate the terms of $\rho_f^{a,b}$ seen in Eq. (6.26). Therefore, write $\rho_f^\psi = (1 - \lambda^\psi)\sigma^\psi + \lambda^\psi\tau^\psi$, where $\tau^\psi \stackrel{\text{def}}{=} \text{Tr}(U(|\psi\rangle\langle\psi| \otimes \frac{\mathbb{1}}{d})U^\dagger)$ and

$$\lambda^\psi \stackrel{\text{def}}{=} \frac{d}{d + \|P|\psi\rangle\|^2} \geq \frac{1}{d + 1} \quad (6.32)$$

where this inequality is a result of Eq. (6.27). Using strong concavity and monotonicity of fidelity under quantum operations [NC00] it is seen that

$$\begin{aligned} F(\rho_f^a, \rho_f^b) &\geq \sqrt{\lambda^a \lambda^b} F(\tau^a, \tau^b) \\ &\geq \sqrt{\lambda^a \lambda^b} F(|a\rangle, |b\rangle) \geq \frac{1}{d + 1} |\langle a|b\rangle|. \end{aligned} \quad (6.33)$$

Finally, observe that by Eqs. (6.31–6.33)

$$D(\rho_f^a, \rho_f^b) \leq \sqrt{1 - \frac{|\langle a|b\rangle|^2}{(d + 1)^2}} \leq 1, \quad (6.34)$$

with equality to unity only possible if $\langle a|b\rangle = 0$.

This proves that the output states of a T-CTC circuit are only perfectly distinguishable from one another in a single measurement if the input states were. However, as proved in Ref. [BW17], T-CTCs can approximately simulate P-CTCs to arbitrary precision efficiently. It therefore follows that while T-CTCs cannot *perfectly* distinguish non-orthogonal quantum states, they can distinguish any set of linearly independent non-orthogonal states to arbitrary precision by simulating the corresponding P-CTC circuit.

It is comparatively very simple to observe that T-CTCs are incapable of cloning pure states. A pure state cloning machine always outputs a pure state. Since any T-CTC outputting a pure state can be simulated exactly by an ordinary quantum operation, then the no-cloning theorem for T-CTCs is simply a result of the no-cloning theorem in ordinary quantum theory. In exactly the same way, it also immediately follows that T-CTCs are incapable of deleting arbitrary pure states. However, the question as to whether mixed states can be broadcast [BCF⁺96] is left open.

6.4.5.6 Relation to Other Alternatives

Before discussing these results, it should be noted that some of the results proved for T-CTCs above can be easily modified to apply to some of the closely related models introduced in Sec. 6.4.3.

First, consider the modification to T-CTCs where, instead of integrating over pure CV states, mixed CV states are integrated over. This represents a class of models since there is no unique natural choice for the integration measure, but it is still possible to deduce some general properties. These models are only renormalisation non-linear and so continuity follows immediately. It is also possible to show that these models always define a unique non-zero ρ_f for every ρ_i and U , so that the models suffer neither dynamical consistency nor information paradoxes¹⁰.

Now consider the modification to T-CTCs expressed in Eq. (6.16), where, instead of separating CV and CR systems by a projection, they are separated by a partial trace. This model always defines a unique non-zero¹¹ ρ_f , thus avoiding both dynamical consistency and information paradoxes. It is not, however, renormalisation non-linear but polynomial non-linear.

There are, of course, further variations which could be considered. For example, one could use an alternative generalisation of transition probability for mixed states, as noted in Sec. 6.4.3. The purpose of this discussion is to show that whilst T-CTCs were focussed on above, the other theories mentioned in Sec. 6.4.3 also have reasonable properties and may be worthy of further development.

6.5 Summary and Discussion

This chapter approached ontology and causality in quantum theory from an unusual angle. By assuming that time travel to the past is possible, one is confronted with an interesting cross-section of problems from ontology and causality. In particular, one is drawn towards non-linear extensions of quantum theory which cast light on certain ontological issues in quantum theory.

Having discussed the various motivations for studying quantum time travel in Sec. 6.1, this chapter followed the tradition of circuit-based approaches to time travel to the past.

¹⁰The proof that ρ_f is non-zero follows by showing that the integrand is positive semi-definite and that there exist some τ for which the integrand is non-zero. Importantly, an appropriate assumption would be needed about the positivity of the chosen measure.

¹¹The proof that ρ_f is always non-zero follows similarly to the previous case, by proving that the integrand is always positive semi-definite and there always exist CV states $|\phi\rangle$ for which both the integrand and $d[\phi]$ are non-zero.

This is a convenient abstract approach that can be summarised using standard form circuits introduced in Sec. 6.2.1. Of course, time travel to the past always raises the possibility of paradoxes, which were systematically introduced in Sec. 6.2.2 in the context of classical time travel.

Following this, previous work on quantum time travel in the circuit approach was outlined, focussing on the model of D-CTCs in Sec. 6.3.1 and the model of P-CTCs in Sec. 6.3.2. The roles of non-linearity and ontology in these models were discussed in Secs. 6.3.3, 6.3.4 where it was noted that not only do one's ontological preferences influence time travel models, but that the types of reasonable non-linear extensions to quantum theory may hint at approaches to quantum ontology.

The main results of this chapter were in Sec. 6.4. The strengths and shortcomings of D- and P-CTCs were used to identify two classes of new quantum time travel models. In particular, the model of T-CTCs was developed in full and its properties were derived.

Non-linear extensions of quantum theory are subtle since the long-standing plurality of co-existing interpretations is broken. When considering time travel this manifests itself in at least two ways. The first is in the development and motivation of various possible models: ontological bias will affect decisions made. The second is in using those models: mixed states with ontological differences but the same density operator may behave differently, as discussed in Sec. 6.3.3. Neither of these issues arise when considering time travel classically, since ontology is generally clear and non-linear evolutions are commonplace.

This uniquely quantum issue has both positive and negative effects on the resulting models. The way in which quantum theory works allows models of time travel that do not suffer from the paradoxes that are present classically, but which generally break some of the central structure of quantum theory. Distinguishability of non-orthogonal states, state cloning/deleting, and the spectre of superluminal signalling all present themselves. It also appears that computational power is greatly increased even beyond that of quantum computers.

The existing models of quantum time travel are not without their shortcomings. Most troubling is that both D-CTCs and P-CTCs suffer from paradoxes (of the information and dynamical consistency types respectively). While both can be eliminated by arbitrarily small noise, the models themselves remain paradoxical. The two classes of new models presented in Sec. 6.4.3 were designed to avoid these paradoxes and hopefully also satisfy many of the other desiderata laid out in Sec. 6.4.2.

Of the new models, that of T-CTCs was selected primarily due to its physical motivation. To illustrate the strength of the physical story told in Sec. 6.4.5, consider applying the same reasoning in a classical context.

A CR observer watching a classical time travel circuit sees a CV system in an unknown ontic state $\tilde{\tau}_i$ emerge from the future, interact with a CR system in a known state, and then disappear back to the past in the ontic state $\tilde{\tau}_f$. For each $\tilde{\tau}_i$, the observer knows that when it heads back in time it must be found to be in the same state. Whilst in quantum theory the probability of finding a system in a given state is given by the transition probability, the corresponding probability classically is either unity or zero: Either $\tilde{\tau}_i = \tilde{\tau}_f$ or $\tilde{\tau}_i \neq \tilde{\tau}_f$. So when the CR observer takes a probability distribution over all possible CV states $\tilde{\tau}_i$, the only states to which non-zero probabilities are assigned are precisely those states for which $\tilde{\tau}_i = \tilde{\tau}_f$ after the interaction. What is missing from this account is a way of specifying the relative probabilities assigned to each CV state. One might choose to use the principle of indifference and weight each possibility equally, but this is not necessary.

A very similar argument could be used to argue in favour of D-CTCs, for example, by demanding exact equality of reduced density operators rather than consistency via the transition probability. In this case, using the principle of indifference would lead to a result that is equivalent to the uniform weighted D-CTCs mentioned in Sec. 6.4.3. However, by accepting the interpretation of transition probabilities and supposing that only pure states are primitive, T-CTCs do have a clear physical motivation.

Of course, this was never meant to be a cast-iron argument for T-CTCs and there is a certain amount of vagueness in the description given in Sec. 6.4.5. Similar interpretational vagueness is found with both D-CTCs and P-CTCs and should be expected when attempting to extend quantum theory (which lacks consensus on interpretation) to a non-linear regime that is so alien to it.

These arguments for the classical model and uniform weighted D-CTCs differ from the argument for T-CTCs in an important respect. The probabilities assigned to the different possible histories with T-CTCs are physically determined: they are proportional to the transition probabilities. On the other hand, in the above discussion of the classical model and D-CTCs using the same narrative there is no physical assignment of probabilities. The use of principle of indifference is an epistemic move, not a physical one.

Consider the desirable features listed in Sec. 6.4.2 in the light of the model of T-CTCs. The above argument aims to satisfy feature (1) at least as far as with D- or P-CTCs. As with D- and P-CTCs, feature (2) is satisfied so long as an appropriate ontology of measurement is chosen [Sec. 6.3.3]. Features (3), (4), and (5) are satisfied without condition, as discussed in Sec. 6.4.5.3. Feature (6) is potentially partially satisfied, in that both τ and ρ_f might be considered ontologically pure, but since a proper mixture is taken over so many possibilities, the mathematical form of either is very rarely pure. In Sec. 6.4.5.5 it is shown that feature

(7) is satisfied strictly speaking, although if one allows for arbitrarily small error then it is violated.

Does this mean that T-CTCs are a better model for quantum theory with time travel or, more specifically, for quantum theory in the presence of CTCs? Not necessarily. The physical motivation for T-CTCs is far from a “first principles” argument and there is still the question as to how, and when, the state projection occurs. However, both D-CTCs and P-CTCs have incomplete physical motivations and both leave questions as to exactly how, and when, a proposed physical change occurs—all three models share the unobservable dynamical ambiguity.

What has been comprehensively shown is that there is a whole landscape of other theories out there. The quantum circuit approach to quantum theory with time travel may be very attractive in that it abstracts away from knotty problems with spacetime geometry or any other exact mechanism for time travel, but it is perhaps too general for the problem at hand. In order to identify a more robustly physical solution to quantum theory with time travel it may be necessary to use a different approach, such as path integral or field theoretic ideas. Alternatively, by very carefully committing to a specific ontology for quantum theory it may be possible to identify the corresponding theory of time travel. When non-linearity is present vagueness on this point is problematic.

7

Conclusions and Further Work

This thesis has considered the possibilities and impossibilities for ontologies of quantum systems in three broad ways. First, by taking a rather traditional approach to the ontology of quantum states. Second, by considering how to properly analyse causality in quantum systems. Third, by considering how one might account for time travel to the past in a quantum universe and how this interacts with possible ontologies. The conclusions for each will now be discussed in turn, together with what they suggest in terms of further work.

7.1 State Ontology and Macro-realism

The limitations that quantum theory places on the ontology of states and the possibility of macro-realist ontologies were discussed in Chaps. 2 and 3 respectively. In the first instance, these chapters were concerned with which ontological features can or cannot be compatible with quantum theory, with extensions concerning compatibility with experiments being developed later.

The primary result of Chap. 2 was Thm. 2.2 which proved that almost every quantum superposition state must be ontic in $d > 3$ dimensions (equivalently, they cannot be epistemic). The conclusion for any prospective epistemic realist must be that even though a superposition inherits all of its properties and dynamics from underlying basis states, one cannot use this fact to construct a simpler ontology where the superposition also inherits its ontic states from underlying basis states. Note that the theorem gives almost

the strongest version of this statement possible, holding for almost every superposition with respect to any given orthonormal basis.

To illustrate, consider Bohmian mechanics of a single particle and its implied ontology where ontic states are of the form $\lambda = (\vec{r}, |\psi\rangle)$, as described in Sec. 3.2.4. Here, the position basis of the particle plays a special role, with its “true value” being represented by the position vector \vec{r} . A less extravagant variation on this ontology might aim to make use of this special role and have other quantum states inherit their ontology from the ontic states of the position basis. However, Thm. 2.2 shows that this can never be achieved—superpositions will always require novel ontic states not accessible to the underlying basis in order to be compatible with quantum theory.

One avenue for developing on this result is to consider *ontic independence* more generally. By analogy with linear independence in geometry, one might define a quantum state $|\psi\rangle$ to have ontic independence with respect to some set of other quantum states if it can access ontic states not covered by preparations of that set. In this language, Thm. 2.2 requires almost all quantum states to have ontic independence with respect to any given basis to be compatible with quantum theory in $d > 3$ dimensions. It would be interesting to see if more powerful ontic independence results could be proved, showing that certain quantum states necessarily have ontic independence with respect to other sets that aren’t necessarily bases.

The techniques of Thm. 2.2 then formed the basis of the rest of the results in Chaps. 2, 3. One of the primary motivations for epistemic realist ontologies with ontic overlaps is to explain the indistinguishability of non-orthogonal quantum states in terms of these overlaps. Much like how Ref. [BCLM14] proved that such overlaps cannot fully explain *all* indistinguishabilities (by a failure of being maximally ψ -epistemic), Thm. 2.3 proved that the indistinguishability between *any pair* of quantum states cannot be fully explained in this way (for $d > 3$ dimensions). Moreover, it was proved that no individual quantum state can be maximally ψ -epistemic at all. By considering the effect of higher dimensions, this was adapted in Thm. 2.4 to prove that many pairs of identifiable quantum states have necessarily small ontic overlap in large dimensions. In particular, as dimension $d \rightarrow \infty$ the ontic overlap between many such pairs must approach zero, while maintaining finite Born rule overlap.

These conclusions make it very difficult for the epistemic realist to use ontic overlaps to explain indistinguishability in particular, but also any other phenomena that might seem to naturally gel with ontic overlaps, such as no-cloning for example. The situation becomes very difficult in large-dimensional systems, where the ontic overlaps between any pair of quantum states satisfying $|\langle\phi|\psi\rangle|^2 < \frac{1}{4}$ must be very small (by Thm. 2.4). Importantly,

these results apply to many *identifiable* pairs of quantum states, *viz.* you give me a pair of quantum states and I can tell you whether the results apply to that pair in particular. This closes an important loophole common to previous similar theorems [Mar12, BCLM14, Bra14, Lei14b, Bal14], which only proved the existence of pairs of quantum states with bounded overlaps. Theorem 2.4 also addressed another shortcoming of those previous theorems, in that its limiting case applies to states with finite Born rule overlap, not states that approach orthogonality (as discussed in Sec. 2.1.4).

These results improve on the current state of no-go theorems for epistemic realist ontologies in several ways, but there are plenty of opportunities for further improvement. Perhaps the ultimate aim for such results would be a proof that any ontological model for a quantum system must be *sometimes ψ -ontic* [Lei14a]. An ontological model is *sometimes ψ -ontic* if for each quantum state there is some finite-measure set of ontic states that can only be obtained by preparing that particular quantum state—*i.e.* each quantum state keeps a region of the ontic state space to itself. This is very similar to ψ -ontic, but is not ruled out by the existence of ψ -epistemic ontological models [LJBR12, ABCL13]. Perhaps the greatest advantage to such a proof would be that it implies several other important foundational results, including Bell’s theorem, and therefore could act to unify those results [Lei14a].

Chapter 3 shifted focus to macro-realism—another possible property of quantum state ontology. The main result was Thm. 3.1, which used Thm. 2.2 to prove that two of three types of macro-realism, ESMR and EMMR, are incompatible with quantum theory in $d > 3$ dimensions. This is more powerful than the Leggett-Garg argument, which is only able to rule out EMMR, while both leave the possibility of a third type of macro-realism, SSMR. Since Bohmian mechanics reproduces quantum predictions and satisfies SSMR, then no theorem can prove that quantum theory is incompatible with every SSMR ontology. However, Bohmian mechanics is ψ -ontic (a very restrictive condition) and it therefore may still be possible for further work to prove that quantum theory is incompatible with all ψ -epistemic (opposite of ψ -ontic) SSMR ontologies. More generally, it would be good for further work to clarify whether any significant subset of SSMR ontologies is also incompatible with quantum theory.

These main results are all unfortunately intolerant to error. That is, the proofs do not directly generalise to the case where quantum probabilities are only assumed to be approximately correct. The core reason for error-intolerance is that they are all based on the asymmetric overlap, a quantity that loses meaning in the presence of finite error. To address this, Thm. 2.5 adapted the proof Thm. 2.4 to use the alternative symmetric overlap, resulting in a somewhat weaker yet crucially error-tolerant result. The conclusion is that

an appropriate set of experiments that reproduce quantum predictions to within some small $\pm\epsilon$ could in principle be used to experimentally bound the ontic overlap between an appropriate specific pair of quantum states.

However, Thm. 2.5 is something of a proof-of-concept result, as the proof is not very well matched to the symmetric overlap and the resulting error term does not scale very well. An alternative approach was taken in Sec. 3.4 where the ϵ -asymmetric overlap was introduced as an error-tolerant generalisation of the asymmetric overlap. This was used to prove Thm. 3.2, which can be used to state error-tolerant variations on Thms. 2.2, 3.1. Similarly, the conclusion is that appropriate sets of experiments could be used to rule out epistemic superpositions and ESMR/EMMR macro-realism to within given precisions.

Both of these error-tolerant results show the way for further experimental work in this area. However in order to achieve that, a more thorough analysis of the exact measurements needed for each would be required. The introduction of the ϵ -asymmetric overlap also suggests that other results using the asymmetric overlap could be made error-tolerant and amenable to experimental investigation in this way, including overlap results already present in the literature [Bal14, LM13, Mar12]. An experimental refutation of ESMR macro-realist ontologies would be particularly useful, as approaches based on the Leggett-Garg argument are incapable of achieving this.

Section 2.5 applied these methods to information theory and classical simulations of quantum channels in particular. The result was Thm. 2.6, proving that to perfectly simulate an n_q -qubit noiseless quantum channel with one-way classical communication a noiseless classical channel of at least $2^{n_q + \mathcal{O}(1)} - 1$ bits is needed, even when the sender and receiver have access to arbitrarily large shared random data. This may be seen as a kind of “anti-Holevo” bound [Sec. 2.5.2], since while the Holevo bound requires n_c qubits to store n_c classical bits this requires $\mathcal{O}(2^{n_q})$ classical bits to store n_q qubits. The bound proved in Thm. 2.6 asymptotically matches the best known bounds for the same simulation but, as noted in Sec. 2.5.3, has three key advantages. First, it is much simpler to prove. Second, because of its simplicity it may be easily seen as a result of a certain property of quantum states. Most importantly, third, the proof method can be reused to produce potentially better bounds, given a certain class of classical error-correction codes.

Clearly then, a good avenue for further work is finding examples of classical error-correcting codes to give better and more explicit bounds via Thm. 2.6. In particular, an explicit family of such codes would produce a bound of the form $c2^{n_q} - 1$ for some specific c . If $c > 0.293$, this bound would exceed the best known specific bound [Mon11].

More generally, Thm. 2.6 demonstrates the power of a fact that, to my knowledge, has not been utilised in quantum information before. That is, the existence of sets of

quantum states exponentially large in dimension where every triple from that set is anti-distinguishable. Given the three-way incompatibility displayed by anti-distinguishable triples, it seems likely that this fact may become useful in the study of quantum cryptography.

Recall from Sec. 2.1.4 that the current generation of ontology theorems concentrate on single systems to avoid the issues raised by Bell's theorem and the PBR theorem in multipartite environments. In particular, the preparation independence postulate (PIP) is assumed in the PBR theorem, which uses it to prove that all compatible ontological models are ψ -ontic. It may be interesting to revisit some of these multipartite issues in light of the quantum causal models of Chap. 5. Bell locality and the PIP are both conditions on the way that the ontologies of multiple quantum systems can combine into a global ontology, while causal models naturally describe how local independent systems can interact in general, so it seems likely that some properties of the causal models framework might influence how one treats multiple systems in ontological models.

7.2 Quantum Causality

The appropriate way to describe and analyse causality in quantum theory was the subject of Chaps. 4, 5. In Chap. 4, a definition for quantum conditional independence was given and motivated, enabling a definition for a quantum Reichenbach's principle and a characterisation of quantum common causes. These were used in Chap. 5 to motivate and give a corresponding definition for full quantum causal models, capable of describing any acyclic causal scenario.

Quantum conditional independence has four equivalent definitions, enumerated in Thm. 4.2. The particular strength of this quantum conditional independence comes from the fact that each of these definitions is a natural generalisation of a corresponding definition for classical conditional independence. While the first of these (conditions (1) and (2)) were obtained in Sec. 4.3.2 by assuming fundamentally unitary dynamics, there are many other ways of obtaining the same quantum conditional independence from different classical starting points. For example, an information-theoretic approach may obtain condition (3) as the first definition of quantum conditional independence. Regardless of one's philosophical predisposition, however, the fact that so many natural definitions are equivalent gives them all strength.

Condition (4) of Thm. 4.2 is of particular interest. In Sec. 4.3.5, it was informally argued that channels satisfying this condition can represent two agents acting independently on a single input system. An important contribution would be to formalise this concept. That is,

from a principled (perhaps operational) definition of what it means to “act independently” on a common input, prove that all channels that do so are of the form required by condition (4). This would provide another good way of justifying quantum conditional independence via condition (4). Assuming this is possible, it would establish channels of this form as very important generalisations of factorised quantum channels, potentially having important uses in quantum information, quantum theory in relativistic causal settings, and beyond.

In Sec. 4.3.5, a new symbol, \oplus , was added to quantum circuit diagrams to concisely depict channels satisfying condition (4). There is much potential for formalising the role of \oplus in quantum circuits, as well as other diagrammatic formulations of quantum theory [CK17, CDP10, CDP11, CDP16, Har01, Har11]. This would be especially useful given a full characterisation of these channels as discussed above. Moreover, simply deriving rules for how diagrams involving \oplus can be manipulated could greatly simplify many calculations in quantum causal models.

Quantum conditional independence characterises channels where the input can act as a complete common cause for the outputs. This enables a definition of quantum Reichenbach’s principle. Classically, the framework of causal models may be seen as a generalisation of Reichenbach’s principle. By way of analogy, Sec. 5.3 proposed a definition of quantum causal models that most simply generalises quantum Reichenbach and then illustrated its utility with examples.

However, as noted in Sec. 5.3.1, it would be better to justify these causal models by an argument from fundamental unitarity, mirroring the justification for quantum conditional independence given in Sec. 4.3.2. Filling this gap is an important piece of further work, one that would no doubt benefit from a more thorough understanding of the behaviour of \oplus . Not only that, it would be of particular relevance to those who believe that quantum dynamics *is* fundamentally unitary just as the argument in Sec. 4.3.2 is.

In quantum causal models, the process of “linking out” nodes was introduced in Sec. 5.3.3 as the quantum causal model analogue of marginalising over nodes in a classical causal model—that is, the mathematical procedure corresponding to ignoring that node. As a first step to being able to derive more properties of quantum causal models, including in applications to specific experiments, it would be very useful to investigate the properties of this linking out operation. In particular, a general specification for how linking out affects the model state, the causal structure, and the relationship between them should be found. One simple case is for linking out a leaf node (that is, a node with no children), where it is easy to check that the resulting model state is simply obtained by removing the factor corresponding to that node from Eq. (5.2). A corresponding specification for linking out a general node would be more difficult to find, but potentially much more useful.

Beyond these opportunities for further work on the formalism of quantum causal models itself, there are also many ways in which they might be used and extended.

An important use case for quantum causal models is in the analysis of experiments. Since the output of experiments is, generally speaking, statistical data, a first step in this direction would be a thorough analysis of what characterises statistics obtained from quantum experiments with certain causal structures. Given a set of experimental data, it is known how to easily check whether a certain variable could represent a complete common cause of another pair of variables, for example. Developing similar techniques based on quantum causal models would go a long way towards bringing them to experimental relevance.

Another use case for quantum causal models is in bringing a causal understanding to certain “paradoxes” or other puzzling thought experiments in quantum foundations. For example, the *extended Wigner’s friend* “paradox” [Wig67, Deu85, Bru17, FR16] has a branched causal structure and a foundational analysis of it may benefit from a causal understanding facilitated by quantum causal models, since the exact disanalogy between the classical and quantum cases remains somewhat unclear [FR16, BHW16].

One powerful way to extend the quantum causal models framework would be to find a d-separation theorem, or something similar [Pea09, SGS01, HLP14]. Classically, d-separation is a graphical criterion that applies to causal structures and the d-separation theorem establishes it as both sound and complete for conditional independence in a Markov causal model. This is one of the foundational results in the study of causal models and the basis of many causal discovery algorithms [Pea09, SGS01]. If a similar enough theorem were found for quantum causal models, then it may be possible to use some of these algorithms already derived in the classical setting with minimal changes. Techniques used in Ref. [HLP14] to prove a d-separation theorem for their version of quantum causal models could be useful for developing such an extension.

Finally, it is interesting to consider whether the techniques of Chaps. 4, 5 could be used as a template for developing formulations of causal models in frameworks for beyond-quantum physical theories. One example is the *general probabilistic theory* (GPT) framework [Har01, Har11, CDP10, CDP11], which is a general way to discuss a wide landscape of theories including classical and quantum. Such a GPT causal model formalism would not only be useful for the study of GPTs themselves, but also could potentially unify quantum and classical causal models. As noted in Sec. 5.1.3, Refs. [HLP14, Fri16] present alternative formulations of quantum causal models (inequivalent to that presented here) that apply to broader frameworks in this way and may be useful resources for this effort.

7.3 Time Travel and CTCs

The final approach to ontology in quantum theory taken in this thesis was in Chap. 6, which considered time travel to the past in a quantum universe. This was done by taking the quantum circuit approach, which has yielded the models of D-CTCs and P-CTCs. After reviewing this previous work in Sec. 6.3, the significance of ontology and non-linearity in these models was discussed in Secs. 6.3.3, 6.3.4. This led to the identification of two classes of new models for quantum time travel in Sec. 6.4.3: weighted D-CTCs and transition probability CTCs. From these, a specific transition probability model called T-CTCs was selected, thoroughly investigated, and compared to the previous models.

Both P-CTCs and T-CTCs make use of an operator $P \stackrel{\text{def}}{=} \text{Tr}_{\text{CV}}(U)$ which is the partial trace of a unitary operator [Eq. (6.5)]. Any further work that develops a fuller understanding of the properties of such operators would be in a good place to quickly identify further properties of P-CTCs and T-CTCs. The only general property of such operators so far proved is a tight upper bound on the action of P on vector norms [Sec. 6.4.5.2].

The original paper on D-CTCs discussed how the second law of thermodynamics remains respected in that model (except within the chronology violating region) [Deu91]. No such analysis has yet appeared for either the P-CTC or T-CTC models. This analysis would be interesting for further comparing the properties of these models and may have a bearing on their relative plausibilities.

Of the opportunities for further work based into the properties of D-, P-, and T-CTCs, perhaps the most intriguing is using them to clarify the relationships between non-linearity, computation, and distinguishability of quantum states. In particular, a well-known result in Ref. [AL98] claims that “virtually any” non-linear extension to quantum theory is able to solve NP-complete and $\#P$ problems in polynomial time; however, the method used in the proof of that result does not apply to the non-linear evolutions provided by T-CTCs. While T-CTCs can solve any $PP \supseteq NP$ problem in polynomial time and does therefore not provide a counter-example, it does call for a more thorough analysis of this claim. Reference [Aar13, Chap. 9] suggests, in a conversational manner, that deterministic distinguishability of states is a necessary consequence of non-linearity. However, it has already been established that T-CTCs, while generally non-linear, cannot distinguish non-orthogonal quantum states in a single shot with certainty. In general, results such as these demonstrate a close relationship between non-linearity, computational power, and distinguishability of non-orthogonal states, but one that is not yet fully understood. It is clear that non-linearity is necessary for the latter two, but further study of time travel models may help to clarify the degree to which it is also sufficient.

Beyond these open questions for the study of D-, P-, and T-CTCs, this also suggests fruitful further work for the study of time travel beyond these models.

Perhaps the most obvious avenue for work beyond these models is to identify other particular weighted D-CTCs or transition probability CTCs that may warrant further investigation. A few of these are noted in Secs. 6.4.3, 6.4.5.6 but none other than T-CTCs have received a full treatment. This would be especially interesting if a convincing ontological argument could be given for a particular model.

In Sec. 6.3.4, intrinsic stochasticity and non-separability were identified as features of quantum theory that make it particularly difficult to find an ontologically sound model of time travel, compared to the classical context. The precise roles of these features could be clarified by starting with a generic classical model with time travel and incrementally adding stochastic and non-separable features. By starting from the ontologically uncontroversial classical setting, this could reveal concrete suggestions of how to consistently model quantum time travel. A complementary task would be to develop a model of time travel for Spekkens' toy model [Spe07] (since that is a toy model for quantum theory that has an explicit ontology) and use it as an inspiration for time travel models in quantum theory.

The analysis of computation in the presence of time travel has potentially interesting implications for computation in generalised probabilistic theories (GPTs) [LB15, BBHL17]. Since the GPT framework is designed to characterise an exceedingly large class of theories that might be applicable to physics, it seems likely that D-, P-, and T-CTC models should be expressible as GPTs. Expressing them in this way should help with comparisons between them, to classical theories, and to standard quantum theory. In particular, a general framework of computation in GPTs is developed in Ref. [LB15] using several plausible assumptions, including those of “tomographic locality” and a “uniformity condition”. There it was proved that no GPT computation in that framework is capable of solving problems outside $AWPP \subseteq PP$ in polynomial time. This is puzzling in light of the known results about D-CTCs, which are capable of solving all problems in PSPACE in polynomial time, since it is believed that PSPACE is strictly larger than PP [AB09] (and therefore also AWPP). This puzzle reveals a tension between the model of GPT computation of Ref. [LB15] and the model of computation used in Ref. [AW09] to prove that D-CTCs have the computational power of PSPACE.

By explicitly expressing D-, P-, and T-CTCs as GPTs it should be possible to compare them to the assumptions of Ref. [LB15]. Doing this would certainly clarify the discrepancy between the results of Ref. [AW09] and Ref. [LB15] and may even suggest ways in which they might be improved. *A priori*, D-, P-, and T-CTCs seem like they should all define

plausible enough GPTs which have similarly plausible notions of computation. Therefore, it would be interesting to see, especially in the case of D-CTCs, whether the assumptions used in Ref. [LB15] hold for these theories, and whether any subtleties are brought to light. In particular, it is not obvious whether the GPTs corresponding to D-, P-, or T-CTCs would be either tomographically local or satisfy the uniformity condition. In general, the study of GPTs could benefit from some more fully worked-out examples and time travel models could benefit from some new settings in which to examine them.

Finally, the framework of quantum causal models developed in Chap. 5 could perhaps also be of use for in the study of quantum time travel. The stipulation of time travel to the past is, after all, a causal one. This will not be as simple as writing a causal model for a CTC, since causal models rule out such causal loops by fiat. However, certain intuitions and understandings from causal models, including the characterisation of channels acting independently on a single system, may be useful in trying to understand smaller parts of the time travel process.

7.4 To Conclude

It is undoubtedly ambitious to attempt to describe what “fundamental reality” can and cannot be like. The way to even partially achieve this, while not losing your head or embarrassing yourself too much, is to be careful and precise with your statements and arguments. This makes concluding a bothersome business, as substantial statements tend to need so many caveats and conditions you might as well just re-state the theorems. Perhaps it is best then to simply say this. If you want real states of affairs respecting probability, then you probably need to accept a certain similarity to quantum states (including superpositions). If you want to correctly describe causal influences, then you need to understand quantum conditional independence and all that it implies. If you want to make sense of time travel to the past, then there are many ontological choices to be made and a wide variety of models to explore. If you want to know what quantum theory means for ontology, then there is plenty more work to be done.

Appendix A

Proofs of Theorems 2.3 and 2.4

Restated Theorem 2.3. Consider a $d > 3$ dimensional quantum system and any pair of quantum states $|\psi\rangle, |0\rangle \in \mathcal{P}(\mathcal{H})$ that satisfy $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$. Assume that preparations of $|\psi\rangle$ are non-contextual with respect to stabiliser unitaries of $|\psi\rangle$. For any preparation measure $\mu \in \Delta_{|\psi\rangle}$, the asymmetric overlap is bounded by

$$\varpi(|0\rangle | \mu) \leq |\langle 0|\psi\rangle|^2 \left(\frac{1}{2} + |\langle 0|\psi\rangle|^2 \right) < |\langle 0|\psi\rangle|^2. \quad (\text{A.1})$$

In particular, this implies that the asymmetric overlap $\varpi(|0\rangle | \mu)$ is strictly less than maximal. Therefore $|\psi\rangle$ is not maximally ψ -epistemic with respect to $|0\rangle$ for any such pair of states in $d > 3$ dimensions satisfying $|\langle 0|\psi\rangle|^2 \in (0, \frac{1}{2})$.

Proof. Define an orthonormal basis \mathcal{B} that contains $|0\rangle$ such that

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \tau|2\rangle, \quad (\text{A.2})$$

where $\alpha \in (0, \frac{1}{\sqrt{2}})$ and $\beta \stackrel{\text{def}}{=} \sqrt{2}\alpha^2$. This is always possible since a choice of global phase for $|0\rangle$ can ensure $\alpha \in \mathbb{R}^+$ and since $|\alpha|^2 + |\beta|^2 = \alpha^2(1 + 2\alpha^2) < 1$ for every such α . With respect to the same \mathcal{B} define

$$|\phi\rangle \stackrel{\text{def}}{=} \delta|0\rangle + \eta|1\rangle + \kappa|3\rangle, \quad (\text{A.3})$$

where $\delta \stackrel{\text{def}}{=} 1 - 2\alpha^2$ and $\eta \stackrel{\text{def}}{=} \sqrt{2}\alpha$. This is always possible since $|\delta|^2 + |\eta|^2 = (1 - 2\alpha^2)^2 + 2\alpha^2 < 1$ for all α and thus an appropriate κ always exists.

The above construction has been chosen such that

- $|\langle 0|\psi\rangle|^2 = \alpha^2 = |\langle \phi|\psi\rangle|^2$ so that there exists a stabiliser unitary $U \in \mathcal{S}_{|\psi\rangle}$ for which $U|0\rangle = |\phi\rangle$; and
- the inner products $|\langle 0|\psi\rangle|^2$, $|\langle \phi|\psi\rangle|^2$, $|\langle 0|\phi\rangle|^2$ satisfy Eq. (2.29) and, therefore, the triple $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ is anti-distinguishable.

For any preparation measure $\mu \in \Delta_{|\psi\rangle}$ consider $\varpi(|0\rangle | \mu)$. For the unitary U , this satisfies $\varpi(|\phi\rangle | \mu') \geq \varpi(|0\rangle | \mu)$ by Lem. 2.3, where $\mu \xrightarrow{\gamma} \mu'$ for some $\gamma \in \Gamma_U$. Moreover, the assumption of non-contextuality with respect to stabiliser unitaries of $|\psi\rangle$ gives $\mu = \mu'$ and so

$$\varpi(|\phi\rangle | \mu) \geq \varpi(|0\rangle | \mu). \quad (\text{A.4})$$

Consider a preparation of the state $|\psi\rangle$ via μ followed by a basis measurement M in the \mathcal{B} basis. Since $|0\rangle$ and $|1\rangle$ are the only outcomes compatible with preparations of both $|\psi\rangle$ and either $|0\rangle$ or $|\phi\rangle$, then by Lem. 2.5 $\mathbb{P}_M(|0\rangle \vee |1\rangle | |\psi\rangle) \geq \varpi(|0\rangle, |\phi\rangle | \mu)$. Further, the anti-distinguishability of $\{|\psi\rangle, |\phi\rangle, |0\rangle\}$ and Lem. 2.6 gives

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | |\psi\rangle) \geq \varpi(|0\rangle, |\phi\rangle | \mu) = \varpi(|0\rangle | \mu) + \varpi(|\phi\rangle | \mu) \quad (\text{A.5})$$

$$\geq 2\varpi(|0\rangle | \mu) \quad (\text{A.6})$$

where the final line comes from Eq. (A.4).

In order to reproduce quantum predictions and satisfy Eq. (2.3)

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | |\psi\rangle) = |\langle 0|\psi\rangle|^2 + |\langle 1|\psi\rangle|^2 = \alpha^2 + 2\alpha^4. \quad (\text{A.7})$$

Combining this with Eq. (A.6), one obtains

$$\varpi(|0\rangle | \mu) \leq \alpha^2 \left(\frac{1}{2} + \alpha^2 \right), \quad (\text{A.8})$$

which is the desired result. \square

Restated Theorem 2.4. *Consider a $d > 3$ dimensional quantum system and any pair $|\psi\rangle, |0\rangle \in \mathcal{P}(\mathcal{H})$ that satisfy $\alpha \stackrel{\text{def}}{=} |\langle 0|\psi\rangle| \in (0, \frac{1}{4})$. Assume that preparations of $|\psi\rangle$ are non-contextual with respect to stabiliser unitaries of $|\psi\rangle$. For any preparation measure $\mu \in \Delta_{|\psi\rangle}$, the asymmetric overlap must satisfy*

$$\varpi(|0\rangle | \mu) \leq \alpha^2 \left(\frac{1 + 2\alpha}{d - 2} \right) \quad (\text{A.9})$$

$$\lim_{d \rightarrow \infty} \varpi(|0\rangle | \mu) = 0 \quad (\text{A.10})$$

and so becomes arbitrarily small as d increases even as α is held constant.

Proof of Thm. 2.4. The proof strategy closely follows that of Thm. 2.3, but modified to take advantage of higher dimensions.

Any such $|\psi\rangle$ may be written

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \tau|2\rangle, \quad (\text{A.11})$$

where $\beta = \sqrt{2}\alpha^{\frac{3}{2}}$ in some orthonormal basis \mathcal{B} containing $|0\rangle$. This is because the global phase of $|0\rangle$ may be chosen to keep $\alpha \in \mathbb{R}^+$ and since the above can always be normalised by an appropriate τ .

Define a set of states $\{|\phi_i\rangle\}_{i=3}^{d-1}$ with respect to the same basis \mathcal{B} such that

$$|\phi_i\rangle \stackrel{\text{def}}{=} \delta|0\rangle + \eta|1\rangle + \kappa|i\rangle, \quad (\text{A.12})$$

with $\delta \stackrel{\text{def}}{=} 1 - 2\alpha^2$ and $\eta \stackrel{\text{def}}{=} \sqrt{2}\alpha^{\frac{3}{2}}$. Again, this is possible since an appropriate κ can always be found to normalise each $|\phi_i\rangle$. Note how the definitions of β and η are different from the construction used in the proof of Thm. 2.3.

It is easy to verify, using Eq. (2.29), that the triples $\{|0\rangle, |\psi\rangle, |\phi_i\rangle\}$ and $\{|\psi\rangle, |\phi_i\rangle, |\phi_j\rangle\}$ are anti-distinguishable for all $i \neq j$. Moreover, $|\langle\phi_i|\psi\rangle|^2 = \alpha^2 = |\langle 0|\psi\rangle|^2$ for all i , so there exist stabiliser unitaries $\{U_i\}_{i=3}^{d-1} \subset \mathcal{S}_{|\psi\rangle}$ for which $U_i|0\rangle = |\phi_i\rangle$. By assumption of non-contextuality with respect to stabiliser unitaries of $|\psi\rangle$, μ remains unchanged by any action of any U_i : $\mu \xrightarrow{\gamma} \mu$ for any $\gamma \in \Gamma_{U_i}$. Therefore, by Lem. 2.3

$$\varpi(|\phi_i\rangle | \mu) \geq \varpi(|0\rangle | \mu) \quad \forall i. \quad (\text{A.13})$$

Consider a preparation of the state $|\psi\rangle$ via μ followed by a basis measurement M in the \mathcal{B} basis. Since $|0\rangle$ and $|1\rangle$ are the only outcomes compatible with preparations of both $|\psi\rangle$ and any state from $\{|0\rangle\} \cup \{|\phi_i\rangle\}_i$, then by an easy generalisation¹ of Lem. 2.5 $\mathbb{P}_M(|0\rangle \vee |1\rangle | |\psi\rangle) \geq \varpi(|0\rangle, |\phi_3\rangle, \dots, |\phi_{d-1}\rangle | \mu)$. Further, using that $\{|\psi\rangle, |0\rangle, |\phi_i\rangle\}$ and $\{|\psi\rangle, |\phi_i\rangle, |\phi_j\rangle\}$ are anti-distinguishable for all $i \neq j$, a similar generalisation² of Lem. 2.6

¹In Lem. 2.5, it is shown that $\mathbb{P}(|0\rangle \vee |1\rangle | |\psi\rangle) \geq \varpi(|0\rangle, |\phi_i\rangle | \mu)$, for example. By following the same proof through using the whole set $\{|0\rangle\} \cup \{|\phi_i\rangle\}_i$, one obtains this more general result. Intuitively, they are getting at the same idea, since any of the quantum states in this set can only return outcomes $|0\rangle$ or $|1\rangle$ in a measurement of \mathcal{B} .

²Intuitively, this is easy to see. Anti-distinguishability of a triple, roughly, requires that no ontic state can be prepared by all three quantum states by Lem. 2.6. There are no true regions of ‘‘tripartite overlap’’. So if all triples are anti-distinguishable, then there can be no tripartite overlap regions in the ontic state space, nor any regions with more-than-two quantum states overlapping, giving the result quoted in the text. The full proof of this simply generalises from that of Lem. 2.6, but becomes very long without introducing any new ideas and so has been omitted.

gives that

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | |\psi\rangle) \geq \varpi(|0\rangle, |\phi_3\rangle, \dots, |\phi_{d-1}\rangle | \mu) \quad (\text{A.14})$$

$$= \varpi(|0\rangle | \mu) + \sum_{i=3}^{d-1} \varpi(|\phi_i\rangle | \mu) \quad (\text{A.15})$$

$$\geq (d-2)\varpi(|0\rangle | \mu). \quad (\text{A.16})$$

The third line follows from Eq. (A.13).

In order to reproduce quantum predictions and satisfy Eq. (2.3)

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | |\psi\rangle) = |\langle 0|\psi\rangle|^2 + |\langle 1|\psi\rangle|^2 = \alpha^2 + 2\alpha^3. \quad (\text{A.17})$$

Combining this with Eq. (A.16), one obtains

$$\varpi(|0\rangle | \mu) \leq \alpha \left(\frac{1+2\alpha}{d-2} \right), \quad (\text{A.18})$$

which is the desired result. □

Appendix B

Proof of Theorem 2.5

The proof presented here uses the simpler non-measure-theoretic version of ontological models that is often seen in the literature. This is the same as the version of ontological models defined in Sec. 2.1.1 except rather than preparation *measures*, preparation *distributions* are used. This change has been made in order to make the proof shorter and easier to read. There should be no conceptual hurdles to extending this proof to the more general ontological models by using the methods of other proofs in the thesis, the hurdles are simply practical.

To be specific, in this appendix it is assumed that there is some canonical measure $d\lambda$ over Λ so that for preparation measure μ , $d\mu(\lambda) = \mu(\lambda) d\lambda$ is the probability that an ontic state is prepared from some infinitesimal neighbourhood of λ . Sets $\Delta_{|\psi\rangle}$ contain these preparation distributions just as they held preparation measures. Other than this, the notation and treatment of ontological models remains the same.

Proof. This proof uses the the assumptions, notation, and constructions from Thm. 2.4 and its proof in Appendix A, except this time it is only assumed that the ontological model reproduces quantum probabilities to within some additive error $\epsilon \in (0, 1]$, as in Eq. (2.4). It will also be necessary to define the tripartite symmetric overlap between three probability distributions μ, ν, χ [Lei14a]

$$\omega(\mu, \nu, \chi) \stackrel{\text{def}}{=} \int_{\Lambda} d\lambda \min\{\mu(\lambda), \nu(\lambda), \chi(\lambda)\}. \quad (\text{B.1})$$

Consider any pair of preparation distributions $\mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|0\rangle}$. From the proof of Thm. 2.4 in Appendix A it is known that there exist $U_i \in \mathcal{S}_{|\psi\rangle}$ such that $U_i|0\rangle = |\phi_i\rangle$. For

each U_i pick any corresponding stochastic map $\gamma_i \in \Gamma_{U_i}$. By assumption, preparations of $|\psi\rangle$ are non-contextual with respect to stabiliser unitaries so each γ_i maps μ to itself, $\mu \xrightarrow{\gamma_i} \mu$. Each γ_i defines some $\chi_i \in \Delta_{|\phi_i\rangle}$ such that $\nu \xrightarrow{\gamma_i} \chi_i$. For notational convenience, let $|\phi_0\rangle \stackrel{\text{def}}{=} |0\rangle$ and $\chi_0 \stackrel{\text{def}}{=} \nu$ then define the index sets $\tilde{I} \stackrel{\text{def}}{=} \{3, \dots, d-1\}$ and $I \stackrel{\text{def}}{=} \{0\} \cup \tilde{I}$.

Recall that ω is defined as twice the probability of making a mistake when attempting to distinguish between measures when using the optimal strategy. Clearly, if some stochastic map γ acts as $\mu \xrightarrow{\gamma} \mu'$ and $\nu \xrightarrow{\gamma} \nu'$, then any way of distinguishing μ' and ν' is also a way to distinguish μ and ν . Therefore, symmetric overlaps cannot decrease under the action of stochastic maps, just as with asymmetric overlaps in Lem. 2.3. In terms of the above definitions, this therefore implies

$$\omega(\mu, \chi_i) \geq \omega(\mu, \nu), \quad \forall i \in \tilde{I}. \quad (\text{B.2})$$

Consider a preparation of $|\psi\rangle$ via μ , followed by a measurement M in the basis \mathcal{B} . Similarly to Thm. 2.4, the aim is to bound $\omega(\mu, \nu)$ by considering the probability of obtaining either of the measurement outcomes $|0\rangle$ or $|1\rangle$, given by Eq. (2.4) as

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) \leq \alpha^2 + \beta^2 + \epsilon. \quad (\text{B.3})$$

The trick is to do this in such a way that all possible errors are accounted for.

In order to link this quantum probability to symmetric overlaps, consider the following subsets of Λ .

- For each $i \in I$ consider $\Omega_i \stackrel{\text{def}}{=} \{\lambda \in \Lambda : 0 < \mu(\lambda) \leq \chi_i(\lambda)\}$. Roughly, Ω_i is the region of the overlap between μ and χ_i for which μ is smaller than χ_i .
- For each $i \in I$ consider $\Theta_i \stackrel{\text{def}}{=} \{\lambda \in \Lambda : 0 < \chi_i(\lambda) < \mu(\lambda); \forall j < i, \chi_j(\lambda) \leq \chi_i(\lambda); \forall j > i, 0 < \chi_j(\lambda) < \chi_i(\lambda)\}$. Roughly, this is the region of the overlap between μ and χ_i for which χ_i is greater than all other $\chi_{j \neq i}$, but smaller than μ .
- For each $i < j \in I$ consider $\Theta_i^j \stackrel{\text{def}}{=} \{\lambda \in \Lambda : 0 < \chi_i(\lambda) \leq \chi_j(\lambda); \chi_i(\lambda) < \mu(\lambda)\}$. Roughly, this is the region of the tripartite overlap of μ, χ_i, χ_j in which χ_i is the minimum of the three.
- Similarly, for each $i > j \in I$ consider $\Theta_i^j \stackrel{\text{def}}{=} \{\lambda \in \Lambda : 0 < \chi_i(\lambda) < \chi_j(\lambda); \chi_i(\lambda) < \mu(\lambda)\}$.
- For every unequal pair $i, j \in I$, let $\Omega_{ij} = \Omega_i \cap \Omega_j$.

Note that these sets are defined to be disjoint, for $i \neq j$: $\Theta_i \cap \Omega_j = \Theta_i \cap \Theta_j = \Theta_i \cap \Theta_i^j = \emptyset$.

The point of these subsets is the way in which they relate to symmetric overlaps. From the definitions of symmetric overlaps it is not difficult to verify that

$$\omega(\mu, \chi_i) = \int_{\Omega_i} d\lambda \mu(\lambda) + \int_{\Theta_i \cup [\cup_{j \neq i} \Theta_i^j]} d\lambda \chi_i(\lambda) \quad (\text{B.4})$$

$$\omega(\mu, \chi_i, \chi_{j \neq i}) = \int_{\Omega_{ij}} d\lambda \mu(\lambda) + \int_{\Theta_i^j} d\lambda \chi_i(\lambda) + \int_{\Theta_j^i} d\lambda \chi_j(\lambda). \quad (\text{B.5})$$

Proceed by separating the probability Eq. (B.3) according to subsets in which λ may obtain. The sets Ω_i collectively capture everywhere in Λ where μ is not the maximum of the measures, while Θ_i collectively capture everywhere that μ is the maximum. So, using a Bonferroni inequality [RS11]

$$\begin{aligned} \mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) &\geq \sum_{i \in I} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_i | \mu) \\ &\quad + \sum_{i \in I} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Theta_i | \mu) \\ &\quad - \sum_{i, j < i} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_{ij} | \mu), \end{aligned} \quad (\text{B.6})$$

$$\begin{aligned} &\geq \sum_{i \in I} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_i | \mu) \\ &\quad + \sum_{i \in I} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Theta_i | \mu) \\ &\quad - \sum_{i, j < i} \int_{\Omega_{ij}} d\lambda \mu(\lambda). \end{aligned} \quad (\text{B.7})$$

The final line follows simply because $\mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_{ij} | \mu) \leq \mathbb{P}_M(\lambda \in \Omega_{ij} | \mu) = \int_{\Omega_{ij}} d\lambda \mu(\lambda)$.

For the $i = 0$ term in the first line of Eq. (B.7), define the function $\xi(\lambda) \stackrel{\text{def}}{=} 1 - \mathbb{P}_M(|0\rangle | \lambda)$ so that

$$\mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_0 | \mu) = \int_{\Omega_0} d\lambda \mu(\lambda) \{ \mathbb{P}_M(|0\rangle | \lambda) + \mathbb{P}_M(|1\rangle | \lambda) \} \quad (\text{B.8})$$

$$\geq \int_{\Omega_0} d\lambda \mu(\lambda) \mathbb{P}_M(|0\rangle | \lambda) \quad (\text{B.9})$$

$$= \int_{\Omega_0} d\lambda \mu(\lambda) - \int_{\Omega_0} d\lambda \mu(\lambda) \xi(\lambda) \quad (\text{B.10})$$

$$\geq \int_{\Omega_0} d\lambda \mu(\lambda) - \int_{\Omega_0} d\lambda \nu(\lambda) \xi(\lambda). \quad (\text{B.11})$$

This can be simplified by noting that, for any $\Omega \subseteq \Lambda$,

$$\int_{\Omega} d\nu(\lambda)\xi(\lambda) = \int_{\Omega} d\lambda \nu(\lambda) - \int_{\Lambda} d\lambda \nu(\lambda)\mathbb{P}_M(|0\rangle|\lambda) \quad (\text{B.12})$$

$$+ \int_{\Lambda \setminus \Omega} d\lambda \nu(\lambda)\mathbb{P}_M(|0\rangle|\lambda) \leq \int_{\Omega} d\lambda \nu(\lambda) - 1 + \epsilon + \int_{\Lambda \setminus \Omega} d\lambda \nu(\lambda) \quad (\text{B.13})$$

$$= \epsilon \quad (\text{B.14})$$

so that

$$\mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_0 | \mu) \geq \int_{\Omega_0} d\lambda \mu(\lambda) - \epsilon. \quad (\text{B.15})$$

The $i \in \tilde{I}$ terms of the first line of Eq. (B.7) simplify in a similar way. Define $\zeta_i(\lambda) \stackrel{\text{def}}{=} 1 - \mathbb{P}_M(|0\rangle|\lambda) - \mathbb{P}_M(|1\rangle|\lambda) - \mathbb{P}_M(|i\rangle|\lambda)$ so that

$$\mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Omega_i | \mu) = \int_{\Omega_i} d\lambda \mu(\lambda) \{\mathbb{P}_M(|0\rangle|\lambda) + \mathbb{P}_M(|1\rangle|\lambda)\} \quad (\text{B.16})$$

$$= \int_{\Omega_i} d\lambda \mu(\lambda) - \int_{\Omega_i} d\lambda \mu(\lambda) \{\zeta_i(\lambda) + \mathbb{P}_M(|i\rangle|\lambda)\} \quad (\text{B.17})$$

$$\geq \int_{\Omega_i} d\lambda \mu(\lambda) - \int_{\Omega_i} d\lambda \chi_i(\lambda)\zeta_i(\lambda) - \epsilon \quad (\text{B.18})$$

$$\geq \int_{\Omega_i} d\lambda \mu(\lambda) - 2\epsilon. \quad (\text{B.19})$$

Together, Eqs. (B.7, B.15, B.19) produce

$$\begin{aligned} \mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) &\geq \sum_{i \in I} \int_{\Omega_i} d\lambda \mu(\lambda) - \sum_{i,j < i} \int_{\Omega_{ij}} d\lambda \mu(\lambda) - (2d - 5)\epsilon \\ &\quad + \sum_{i \in I} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Theta_i | \mu). \end{aligned} \quad (\text{B.20})$$

The $i = 0$ term of the second line of Eq. (B.20) can be bounded as follows

$$\mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Theta_0 | \mu) \geq \int_{\Theta_0} d\lambda \mu(\lambda)\mathbb{P}_M(|0\rangle|\lambda) \quad (\text{B.21})$$

$$\geq \int_{\Theta_0} d\lambda \nu(\lambda)\mathbb{P}_M(|0\rangle|\lambda) \quad (\text{B.22})$$

$$= \int_{\Lambda} d\lambda \nu(\lambda)\mathbb{P}_M(|0\rangle|\lambda) - \int_{\Lambda \setminus \Theta_0} d\lambda \nu(\lambda)\mathbb{P}_M(|0\rangle|\lambda) \quad (\text{B.23})$$

$$\geq 1 - \epsilon - \int_{\Lambda \setminus \Theta_0} d\lambda \nu(\lambda) = \int_{\Theta_0} d\lambda \nu(\lambda) - \epsilon. \quad (\text{B.24})$$

The $i \in \tilde{I}$ terms of the second line of Eq. (B.20) can be similarly bounded

$$\begin{aligned} \mathbb{P}_M(|0\rangle \vee |1\rangle, \lambda \in \Theta_i | \mu) &= \int_{\Theta_i} d\lambda \mu(\lambda) \{ \mathbb{P}_M(|0\rangle | \lambda) + \mathbb{P}_M(|1\rangle | \lambda) \\ &\quad + \mathbb{P}_M(|i\rangle | \lambda) \} - \int_{\Theta_i} d\lambda \mu(\lambda) \mathbb{P}_M(|i\rangle | \lambda) \end{aligned} \quad (\text{B.25})$$

$$\begin{aligned} &\geq \int_{\Theta_i} d\lambda \chi_i(\lambda) \{ \mathbb{P}_M(|0\rangle | \lambda) + \mathbb{P}_M(|1\rangle | \lambda) \\ &\quad + \mathbb{P}_M(|i\rangle | \lambda) \} - \epsilon \end{aligned} \quad (\text{B.26})$$

$$\geq (1 - \epsilon) - \int_{\Lambda \setminus \Theta_i} d\lambda \chi_i(\lambda) - \epsilon \quad (\text{B.27})$$

$$= \int_{\Theta_i} d\lambda \chi_i(\lambda) - 2\epsilon. \quad (\text{B.28})$$

So now combining Eqs. (B.20, B.24, B.28) it is found that

$$\begin{aligned} \mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) &\geq \sum_{i \in I} \left\{ \int_{\Omega_i} d\lambda \mu(\lambda) + \int_{\Theta_i} d\lambda \chi_i(\lambda) \right\} \\ &\quad - \sum_{i, j < i} \int_{\Omega_{ij}} d\lambda \mu(\lambda) - 2(2d - 5)\epsilon. \end{aligned} \quad (\text{B.29})$$

Equation (B.29) can be further reduced by adding any negative quantity. In particular, consider Boole's inequality

$$\sum_{i \in I} \int_{\cup_{j \neq i} \Theta_i^j} d\lambda \chi_i(\lambda) - \sum_{i, j \neq i} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) \leq 0. \quad (\text{B.30})$$

Therefore, Eq. (B.29) reduces to

$$\begin{aligned} \mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) &\geq \sum_{i \in I} \left\{ \int_{\Omega_i} d\lambda \mu(\lambda) + \int_{\Theta_i \cup [\cup_{j \neq i} \Theta_i^j]} d\lambda \chi_i(\lambda) \right\} \\ &\quad - \sum_{i, j < i} \int_{\Omega_{ij}} d\lambda \mu(\lambda) - \sum_{i, j \neq i} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) \\ &\quad - 2(2d - 5)\epsilon. \end{aligned} \quad (\text{B.31})$$

This can be further simplified by noting

$$\sum_{i, j \neq i} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) = \sum_{i, j < i} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) + \sum_{i, j > i} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) \quad (\text{B.32})$$

$$= \sum_{i, j < i} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) + \sum_{j, i < j} \int_{\Theta_i^j} d\lambda \chi_i(\lambda) \quad (\text{B.33})$$

$$= \sum_{i, j < i} \left\{ \int_{\Theta_i^j} d\lambda \chi_i(\lambda) + \int_{\Theta_j^i} d\lambda \chi_j(\lambda) \right\}, \quad (\text{B.34})$$

so that Eq. (B.31) becomes

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) \geq \sum_{i \in I} \omega(\mu, \chi_i) - \sum_{i, j < i} \omega(\mu, \chi_i, \chi_j) - 2(2d - 5)\epsilon \quad (\text{B.35})$$

having used Eqs. (B.4, B.5).

As a final step, consider how the tripartite symmetric overlaps are bounded by ϵ . Consider the measurement $M' = \{E_{-\psi}, E_{-i}, E_{-j}\}$ which anti-distinguishes $\{|\psi\rangle, |\phi_i\rangle, |\phi_j\rangle\}$, so that

$$\int_{\Lambda} d\mu(\lambda) \mathbb{P}_{M'}(E_{-\psi} | \lambda) \leq \epsilon, \quad (\text{B.36})$$

$$\int_{\Lambda} d\chi_i(\lambda) \mathbb{P}_{M'}(E_{-i} | \lambda) \leq \epsilon, \quad (\text{B.37})$$

$$\int_{\Lambda} d\chi_j(\lambda) \mathbb{P}_{M'}(E_{-j} | \lambda) \leq \epsilon. \quad (\text{B.38})$$

Conservation of probability requires that

$$\mathbb{P}_{M'}(E_{-\psi} | \lambda) + \mathbb{P}_{M'}(E_{-i} | \lambda) + \mathbb{P}_{M'}(E_{-j} | \lambda) = 1 \quad (\text{B.39})$$

for all $\lambda \in \Lambda$. Consider the tripartite symmetric overlap, Eqs. (B.1, B.5), for μ, χ_i, χ_j . Then

$$\begin{aligned} \omega(\mu, \chi_i, \chi_j) &= \int_{\Lambda} d\lambda \min\{\mu(\lambda), \chi_i(\lambda), \chi_j(\lambda)\} [\mathbb{P}_{M'}(E_{-\psi} | \lambda) + \mathbb{P}_{M'}(E_{-i} | \lambda) \\ &\quad + \mathbb{P}_{M'}(E_{-j} | \lambda)] \end{aligned} \quad (\text{B.40})$$

$$\begin{aligned} &\leq \int_{\Lambda} d\mu(\lambda) \mathbb{P}_{M'}(E_{-\psi} | \lambda) + \int_{\Lambda} d\chi_i(\lambda) \mathbb{P}_{M'}(E_{-i} | \lambda) \\ &\quad + \int_{\Lambda} d\chi_j(\lambda) \mathbb{P}_{M'}(E_{-j} | \lambda) \end{aligned} \quad (\text{B.41})$$

$$\leq 3\epsilon. \quad (\text{B.42})$$

Applying Eq. (B.42) to Eq. (B.35), one finds that

$$\mathbb{P}_M(|0\rangle \vee |1\rangle | \mu) \geq \sum_{i \in I} \omega(\mu, \chi_i) - \frac{3}{2}(d-3)(d-2)\epsilon - 2(2d-5)\epsilon \quad (\text{B.43})$$

$$\geq (d-2)\omega(\mu, \nu) - \frac{1}{2}(3d^2 - 7d - 2)\epsilon \quad (\text{B.44})$$

having used Eq. (B.2). Combining Eqs. (B.3, B.44) one obtains an upper bound on $\omega(\mu, \nu)$ for any $\mu \in \Delta_{|\psi\rangle}, \nu \in \Delta_{|0\rangle}$, which must be greater than or equal to the least upper bound, Eq. (2.64), finally giving

$$\omega(|\psi\rangle, |0\rangle) \leq \alpha^2 \left(\frac{1+2\alpha}{d-2} \right) + \frac{(3d^2 - 7d)}{2(d-2)} \epsilon. \quad (\text{B.45})$$

This completes the proof. \square

A note is in order about the tightness of this bound, assuming arbitrarily small ϵ . At $d = 4$, this bound cannot improve upon that of Eq. (2.65) for any $\alpha \in (0, \frac{1}{4})$. At $d = 5$, an improvement is possible for some values of α . It is only for $d > 5$ that this bound is capable of improving upon Eq. (2.65) for all values of $\alpha \in (0, \frac{1}{4})$. This is because the theorem extends the methods of Thms. 2.2, 2.4 (which are closely linked to the asymmetric overlap) to the symmetric overlap.

Clearly the error model used here is very simplistic: it has been assumed that some $\epsilon > 0$ can be used to bound the deviation of all probabilities from the quantum predictions. Another source of possible error is in the use of stabiliser unitaries for $|\psi\rangle$. To obtain Eq. (B.44) one uses Eq. (B.2) which requires that the χ_i are obtained from μ by a transformation implementing a stabiliser unitary. Any experiment would also have to engage with the problem of how to account for errors in the implementation of the stabiliser unitary. However, an argument along the lines of Sec. 2.4.4 might also help here.

It may be possible to improve on the error term above by more carefully using higher Bonferroni inequalities than the one used in Eq. (B.6) above. Considering quadpartite and higher-order overlaps (rather than stopping at the tripartite overlap, as done here) may give an improvement in the error term, possibly improving the scaling with d .

Bibliography

- [Aar05] S. Aaronson, *Quantum Computing, Postselection, and Probabilistic Polynomial-time*. Proceedings of the Royal Society A **461**(2063), 3473 (2005). arXiv:quant-ph/0412187 [quant-ph]
- [Aar13] S. Aaronson, *Quantum Computing since Democritus*. Cambridge University Press, Cambridge (2013). <http://www.scottaaronson.com/democritus/>
- [AB09] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*. Cambridge University Press, Cambridge (2009)
- [ABCL13] S. Aaronson, A. Bouland, L. Chua, and G. Lowther, *ψ -Epistemic Theories: The Role of Symmetry*. Physical Review A **88**(3), 032111 (2013). arXiv:1303.2834 [quant-ph]
- [ABH⁺17] J.-M. A. Allen, J. Barrett, D. C. Horsman, C. M. Lee, and R. W. Spekkens, *Quantum Common Causes and Quantum Causal Models*. Physical Review X **In preparation** (2017). arXiv:1609.09487 [quant-ph]
- [ACB14] M. Araújo, F. Costa, and Časlav Brukner, *Computational Advantage from Quantum-Controlled Ordering of Gates*. Physical Review Letters **113**(25), 250402 (2014). arXiv:1401.8127 [quant-ph]
- [AL98] D. S. Abrams and S. Lloyd, *Nonlinear Quantum Mechanics Implies Polynomial-time Solution for NP-complete and #P Problems*. Physical Review Letters **81**(18), 3992 (1998). arXiv:quant-ph/9801041 [quant-ph]
- [Alb10] D. Albert, *Probability in the Everett Picture*. In *Many Worlds? Everett, Quantum Theory, & Reality*, eds. S. Saunders, J. Barrett, A. Kent, and D. Wallace, chap. 11, Oxford University Press, Oxford (2010)

- [Ali95] R. Alicki, *Comment on “Reduced Dynamics Need Not Be Completely Positive”*. Physical Review Letters **75**(16), 3020 (1995)
- [All14] J.-M. A. Allen, *Treating Time Travel Quantum Mechanically*. Physical Review A **90**(4), 042107 (2014). arXiv:1401.4933 [quant-ph]
- [All16] J.-M. A. Allen, *Quantum Superpositions Cannot be Epistemic*. Quantum Studies: Mathematics and Foundations **3**(2), 161–177 (2016). arXiv:1501.05969 [quant-ph]
- [AM13] F. Arntzenius and T. Maudlin, *Time Travel and Modern Physics*. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, winter 2013 ed. (2013)
- [AMG16] J.-M. A. Allen, O. J. E. Maroney, and S. Gogioso, *A Stronger Theorem Against Macro-realism* (2016). arXiv:1610.00022 [quant-ph]
- [AMRM12] D. Ahn, C. R. Myers, T. C. Ralph, and R. B. Mann, *Quantum State Cloning in the Presence of a Closed Timelike Curve* (2012). arXiv:1207.6062 [quant-ph]
- [And95] A. Anderson, *Unitarity Restoration in the Presence of Closed Timelike Curves*. Physical Review D **51**(10), 5707 (1995). arXiv:gr-qc/9405058 [gr-qc]
- [APT14] Y. Aharonov, S. Popescu, and J. Tollaksen, *Each Instant of Time a New Universe*. In *Quantum Theory: A Two-Time Success Story: Yakir Aharonov Festschrift*, eds. D. C. Struppa and J. M. Tollaksen, pp. 21–36, Springer, Berlin (2014). arXiv:1305.1615 [quant-ph]
- [APTV09] Y. Aharonov, S. Popescu, J. Tollaksen, and L. Vaidman, *Multiple-Time States and Multiple-Time Measurements in Quantum Mechanics*. Physical Review A **79**(5), 052110 (2009). arXiv:0712.0320 [quant-ph]
- [AV07] Y. Aharonov and L. Vaidman, *The Two-State Vector Formalism: An Updated Review*. In *Time in Quantum Mechanics*, eds. G. Muga, R. S. Mayato, and I. Egusquiza, pp. 399–447, Springer, Berlin (2007)
- [AW09] S. Aaronson and J. Watrous, *Closed Timelike Curves make Quantum and Classical Computing Equivalent*. Proceedings of the Royal Society A **465**(2102), 631 (2009). arXiv:0808.2669 [quant-ph]
- [Bac15] G. Bacciagaluppi, *Leggett-Garg Inequalities, Pilot Waves and Contextuality*. International Journal of Quantum Foundations **1**, 1–17 (2015). arXiv:1409.4104 [quant-ph]

- [Bal87] L. E. Ballentine, *Realism and Quantum Flux Tunneling*. Physical Review Letters **59**(14), 1493 (1987)
- [Bal14] L. Ballentine, *Ontological Models in Quantum Mechanics: What do they Tell Us?* (2014). arXiv:1402.5689 [quant-ph]
- [BB95] E. G. Beltrametti and S. Bugajski, *A Classical Extension of Quantum Mechanics*. Journal of Physics A **28**(12), 3329 (1995)
- [BBHL17] J. Barrett, N. de Beaudrap, M. J. Hoban, and C. M. Lee, *The Computational Landscape of General Physical Theories* (2017). arXiv:1702.08483 [quant-ph]
- [BCF⁺96] H. Barnum, C. M. Caves, C. A. Fuchs, R. Jozsa, and B. Schumacher, *Noncommuting Mixed States Cannot Be Broadcast*. Physical Review Letters **76**(15), 2818–2821 (1996). arXiv:quant-ph/9511010 [quant-ph]
- [BCLM14] J. Barrett, E. G. Cavalcanti, R. Lal, and O. J. E. Maroney, *No ψ -Epistemic Model Can Fully Explain the Indistinguishability of Quantum States*. Physical Review Letters **112**(25), 250403 (2014). arXiv:1310.8302 [quant-ph]
- [BCM^W10] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf, *Non-locality and Communication Complexity*. Reviews of Modern Physics **82**(1), 665–698 (2010). arXiv:0907.3584 [quant-ph]
- [BCT99] G. Brassard, R. Cleve, and A. Tapp, *Cost of Exactly Simulating Quantum Entanglement with Classical Communication*. Physical Review Letters **83**(9), 1874 (1999). arXiv:quant-ph/9901035 [quant-ph]
- [BCW98] H. Buhrman, R. Cleve, and A. Wigderson, *Quantum vs. Classical Communication and Computation*. In *STOC 98 Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pp. 63–68, ACM (1998). arXiv:quant-ph/9802040 [quant-ph]
- [BCWW01] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, *Quantum Fingerprinting*. Physical Review Letters **87**(16), 167902 (2001). arXiv:quant-ph/0102001 [quant-ph]
- [BD82] N. D. Birrell and P. C. W. Davies, *Quantum Fields in Curved Space*. Cambridge University Press, Cambridge (1982)

- [Bel66] J. S. Bell, *On the Problem of Hidden Variables in Quantum Mechanics*. *Reviews of Modern Physics* **38**(3), 447–452 (1966)
- [Bel87] J. S. Bell, *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*. Cambridge University Press, Cambridge (1987)
- [BGNP01] D. Beckman, D. Gottesman, M. A. Nielsen, and J. Preskill, *Causal and Localizable Quantum Operations*. *Physical Review A* **64**(5), 052309 (2001). arXiv:quant-ph/0102043 [quant-ph]
- [BGP10] C. Branciard, N. Gisin, and S. Pironio, *Characterizing the Nonlocal Correlations Created via Entanglement Swapping*. *Physical Review Letters* **104**(17), 170401 (2010)
- [BH95] D. Bohm and B. J. Hiley, *The Undivided Universe: An Ontological Interpretation of Quantum Theory*. Taylor & Francis (1995)
- [BH13] F. G. S. L. Brandao and A. W. Harrow, *Quantum de Finetti Theorems under Local Measurements with Applications*. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, eds. D. Boneh, T. Roughgarden, and J. Feigenbaum, pp. 861–870, STOC, ACM (2013). arXiv:1210.6367 [quant-ph]
- [BHK05] J. Barrett, L. Hardy, and A. Kent, *No Signaling and Quantum Key Distribution*. *Physical Review Letters* **95**(1), 010503 (2005). arXiv:quant-ph/0405101 [quant-ph]
- [BHL⁺05] C. H. Bennett, P. Hayden, D. W. Leung, P. W. Shor, and A. Winter, *Remote Preparation of Quantum States*. *IEEE Transactions on Information Theory* **51**(1), 56–74 (2005). arXiv:quant-ph/0307100 [quant-ph]
- [BHW09] T. A. Brun, J. Harrington, and M. M. Wilde, *Localized Closed Timelike Curves Can Perfectly Distinguish Quantum States*. *Physical Review Letters* **102**(21), 210402 (2009). arXiv:1209.0811 [quant-ph]
- [BHW16] V. Baumann, A. Hansen, and S. Wolf, *The Measurement Problem is the Measurement Problem* (2016). arXiv:1611.01111 [quant-ph]

- [BK04] J. Barrett and A. Kent, *Non-Contextuality, Finite Precision Measurement and the Kochen-Specker Theorem*. *Studies in History and Philosophy of Science Part B* **35**(2), 151–176 (2004). arXiv:quant-ph/0309017 [quant-ph]
- [BL09] J. Barrett and M. Leifer, *The de Finetti Theorem for Test Spaces*. *New Journal of Physics* **11**, 033024 (2009). arXiv:0712.2265 [quant-ph]
- [BLSS09] C. H. Bennett, D. Leung, G. Smith, and J. A. Smolin, *Can Closed Timelike Curves or Nonlinear Quantum Mechanics Improve Quantum State Discrimination or Help Solve Hard Problems?* *Physical Review Letters* **103**(17), 170502 (2009). arXiv:0908.3023 [quant-ph]
- [BLW15] M. Berta, M. Lemm, and M. M. Wilde, *Monotonicity of Quantum Relative Entropy and Recoverability*. *Quantum Information and Computation* **15**(15–16), 1333–1354 (2015). arXiv:1412.4067 [quant-ph]
- [Boh52a] D. Bohm, *A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. I*. *Physical Review* **85**(2), 166–179 (1952)
- [Boh52b] D. Bohm, *A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. II*. *Physical Review* **85**(2), 180–193 (1952)
- [Bon80] W. B. Bonnor, *The Rigidly Rotating Relativistic Dust Cylinder*. *Journal of Physics A* **13**(6), 2121 (1980)
- [Bra14] C. Branciard, *How ψ -Epistemic Models Fail at Explaining the Indistinguishability of Quantum States*. *Physical Review Letters* **113**(2), 020409 (2014). arXiv:1407.3005 [quant-ph]
- [BRGP12] C. Branciard, D. Rosset, N. Gisin, and S. Pironio, *Bilocal versus Nonbilocal Correlations in Entanglement-Swapping Experiments*. *Physical Review A* **85**(3), 032119 (2012). arXiv:1112.4502 [quant-ph]
- [Bro09] L. de Broglie, *The New Dynamics of Quanta*. In *Quantum Theory at the Crossroads: Reconsidering the 1927 Solvay Conference*, eds. G. Bacciagaluppi and A. Valentini, pp. 373–406, Cambridge University Press, Cambridge (2009), ISBN 9780521814218. arXiv:quant-ph/0609184 [quant-ph]
- [BRS12] S. D. Bartlett, T. Rudolph, and R. W. Spekkens, *Reconstruction of Gaussian Quantum Mechanics from Liouville Mechanics with an Epistemic Restriction*. *Physical Review A* **86**(1), 012103 (2012). arXiv:1111.5057 [quant-ph]

- [Bru17] Āaslav Brukner, *On the Quantum Measurement Problem*. In *Quantum (Un)Speakables II: Half a Century of Bell’s Theorem*, eds. R. A. Bertlmann and A. Zeilinger, pp. 95–117, Springer, Berlin (2017). arXiv:1507.05255 [quant-ph]
- [BS05] C. H. Bennett and B. Schumacher, *Talk on “Teleportation, simulated time travel, and how to flirt with someone who has fallen into a black hole”* (2005). Slides available at <http://web.archive.org/web/20110514013911/http://www.research.ibm.com/people/b/bennetc/QUPONBshort.pdf> accessed 2017-03-02
- [BW12] T. A. Brun and M. M. Wilde, *Perfect State Distinguishability and Computational Speedups with Postselected Closed Timelike Curves*. *Foundations of Physics* **42**(3), 341 (2012). arXiv:1008.0433 [quant-ph]
- [BW17] T. A. Brun and M. M. Wilde, *Simulations of Closed Timelike Curves*. *Foundations of Physics* **47**(3), 375–391 (2017). arXiv:1504.05911 [quant-ph]
- [BWW13] T. A. Brun, M. M. Wilde, and A. Winter, *Quantum State Cloning Using Deutschian Closed Timelike Curves*. *Physical Review Letters* **111**(19), 190401 (2013). arXiv:1306.1795 [quant-ph]
- [BZ02] R. A. Bertlmann and A. Zeilinger, eds., *Quantum (Un)speakables*. Springer, Berlin (2002)
- [BZ06] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge University Press, Cambridge (2006)
- [BZ17] R. A. Bertlmann and A. Zeilinger, eds., *Quantum (Un)Speakables II: Half a Century of Bell’s Theorem*. Springer, Berlin (2017)
- [CD02] M. Czachor and H. D. Doebner, *Correlation Experiments in Nonlinear Quantum Mechanics*. *Physics Letters A* **301**(3–4), 139 (2002). arXiv:quant-ph/0110008 [quant-ph]
- [CDP09] G. Chiribella, G. M. D’Ariano, and P. Perinotti, *Theoretical Framework for Quantum Networks*. *Physical Review A* **80**(2), 022339 (2009). arXiv:0904.4483 [quant-ph]

- [CDP10] G. Chiribella, G. M. D’Ariano, and P. Perinotti, *Probabilistic Theories with Purification*. Physical Review A **81**(6), 062348 (2010). arXiv:0908.1583 [quant-ph]
- [CDP11] G. Chiribella, G. M. D’Ariano, and P. Perinotti, *Informational Derivation of Quantum Theory*. Physical Review A **84**(1), 012311 (2011). arXiv:1011.6451 [quant-ph]
- [CDP16] G. Chiribella, G. M. D’Ariano, and P. Perinotti, *Quantum from Principles*. In *Quantum Theory: Informational Foundations and Foils*, eds. G. Chiribella and R. W. Spekkens, vol. 181 of *Fundamental Theories of Physics*, pp. 171–221, Springer, Berlin (2016). arXiv:1506.00398 [quant-ph]
- [CDPV13] G. Chiribella, G. M. D’Ariano, P. Perinotti, and B. Valiron, *Quantum Computations without Definite Causal Structure*. Physical Review A **88**(2), 022318 (2013). arXiv:0912.0195 [quant-ph]
- [CFS02a] C. M. Caves, C. A. Fuchs, and R. Schack, *Conditions for Compatibility of Quantum-State Assignments*. Physical Review A **66**(6), 062111 (2002). arXiv:quant-ph/0206110 [quant-ph]
- [CFS02b] C. M. Caves, C. A. Fuchs, and R. Schack, *Quantum Probabilities as Bayesian Probabilities*. Physical Review A **65**(2), 022305 (2002). arXiv:quant-ph/0106133 [quant-ph]
- [CFS02c] C. M. Caves, C. A. Fuchs, and R. Schack, *Unknown Quantum States: The Quantum de Finetti Representation*. Journal of Mathematical Physics **43**(9), 4537–4559 (2002). arXiv:quant-ph/0104088 [quant-ph]
- [Cha16] R. Chaves, *Polynomial Bell inequalities*. Physical Review Letters **116**(1), 010402 (2016). arXiv:1506.04325 [quant-ph]
- [Chi12] G. Chiribella, *Perfect Discrimination of No-Signalling Channels via Quantum Superposition of Causal Structures*. Physical Review A **86**(4), 040301 (2012). arXiv:1109.5154 [quant-ph]
- [Cho75] M. D. Choi, *Completely Positive Linear Maps on Complex Matrices*. Linear Algebra and its Applications **10**(3), 285–290 (1975)

- [CK00] R. Clifton and A. Kent, *Simulating Quantum Mechanics by Non-Contextual Hidden Variables*. Proceedings of the Royal Society A **456**(2001), 2101–2114 (2000). arXiv:quant-ph/9908031 [quant-ph]
- [CK15] L. Clemente and J. Kofler, *Necessary and Sufficient Conditions for Macroscopic Realism from Quantum Mechanics*. Physical Review A **91**(6), 062103 (2015). arXiv:1501.07517 [quant-ph]
- [CK16] L. Clemente and J. Kofler, *No Fine Theorem for Macrorealism: Limitations of the Leggett-Garg Inequality*. Physical Review Letters **116**(15), 150401 (2016)
- [CK17] B. Coecke and A. Kissinger, *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, Cambridge (2017)
- [CKBG15] R. Chaves, R. Kueng, J. B. Brask, and D. Gross, *Unifying Framework for Relaxations of the Causal Assumptions in Bell’s Theorem*. Physical Review Letters **114**(14), 140403 (2015). arXiv:1411.4648 [quant-ph]
- [CKMR07] M. Christandl, R. Koenig, G. Mitchison, and R. Renner, *One-and-a-Half Quantum de Finetti Theorems*. Communications in Mathematical Physics **273**(2), 473–498 (2007). arXiv:quant-ph/0602130 [quant-ph]
- [CKW00] V. Coffman, J. Kundu, and W. K. Wootters, *Distributed Entanglement*. Physical Review A **61**(5), 052306 (2000). arXiv:quant-ph/9907047
- [CL13] B. Coecke and R. Lal, *Causal Categories: Relativistically Interacting Processes*. Foundations of Physics **43**(4), 458–501 (2013). arXiv:1107.6019 [quant-ph]
- [CL14] E. G. Cavalcanti and R. Lal, *On Modifications of Reichenbach’s Principle of Common Cause in Light of Bell’s Theorem*. Journal of Physics A: Mathematical and Theoretical **47**(42), 424018 (2014). arXiv:1311.6852 [quant-ph]
- [CLG14] R. Chaves, L. Luft, and D. Gross, *Causal Structures from Entropic Information: Geometry and Novel Scenarios*. New Journal of Physics **16**(4), 043001 (2014). arXiv:1310.0284 [quant-ph]
- [CLM⁺14] R. Chaves, L. Luft, T. O. Maciel, D. Gross, D. Janzing, and B. Schölkopf, *Inferring Latent Structures via Information Inequalities*. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, eds. N. L. Zhang and J. Tian, AUAI Press (2014). arXiv:1407.2256 [quant-ph]

- [CM10] E. G. Cavalcanti and N. C. Menicucci, *Verifiable Nonlinear Quantum Evolution Implies Failure of Density Matrices to Represent Proper Mixtures* (2010). arXiv:1004.1219 [quant-ph]
- [CMG15] R. Chaves, C. Majenz, and D. Gross, *Information-Theoretic Implications of Quantum Causal Structures*. *Nature Communications* **6**, 5766 (2015). arXiv:1407.3800 [quant-ph]
- [CMP12] E. G. Cavalcanti, N. C. Menicucci, and J. L. Pienaar, *The Preparation Problem in Nonlinear Extensions of Quantum Theory* (2012). arXiv:1206.2725 [quant-ph]
- [Coe10] B. Coecke, *Quantum Pictorialism*. *Contemporary Physics* **51**(1), 59–83 (2010). arXiv:0908.1787 [quant-ph]
- [CR11] R. Colbeck and R. Renner, *No Extension of Quantum Theory Can Have Improved Predictive Power*. *Nature Communications* **2**, 411 (2011). arXiv:1005.5173 [quant-ph]
- [CR12] R. Colbeck and R. Renner, *Is a System’s Wave Function in One-to-One Correspondence with Its Elements of Reality?* *Physical Review Letters* **108**(15), 150402 (2012). arXiv:1111.6597 [quant-ph]
- [CR13] R. Colbeck and R. Renner, *A System’s Wave Function is Uniquely Determined by its Underlying Physical State*. *New Journal of Physics* **19**(1), 013016 (2013). arXiv:1312.7353 [quant-ph]
- [CR16] R. Colbeck and R. Renner, *The Completeness of Quantum Theory for Predicting Measurement Outcomes*. In *Quantum Theory: Informational Foundations and Foils*, eds. G. Chiribella and R. W. Spekkens, vol. 181 of *Fundamental Theories of Physics*, pp. 497–528, Springer, Berlin (2016). arXiv:1208.4123 [quant-ph]
- [CS16] F. Costa and S. Shrapnel, *Quantum Causal Modelling*. *New Journal of Physics* **18**(6), 063032 (2016). arXiv:1512.07106 [quant-ph]
- [CT09] M. Christandl and B. Toner, *Finite de Finetti Theorem for Conditional Probability Distributions Describing Physical Theories*. *Journal of Mathematical Physics* **50**, 042104 (2009). arXiv:0712.0916 [quant-ph]

- [DE10] M. Dorato and M. Esfeld, *GRW as an Ontology of Dispositions*. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **41**(1), 41–49 (2010)
- [Deu85] D. Deutsch, *Quantum Theory as a Universal Physical Theory*. *International Journal of Theoretical Physics* **24**(1), 1–41 (1985)
- [Deu91] D. Deutsch, *Quantum Mechanics near Closed Timelike Lines*. *Physical Review D* **44**(10), 3197 (1991)
- [DF80] P. Diaconis and D. Freedman, *Finite Exchangeable Sequences*. *The Annals of Probability* **8**(4), 745–764 (1980)
- [DFI10] R. DeJonghe, K. Frey, and T. Imbo, *Discontinuous Quantum Evolutions in the Presence of Closed Timelike Curves*. *Physical Review D* **81**(8), 087501 (2010). arXiv:0908.2655 [quant-ph]
- [DG73] B. S. DeWitt and R. N. Graham, eds., *The Many-Worlds Interpretation of Quantum Mechanics*. Princeton Series in Physics, Princeton University Press, Princeton (1973)
- [EKT91] F. Echeverria, G. Klinkhammer, and K. S. Thorne, *Billiard Balls in Wormhole Spacetimes with Closed Timelike Curves: Classical Theory*. *Physical Review D* **44**(4), 1077 (1991)
- [ELN14] C. Emary, N. Lambert, and F. Nori, *Leggett-Garg Inequalities*. *Reports on Progress in Physics* **77**(1), 016001 (2014). arXiv:1304.5133 [quant-ph]
- [EPR35] A. Einstein, B. Podolsky, and N. Rosen, *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?* *Physical Review* **47**(10), 777–780 (1935)
- [ESSV13] J. Emerson, D. Serbin, C. Sutherland, and V. Veitch, *The Whole is Greater than the Sum of the Parts: On the Possibility of Purely Statistical Interpretations of Quantum Theory* (2013). arXiv:1312.1345 [quant-ph]
- [ESW02] T. Eggeling, D. Schlingemann, and R. F. Werner, *Semicausal Operations are Semilocalizable*. *Europhysics Letters* **57**(6), 782–788 (2002). arXiv:quant-ph/0104027 [quant-ph]

- [FB16] A. Feix and Č. Brukner, *Quantum Superpositions of “Common-Cause” and “Direct-Cause” Causal Structures* (2016). arXiv:1606.09241 [quant-ph]
- [Fey64] R. Feynman, *Probability and Uncertainty: The Quantum Mechanical View of Nature* (1964). “The Character of Physical Law” Lecture 6, Cornell University. Recorded and broadcast by the BBC. Available at <http://www.cornell.edu/video/richard-feynman-messenger-lecture-6-probability-uncertainty-quantum-mechanical-view-nature> accessed 2017-03-02
- [FG99] C. A. Fuchs and J. van de Graaf, *Cryptographic Distinguishability Measures for Quantum-mechanical States*. Information Theory, IEEE Transactions on **45**(4), 1216–1227 (1999). arXiv:quant-ph/9712042 [quant-ph]
- [Fin75] B. de Finetti, *Theory of Probability: A Critical Introductory Treatment*, vol. 1 and 2. Wiley (1975)
- [Fin93] B. de Finetti, *Probabilità e Induzione—Induction and Probability*. Biblioteca di STATISTICA, CLUEB, Bologna (1993)
- [FMS14] C. A. Fuchs, N. D. Mermin, and R. Schack, *An Introduction to QBism with an Application to the Locality of Quantum Mechanics*. American Journal of Physics **82**(8), 749–754 (2014). arXiv:1311.5253 [quant-ph]
- [FR15] O. Fawzi and R. Renner, *Quantum Conditional Mutual Information and Approximate Markov Chains*. Communications in Mathematical Physics **340**(2), 575–611 (2015). arXiv:1410.0664 [quant-ph]
- [FR16] D. Frauchiger and R. Renner, *Single-World Interpretations of Quantum Theory Cannot be Self-Consistent* (2016). arXiv:1604.07422 [quant-ph]
- [Fri12] T. Fritz, *Beyond Bell’s Theorem: Correlation Scenarios*. New Journal of Physics **14**(10), 103001 (2012). arXiv:1206.5115 [quant-ph]
- [Fri16] T. Fritz, *Beyond Bell’s Theorem II: Scenarios with Arbitrary Causal Structure*. Communications in Mathematical Physics **341**(2), 391–434 (2016). arXiv:1404.4812 [quant-ph]
- [Fuc02] C. A. Fuchs, *Quantum Mechanics as Quantum Information (and only a little more)* (2002). arXiv:quant-ph/0205039 [quant-ph]

- [FW95] C. J. Fewster and C. G. Wells, *Unitarity of Quantum Theory and Closed Timelike Curves*. *Physical Review D* **52**(10), 5773 (1995). arXiv:hep-th/9409156 [hep-th]
- [GCKGM⁺12] S. S. George C. Knee, E. M. Gauger, J. J. L. Morton, H. Riemann, N. V. Abrosimov, P. Becker, H.-J. Pohl, K. M. Itoh, M. L. W. Thewalt, G. A. D. Briggs, and S. C. Benjamin, *Violation of a Leggett-Garg Inequality with Ideal Non-invasive Measurements*. *Nature Communications* **3**, 606 (2012). arXiv:1104.0238 [quant-ph]
- [GHZ89] D. M. Greenberger, M. A. Horne, and A. Zeilinger, *Going Beyond Bell's Theorem*. In *Bell's Theorem, Quantum Theory, and Conceptions of the Universe*, ed. M. Kafatos, *Fundamental Theories of Physics*, Springer (1989). arXiv:0712.0921 [quant-ph]
- [Gib87] P. Gibbins, *Particles and Paradoxes: The Limits of Quantum Logic*. Cambridge University Press, Cambridge (1987)
- [GJ79] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, New York (1979)
- [Göd49] K. Gödel, *An Example of a New Type of Cosmological Solution of Einstein's Field Equations of Gravitation*. *Reviews of Modern Physics* **21**(3), 447 (1949)
- [Gol08] G. A. Goldin, *Nonlinear Quantum Mechanics: Results and Open Questions*. *Physics of Atomic Nuclei* **71**, 884 (2008)
- [Gol16] S. Goldstein, *Bohmian Mechanics*. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, fall 2016 ed. (2016)
- [Got91] J. R. Gott, *Closed Timelike Curves Produced by Pairs of Moving Cosmic Strings: Exact Solutions*. *Physical Review Letters* **66**(9), 1126 (1991)
- [GRW86] G. C. Ghirardi, Rimini, and Weber, *Unified dynamics for Microscopic and Macroscopic Systems*. *Physical Review D* **34**(2), 470 (1986)
- [GVW⁺15] M. Giustina, M. A. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J. Åke Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin, B. Wittmann, , and A. Zeilinger, *Significant-Loophole-Free Test of Bell's Theorem with Entangled*

- Photons*. Physical Review Letters **115**(25), 250401 (2015). arXiv:1511.03190 [quant-ph]
- [Hal11] M. J. W. Hall, *Generalisations of the Recent Pusey-Barrett-Rudolph Theorem for Statistical Models of Quantum Phenomena* (2011). arXiv:1111.6304 [quant-ph]
- [Har94] J. B. Hartle, *Unitarity and Causality in Generalised Quantum Mechanics for Nonchronal Spacetimes*. Physical Review D **49**(12), 6543 (1994). arXiv:gr-qc/9309012 [gr-qc]
- [Har01] L. Hardy, *Quantum Theory From Five Reasonable Axioms* (2001). arXiv:quant-ph/0101012 [quant-ph]
- [Har04] L. Hardy, *Quantum Ontological Excess Baggage*. Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics **35**(2), 267–276 (2004)
- [Har09] L. Hardy, *Quantum Gravity Computers: On the Theory of Computation with Indefinite Causal Structure*. In *Quantum Reality, Relativistic Causality, and Closing the Epistemic Circle*, eds. W. C. Myrvold and J. Christian, pp. 379–401, Springer, Berlin (2009). arXiv:quant-ph/0701019 [quant-ph]
- [Har11] L. Hardy, *Reformulating and Reconstructing Quantum Theory* (2011). arXiv:1104.2066 [quant-ph]
- [Har13] L. Hardy, *Are Quantum States Real?* International Journal of Modern Physics B **27**, 1345012 (2013). arXiv:1205.1439 [quant-ph]
- [Haw92] S. W. Hawking, *Chronology Protection Conjecture*. Physical Review D **46**(2), 603 (1992)
- [Haw95] S. W. Hawking, *Quantum Coherence and Closed Timelike Curves*. Physical Review D **52**(10), 5681 (1995). arXiv:gr-qc/9502017 [gr-qc]
- [HBD⁺15] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Loophole-Free Bell Inequality Violation Using Electron Spins Separated by 1.3 Kilometres*. Nature **526**(7575), 682–686 (2015)

- [Hel13] C. Held, *The Kochen-Specker Theorem*. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, spring 2013 ed. (2013)
- [HHP⁺16] D. Horsman, C. Heunen, M. F. Pusey, J. Barrett, and R. W. Spekkens, *Can a Quantum State Over Time Resemble a Quantum State at a Single Time?* (2016). arXiv:1607.03637 [quant-ph]
- [HJPW04] P. Hayden, R. Jozsa, D. Petz, and A. Winter, *Structure of States Which Satisfy Strong Subadditivity of Quantum Entropy with Equality*. *Communications in Mathematical Physics* **246**(2), 359–374 (2004). arXiv:quant-ph/0304007 [quant-ph]
- [HLP14] J. Henson, R. Lal, and M. F. Pusey, *Theory-Independent Limits on Correlations from Generalized Bayesian Networks*. *New Journal of Physics* **16**(11), 113043 (2014). arXiv:1405.2572 [quant-ph]
- [HM04] G. T. Horowitz and J. Maldacena, *The Black Hole Final State*. *Journal of High Energy Physics* **02**, 008 (2004). arXiv:hep-th/0310281 [hep-th]
- [HM16] E. Huffman and A. Mizel, *Leggett-Garg Test of Superconducting Qubit Addressing the Clumsiness Loophole* (2016). arXiv:1609.05957 [quant-ph]
- [HM17] R. Hermens and O. J. E. Maroney, *Constraints on Macroscopic Realism Without Assuming Non-invasive Measurability* (2017). (Forthcoming)
- [Hol73] A. S. Holevo, *Bounds for the Quantity of Information Transmitted by a Quantum Communication Channel*. *Problemy Peredachi Informatsii* **9**(3), 3 (1973)
- [HR07] N. Harrigan and T. Rudolph, *Ontological Models and the Interpretation of Contextuality* (2007). arXiv:0709.4266 [quant-ph]
- [HS10] N. Harrigan and R. W. Spekkens, *Einstein, Incompleteness, and the Epistemic View of Quantum States*. *Foundations of Physics* **40**(2), 125 (2010). arXiv:0706.2661 [quant-ph]
- [Hur97] A. Hurwitz. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* pp. 71–90 (1897)
- [HWB11] M. J. Hoban, J. J. Wallman, and D. E. Browne, *Generalized Bell-Inequality Experiments and Computation*. *Physical Review A* **84**(6), 062107 (2011). arXiv:1108.4798 [quant-ph]

- [ILW08] B. Ibinson, N. Linden, and A. Winter, *Robustness of Quantum Markov Chains*. *Communications in Mathematical Physics* **277**(2), 289–304 (2008). arXiv:quant-ph/0611057 [quant-ph]
- [Jam72] A. Jamiolkowski, *Linear Transformations which Preserve Trace and Positive Semidefiniteness of Operators*. *Reports on Mathematical Physics* **3**(4), 275–278 (1972)
- [JL03] R. Jozsa and N. Linden, *On the Role of Entanglement in Quantum-Computational Speed-up*. *Proceedings of the Royal Society A* **459**(2036), 2011–2032 (2003). arXiv:quant-ph/0201143 [quant-ph]
- [JL15] D. Jennings and M. S. Leifer, *No Return to Classical Reality*. *Contemporary Physics* **57**(1), 60–82 (2015). arXiv:1501.03202 [quant-ph]
- [Jos98] R. Jozsa, *Entanglement and Quantum Computation*. In *The Geometric Universe: Science, Geometry, and the Work of Roger Penrose*, eds. S. A. Huggett, L. J. Mason, K. P. Tod, S. Tsou, and N. M. J. Woodhouse, pp. 369–379, Oxford University Press, Oxford (1998)
- [Joz04] R. Jozsa, *Illustrating the Concept of Quantum Information*. *IBM Journal of Research and Development* **48**(1), 79–85 (2004). arXiv:quant-ph/0305114 [quant-ph]
- [JRS⁺15] M. Junge, R. Renner, D. Sutter, M. M. Wilde, and A. Winter, *Universal Recovery from a Decrease of Quantum Relative Entropy* (2015). arXiv:1509.07127 [quant-ph]
- [KB13] J. Kofler and Āaslav Brukner, *Condition for Macroscopic Realism Beyond the Leggett-Garg Inequalities*. *Physical Review A* **87**(5), 052115 (2013). arXiv:1207.3666 [quant-ph]
- [Ken99] A. Kent, *Non-Contextual Hidden Variables and Physical Measurements*. *Physical Review Letters* **83**(19), 3755–3757 (1999). arXiv:quant-ph/9906006 [quant-ph]
- [Ken05] A. Kent, *Nonlinearity Without Superluminality*. *Physical Review A* **72**(1), 012108 (2005). arXiv:quant-ph/0204106 [quant-ph]

- [KKY⁺16] G. C. Knee, K. Kakuyanagi, M.-C. Yeh, Y. Matsuzaki, H. Toida, H. Yamaguchi, S. Saito, A. J. Leggett, and W. J. Munro, *A Strict Experimental Test of Macroscopic Realism in a Superconducting Flux Qubit*. *Nature Communications* **7**, 13253 (2016). arXiv:1601.03728 [quant-ph]
- [Kne16] G. C. Knee, *Towards Optimal Experimental Tests on the Reality of the Quantum State* (2016). arXiv:1609.01558 [quant-ph]
- [KS67] S. Kochen and E. P. Specker, *The Problem of Hidden Variables in Quantum Mechanics*. *Journal of Mathematics and Mechanics* **17**, 59–87 (1967)
- [LB15] C. M. Lee and J. Barrett, *Computation in Generalised Probabilistic Theories*. *New Journal of Physics* **17**, 083001 (2015). arXiv:1412.8671 [quant-ph]
- [Leg88] A. J. Leggett, *Experimental Approaches to the Quantum Measurement Paradox*. *Foundations of Physics* **18**(9), 939–952 (1988)
- [Leg02a] A. J. Leggett, *Probing Quantum Mechanics Towards the Everyday World: Where do we Stand?* *Physica Scripta* **2002**(T102), 69 (2002)
- [Leg02b] A. J. Leggett, *Testing the Limits of Quantum Mechanics: Motivation, State of Play, Prospects*. *Journal of Physics: Condensed Matter* **14**(15), R415 (2002)
- [Lei06] M. S. Leifer, *Quantum Dynamics as an Analog of Conditional Probability*. *Physical Review A* **74**(4), 042310 (2006). arXiv:quant-ph/0606022 [quant-ph]
- [Lei11] M. S. Leifer, *The Choi-Jamiołkowski Isomorphism: You’re Doing it Wrong!* (2011). <http://mattleifer.info/2011/08/01/the-choi-jamiołkowski-isomorphism-youre-doing-it-wrong/> [Online; accessed 2017-01-05]
- [Lei14a] M. S. Leifer, *Is the Quantum State Real? An Extended Review of ψ -ontology Theorems*. *Quanta* **3**(1), 67–155 (2014). arXiv:1409.1570 [quant-ph]
- [Lei14b] M. S. Leifer, *ψ -Epistemic Models are Exponentially Bad at Explaining the Distinguishability of Quantum States*. *Physical Review Letters* **112**(16), 160404 (2014). arXiv:1401.7996 [quant-ph]
- [LG85] A. J. Leggett and A. Garg, *Quantum Mechanics versus Macroscopic Realism: Is the Flux There when Nobody Looks?* *Physical Review Letters* **54**(9), 857–860 (1985)

- [LG87] A. J. Leggett and A. Garg, *Comment on “Realism and Quantum Flux Tunneling”*. Physical Review Letters **59**(14), 1621 (1987)
- [LJBR12] P. G. Lewis, D. Jennings, J. Barrett, and T. Rudolph, *Distinct Quantum States Can Be Compatible with a Single State of Reality*. Physical Review Letters **109**(15), 150404 (2012). arXiv:1201.6554 [quant-ph]
- [LM13] M. S. Leifer and O. J. E. Maroney, *Maximally Epistemic Interpretations of the Quantum State and Contextuality*. Physical Review Letters **110**(12), 120401 (2013). arXiv:1208.5132 [quant-ph]
- [LMGP⁺11a] S. Lloyd, L. Maccone, R. Garcia-Patron, V. Giovannetti, and Y. Shikano, *Quantum Mechanics of Time Travel through Post-selected Teleportation*. Physical Review D **84**(2), 025007 (2011). arXiv:1007.2615 [quant-ph]
- [LMGP⁺11b] S. Lloyd, L. Maccone, R. Garcia-Patron, V. Giovannetti, Y. Shikano, S. Pirandola, L. A. Rozema, A. Darabi, Y. Soudagar, L. K. Shalm, and A. M. Steinberg, *Closed Timelike Curves via Postselection: Theory and Experimental Test of Consistency*. Physical Review Letters **106**(4), 040403 (2011). arXiv:1005.2219 [quant-ph]
- [LMGP⁺11c] S. Lloyd, L. Maccone, R. Garcia-Patron, V. Giovannetti, Y. Shikano, S. Pirandola, L. A. Rozema, A. Darabi, Y. Soudagar, L. K. Shalm, and A. M. Steinberg, *A Reply to “Problems with Modelling Closed Timelike Curves as Post-Selected Teleportation”* (2011). arXiv:1108.0153 [quant-ph]
- [LP08] M. S. Leifer and D. Poulin, *Quantum Graphical Models and Belief Propagation*. Annals of Physics **323**(8), 1899–1946 (2008). arXiv:0708.1337 [quant-ph]
- [LPZ⁺16] Z.-W. Liu, C. Perry, Y. Zhu, D. E. Koh, and S. Aaronson, *Doubly Infinite Separation of Quantum Information and Communication*. Physical Review A **93**(1), 012347 (2016). arXiv:1507.03546 [quant-ph]
- [LS13] M. S. Leifer and R. W. Spekkens, *Towards a Formulation of Quantum Theory as a Causally Neutral Theory of Bayesian Inference*. Physical Review A **88**(5), 052130 (2013). arXiv:1107.5849 [quant-ph]
- [LS15] C. M. Lee and R. W. Spekkens, *Causal Inference via Algebraic Geometry: Necessary and Sufficient Conditions for the Feasibility of Discrete Causal Models* (2015). arXiv:1506.03880 [stat.ML]

- [LW14] K. Li and A. Winter, *Squashed Entanglement, k -extendibility, Quantum Markov Chains, and Recovery Maps* (2014). arXiv:1410.4184 [quant-ph]
- [Man16] S. Mansfield, *Reality of the Quantum State: Towards A Stronger ψ -ontology Theorem*. Physical Review A **94**(4), 042124 (2016). arXiv:1412.0669 [quant-ph]
- [Mar12] O. J. E. Maroney, *How Statistical are Quantum States?* (2012). arXiv:1207.6906 [quant-ph]
- [Mey99] D. A. Meyer, *Finite Precision Measurement Nullifies the Kochen-Specker theorem*. Physical Review Letters **83**(19), 3751–3754 (1999). arXiv:quant-ph/9905080 [quant-ph]
- [Mon08] A. Montina, *Exponential Complexity and Ontological Theories of Quantum Mechanics*. Physical Review A **77**(2), 022104 (2008). arXiv:0711.4770 [quant-ph]
- [Mon11] A. Montina, *Communication Cost of Classically Simulating a Quantum Channel with Subsequent Rank-1 Projective Measurement*. Physical Review A **84**(6), 060303 (2011). arXiv:1110.5944 [quant-ph]
- [Mon12] A. Montina, *Epistemic View of Quantum States and Communication Complexity of Quantum Channels*. Physical Review Letters **109**(11), 110501 (2012). arXiv:1206.2961 [quant-ph]
- [Mon13] A. Montina, *Exponential Communication Gap Between Weak and Strong Classical Simulations of Quantum Communication*. Physical Review A **87**(4), 042331 (2013). arXiv:1301.3452 [quant-ph]
- [Mon15] A. Montina, *Communication Complexity and the Reality of the Wave Function*. Modern Physics Letters A **30**(1), 1530001 (2015). arXiv:1412.1723 [quant-ph]
- [Mon16a] A. Montanaro, *Quantum States Cannot be Transmitted Efficiently Classically* (2016). arXiv:1612.06546 [quant-ph]
- [Mon16b] A. Montina, *Information-Based Measure of Nonlocality*. New Journal of Physics **18**, 013035 (2016). arXiv:1312.6290 [quant-ph]
- [MRSR16] J.-P. W. MacLean, K. Ried, R. W. Spekkens, and K. J. Resch, *Quantum-Coherent Mixtures of Causal Relations* (2016). arXiv:1606.04523 [quant-ph]

- [MT17] O. J. E. Maroney and C. G. Timpson, *Quantum- vs. Macro- Realism: What does the Leggett-Garg Inequality Actually Test?* *British Journal for the Philosophy of Science* **In Production** (2017). arXiv:1412.6139 [quant-ph]
- [MTY88] M. S. Morris, K. S. Thorne, and U. Yurtsever, *Wormholes, Time Machines, and the Weak Energy Condition*. *Physical Review Letters* **61**(13), 1446 (1988)
- [NC00] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge (2000)
- [NMS⁺15] D. Nigg, T. Monz, P. Schindler, E. A. Martinez, M. Chwalla, M. Hennrich, R. Blatt, M. F. Pusey, T. Rudolph, and J. Barrett, *Can Different Quantum State Vectors Correspond to the Same Physical State? An Experimental Test*. *New Journal of Physics* **18**, 013007 (2015). arXiv:1211.0942 [quant-ph]
- [OCB12] O. Oreshkov, F. Costa, and Časlav Brukner, *Quantum Correlations with no Causal Order*. *Nature Communications* **3**, 1092 (2012). arXiv:1105.4464 [quant-ph]
- [PB00] A. K. Pati and S. L. Braunstein, *Impossibility of Deleting an Unknown Quantum State*. *Nature* **404**(6774), 164–165 (2000). arXiv:quant-ph/9911090 [quant-ph]
- [PB15] J. Pienaar and Časlav Brukner, *A Graph-Separation Theorem for Quantum Causal Models*. *New Journal of Physics* **17**(7), 073020 (2015). arXiv:1406.0430 [quant-ph]
- [PBR12] M. F. Pusey, J. Barrett, and T. Rudolph, *On the Reality of the Quantum State*. *Nature Physics* **8**(6), 475–478 (2012). arXiv:1111.3328 [quant-ph]
- [PCA11] A. K. Pati, I. Chakrabarty, and P. Agrawal, *Purification of Mixed States with Closed Timelike Curve is not Possible*. *Physical Review A* **84**(6), 062325 (2011). arXiv:1003.4221 [quant-ph]
- [Pea09] J. Pearl, *Causality: Models, Reasoning, and Inference*. Second ed., Cambridge University Press, Cambridge (2009)
- [Pen98] R. Penrose, *Quantum Computation, Entanglement and State Reduction*. *Philosophical Transactions of the Royal Society of London A* **356**, 1927–1939 (1998)

- [PHHH06] M. Piani, M. Horodecki, P. Horodecki, and R. Horodecki, *Properties of Quantum Nonsignaling Boxes*. Physical Review A **74**(1), 012305 (2006). arXiv:quant-ph/0505110 [quant-ph]
- [Pie13] J. L. Pienaar, *Causality Violation and Nonlinear Quantum Mechanics*. PhD thesis, School of Mathematics and Physics, University of Queensland (2013). arXiv:1401.0167 [quant-ph]
- [PJO15] C. Perry, R. Jain, and J. Oppenheim, *Communication Tasks with Infinite Quantum-Classical Separation*. Physical Review Letters **115**(3), 030504 (2015). arXiv:1407.8217 [quant-ph]
- [Pol94] H. D. Politzer, *Path Integrals, Density Matrices, and Information Flow with Closed Timelike Curves*. Physical Review D **49**(8), 3981 (1994). arXiv:gr-qc/9310027 [gr-qc]
- [PPM13] M. K. Patra, S. Pironio, and S. Massar, *No-Go Theorems for ψ -Epistemic Models Based on a Continuity Assumption*. Physical Review Letters **111**(9), 090402 (2013). arXiv:1211.1179 [quant-ph]
- [Rag02] M. Raginsky, *Strictly Contractive Quantum Channels and Physically Realizable Quantum Computers*. Physical Review A **65**(3), 032306 (2002). arXiv:quant-ph/0105141 [quant-ph]
- [RBB⁺16] D. Rosset, C. Branciard, T. J. Barnea, G. Pütz, N. Brunner, and N. Gisin, *Nonlinear Bell Inequalities Tailored for Quantum Networks*. Physical Review Letters **116**(1), 010403 (2016). arXiv:1506.07380 [quant-ph]
- [RDB⁺15] M. Ringbauer, B. Duffus, C. Branciard, E. G. Cavalcanti, A. G. White, and A. Fedrizzi, *Measurements on the Reality of the Wavefunction*. Nature Physics **11**, 249–254 (2015). arXiv:1412.6213 [quant-ph]
- [Rei56] H. Reichenbach, *The Direction of Time*. University of California Press, Berkeley (1956)
- [RM10] T. C. Ralph and C. R. Myers, *Information Flow of Quantum States Interacting with Closed Timelike Curves*. Physical Review A **82**(6), 062330 (2010). arXiv:1003.1987 [quant-ph]

- [RR13] T. S. Richardson and J. M. Robins, *Single World Intervention Graphs (SWIGs): A Unification of the Counterfactual and Graphical Approaches to Causality*. Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper **128** (2013)
- [RS11] V. K. Rohatgi and A. K. M. E. Saleh, *An Introduction to Probability and Statistics*, vol. 910 of *Wiley Series in Probability and Statistics*. Second ed., John Wiley & Sons (2011)
- [Rus02] M. B. Ruskai, *Inequalities for Quantum Entropy: A Review with Conditions for Equality*. *Journal of Mathematical Physics* **43**(9), 4358–4375 (2002). arXiv:quant-ph/0205064 [quant-ph]
- [SBKW10] S. Saunders, J. Barrett, A. Kent, and D. Wallace, eds., *Many Worlds?: Everett, Quantum Theory, & Reality*. Oxford University Press, Oxford (2010)
- [Sch30] J. Schauder, *Der Fixpunktsatz in Funktionalräumen*. *Sudia Mathematica* **2**, 171 (1930)
- [SF14] M. Schlosshauer and A. Fine, *No-Go Theorem for the Composition of Quantum Systems*. *Physical Review Letters* **112**(7), 070407 (2014). arXiv:1306.5805 [quant-ph]
- [SFR16] D. Sutter, O. Fawzi, and R. Renner, *Universal Recovery Map for Approximate Markov Chains*. *Proceedings of the Royal Society A* **472**(2186), 20150623 (2016). arXiv:1504.07251 [quant-ph]
- [SGB⁺14] R. Silva, Y. Guryanova, N. Brunner, N. Linden, A. J. Short, and S. Popescu, *Pre- and Post-Selected Quantum States: Density Matrices, Tomography, and Kraus Operators*. *Physical Review A* **89**(1), 012121 (2014). arXiv:1308.2089 [quant-ph]
- [SGS01] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*. Second ed., The MIT Press, Cambridge, MA (2001)
- [Shi13] A. Shimony, *Bell's Theorem*. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, winter 2013 ed. (2013)
- [SMSC⁺15] L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, K. J. Coakley, S. D. Dyer, C. Hodge, A. E. Lita, V. B. Verma, C. Lambrocco, E. Tortorici,

- A. L. Migdall, Y. Zhang, D. R. Kumor, W. H. Farr, F. Marsili, M. D. Shaw, J. A. Stern, C. Abellán, W. Amaya, V. Pruneri, T. Jennewein, M. W. Mitchell, P. G. Kwiat, J. C. Bienfang, R. P. Mirin, E. Knill, and S. W. Nam, *A Strong Loophole-Free Test of Local Realism*. *Physical Review Letters* **115**(25), 250402 (2015). arXiv:1511.03189 [quant-ph]
- [Spe05] R. W. Spekkens, *Contextuality for Preparations, Transformations, and Unsharp Measurements*. *Physical Review A* **71**(5), 052108 (2005). arXiv:quant-ph/0406166 [quant-ph]
- [Spe07] R. W. Spekkens, *Evidence for the Epistemic View of Quantum States: A Toy Theory*. *Physical Review A* **75**(3), 032110 (2007). arXiv:quant-ph/0401052 [quant-ph]
- [Spe14] R. W. Spekkens, *Quasi-quantization: Classical Statistical Theories with an Epistemic Restriction* (2014). arXiv:1409.5041 [quant-ph]
- [Spe16] R. W. Spekkens, *Quasi-quantization: Classical Statistical Theories with an Epistemic Restriction*. In *Quantum Theory: Informational Foundations and Foils*, eds. G. Chiribella and R. W. Spekkens, vol. 181 of *Fundamental Theories of Physics*, pp. 83–135, Springer, Berlin (2016). arXiv:1409.5041 [quant-ph]
- [STH16] D. Sutter, M. Tomamichel, and A. W. Harrow, *Strengthened Monotonicity of Relative Entropy via Pinched Petz Recovery Map*. *IEEE Transactions on Information Theory* **62**(5), 2907–2913 (2016). arXiv:1507.00303 [quant-ph]
- [Sve09] G. Svetlichny, *Effective Quantum Time Travel* (2009). arXiv:0902.4898 [quant-ph]
- [SW05] B. Schumacher and M. D. Westmoreland, *Locality and Information Transfer in Quantum Operations*. *Quantum Information Processing* **4**(1), 13–34 (2005). arXiv:quant-ph/0406223 [quant-ph]
- [SW12] B. Schumacher and M. D. Westmoreland, *Isolation and Information Flow in Quantum Dynamics*. *Foundations of Physics* **42**(7), 926–931 (2012)
- [TB03] B. F. Toner and D. Bacon, *Communication Cost of Simulating Bell Correlations*. *Physical Review Letters* **91**(18), 187904 (2003). arXiv:quant-ph/0304076 [quant-ph]

- [Tim06] C. G. Timpson, *The Grammar of Teleportation*. The British Journal for the Philosophy of Science **57**(3), 587–621 (2006)
- [TSCA14] A. Tavakoli, P. Skrzypczyk, D. Cavalcanti, and A. Acín, *Nonlocal Correlations in the Star-Network Configuration*. Physical Review A **90**(6), 062109 (2014). arXiv:1409.5702 [quant-ph]
- [Tuc95] R. R. Tucci, *Quantum Bayesian Nets*. International Journal of Modern Physics B **9**(3), 295–337 (1995). arXiv:quant-ph/9706039 [quant-ph]
- [Tuc12] R. R. Tucci, *An Introduction to Quantum Bayesian Networks for Mixed States* (2012). arXiv:1204.1550 [quant-ph]
- [Tyc35] A. Tychonoff, *Ein Fixpunktsatz*. Mathematische Annalen **111**, 767 (1935)
- [Uhl11] A. Uhlmann, *Transition Probability (Fidelity) and Its Relatives*. Foundations of Physics **41**(3), 288–298 (2011). arXiv:1106.0979 [quant-ph]
- [Wal10] R. M. Wald, *General Relativity*. University of Chicago Press, Chicago (2010)
- [Wal13] P. Wallden, *Distinguishing Initial State-Vectors from Each Other in Histories Formulations and the PBR Argument*. Foundations of Physics **43**(12), 1502–1525 (2013). arXiv:1211.2084 [quant-ph]
- [WB12] J. J. Wallman and S. D. Bartlett, *Revisiting Consistency Conditions for Quantum States of Systems on Closed Timelike Curves*. Foundations of Physics **42**(5), 656 (2012). arXiv:1005.2438 [quant-ph]
- [WDN⁺12] G. Waldherr, A. C. Dada, P. Neumann, F. Jelezko, E. Andersson, and J. Wrachtrup, *Distinguishing between Nonorthogonal Quantum States of a Single Nuclear Spin*. Physical Review Letters **109**(18), 180501 (2012). arXiv:1206.0453 [quant-ph]
- [Wei89] S. Weinberg, *Testing Quantum Mechanics*. Annals of Physics **194**(2), 336–386 (1989)
- [Wig67] E. P. Wigner, *Remarks on the Mind-Body Question*. In *Symmetries and Reflections*, pp. 171–184, Indiana University Press (1967)
- [Wil15] M. M. Wilde, *Recoverability in Quantum Information Theory*. Proceedings of the Royal Society A **471**(2182), 20150338 (2015). arXiv:1505.04661 [quant-ph]

- [WKZ⁺17] K. Wang, G. C. Knee, X. Zhan, Z. Bian, J. Li, and P. Xue, *Optimal Experimental Demonstration of Error-Tolerant Quantum Witnesses*. *Physical Review A* **95**(3), 032122 (2017). arXiv:1701.04630 [quant-ph]
- [WM12] M. M. Wilde and A. Mizel, *Addressing the Clumsiness Loophole in a Leggett-Garg Test of Macrorealism*. *Foundations of Physics* **42**(2), 256–265 (2012). arXiv:1001.1777 [quant-ph]
- [WS15] C. J. Wood and R. W. Spekkens, *The Lesson of Causal Discovery Algorithms for Quantum Correlations: Causal Explanations of Bell-Inequality Violations Require Fine-Tuning*. *New Journal of Physics* **17**(3), 033002 (2015). arXiv:1208.4119 [quant-ph]
- [WSF16] E. Wolfe, R. W. Spekkens, and T. Fritz, *The Inflation Technique for Causal Inference with Latent Variables* (2016). arXiv:1609.00672 [quant-ph]
- [Zei85] E. Zeidler, *Nonlinear Functional Analysis and its Applications Volume 1: Fixed Point Theorems*. Springer-Verlag, Berlin (1985)
- [ZHLG15] Z.-Q. Zhou, S. F. Huelga, C.-F. Li, and G.-C. Guo, *Experimental Detection of Quantum Coherent Evolution through the Violation of Leggett-Garg-Type Inequalities*. *Physical Review Letters* **115**(11), 113002 (2015). arXiv:1209.2176 [quant-ph]
- [ZS01] K. Życzkowski and H.-J. Sommers, *Induced Measures in the Space of Mixed Quantum States*. *Journal of Physics A* **34**(35), 7111 (2001). arXiv:quant-ph/0012101 [quant-ph]