

Nicholas Gauguin Houghton-Larsen

A Mathematical Framework for
Causally Structured Dilations and its
Relation to Quantum Self-Testing

This PhD thesis has been submitted for assessment to the PhD School of the Faculty of Science, University of Copenhagen, on the 1st of December 2020.

PHD THESIS BY:

Nicholas Gauguin Houghton-Larsen
Department of Mathematical Sciences, University of Copenhagen
Universitetsparken 5, 2100 København Ø, Denmark
nicholas.gauguin@gmail.com

DATE OF SUBMISSION:

1st of December 2020

SUPERVISOR:

Matthias Christandl (professor), University of Copenhagen, Denmark

ASSESSMENT COMMITTEE:

Roger Colbeck (professor), University of York, UK
Tobias Fritz (assistant professor), University of Innsbruck, Austria
Nathalie Wahl (professor), University of Copenhagen, Denmark

ISBN: 978-87-7125-039-8

This final version of my thesis has been corrected in agreement with the official guidelines prior to printing. Specifically, minor errors and typographical mistakes have been rectified, and certain lay-out features have been altered.

A version of the thesis will be maintained at www.arxiv.org .

*Nicholas Gauguin Houghton-Larsen
Copenhagen, January 2021*

Abstract

This is a PhD thesis within the sub-field of mathematical physics that pertains to *quantum information theory*. Most of its results can be interpreted in the mathematical language of category theory, and may as such be of interest also outside of quantum information theory.

In high-level terms, I present a framework in which one can argue mathematically about aspects of the following fundamental question: *How do two given implementations of the same physical process compare to each other?* Though of independent interest, the main motivation for this question comes from the area of *quantum self-testing* ([MY98, MY04]), where one desires to understand all the different ways in which a given set of measurement statistics can be produced by an implementation of local measurements on a multipartite quantum state. The problem which motivated the thesis is that although the traditional envision of quantum self-testing is mathematically precise, the language in which it is cast has no clear operational interpretation.

According to the framework proposed in the thesis, a collection of measurement statistics is regarded as the input-output behaviour of an information channel, and the various implementations of this channel correspond to causally structured computations which may be secretly executed in the environment of the channel during our interaction with it. The main contribution of the thesis is to introduce a formalism which makes the previous sentence precise, and to provide its relation to the usual definition of quantum self-testing. The relation is essentially that quantum self-testing corresponds to the existence of an implementation from which all others can be derived, and which moreover holds no pre-existing information about the outputs of the channel. This constitutes a first step towards recasting quantum self-testing in purely operational (theory-independent) terms.

Chapter 1 reviews a variation on a category-theoretic model for physical theories. This model includes quantum information theory and classical information theory, but also more mathematical examples such as any category with finite products (e.g. the categories of sets or groups), and any partially ordered commutative monoid, when suitably interpreted. The key feature of the model is that it facilitates the notion of *marginals* (as known from e.g. classical probability theory), and the dual notion of *dilations*.

Dilations are the topic of Chapter 2. The results presented there are conceptually independent of quantum self-testing, but rather initiate a systematic study of dilations and constitute an original proof of concept, by demonstrating that several features of information theories can be derived from a handful of principles which reference only the structure of dilations.

Chapter 3 contains some initial thoughts as to how to make an approximate (metric) version of the theory of dilations, and a new metric for quantum channels, the *purified diamond-distance* is introduced. It generalises the purified distance of Refs. [TCR10, Tom12].

Chapter 4 lays out a formalism for arguing about information channels whose outputs are causally contingent on their inputs. This can be seen as a generalised alternative to the framework of quantum combs ([CDP09]), but can also be viewed as generalising the abstract notion of traces in symmetric monoidal categories ([JSV96]). The formalism allows us to make precise the notion of a *causal dilation*, which captures the above-mentioned causally structured side-computations.

Finally, in Chapter 5, the connection to quantum self-testing is established. This chapter also contains simple proofs of a few general results about self-testing, and a novel recharacterisation of the set of quantum behaviours in terms of non-signalling properties of their Stinespring dilations.

Resumé

Dette er en ph.d.-afhandling inden for den gren af matematisk fysik der vedrører *kvanteinformationsteori*. De fleste af dens resultater kan fortolkes i et matematisk kategori-teoretisk sprog og kan som sådan være af interesse også uden for kvanteinformationsteorien.

I overordnede træk præsenteres en teoretisk ramme, i hvilken man matematisk kan tale om aspekter ved følgende grundlæggende spørgsmål: *Hvad er forholdet mellem to givne implementeringer af den samme fysiske proces?* Spørgsmålet er af uafhængig interesse, men dets vigtigste motivation kommer fra feltet ‘*quantum self-testing*’ ([MY98, MY04]), hvor man ønsker at forstå alle de forskellige måder, hvorpå et givent sæt af fordelinger for måleresultater kan fremkomme ved lokale målinger på en kvantetilstand delt mellem flere parter. Det problem der motiverede afhandlingen er, at omend den traditionelle opfattelse af ‘*quantum self-testing*’ er matematisk præcis, så har det sprog i hvilket fænomenet er defineret ikke nogen klar operational fortolkning.

Ifølge den teoretiske ramme der udlægges i afhandlingen betragtes et sæt af fordelinger for måleudfald som input-output-opførslen for en informationskanal, og de mulige implementeringer af denne kanal svarer til kausalt strukturerede processer som hemmeligt udføres i kanalens omgivelser i løbet af vores interaktion med den. Afhandlingens hovedbidrag er at indføre en formalisme der gør forudgående sætning præcis, samt at bestemme formalismens relation til den sædvanlige definition af ‘*quantum self-testing*’. Relationen er essentielt set, at ‘*quantum self-testing*’ svarer til eksistensen af en implementering, hvorfra alle andre kan udledes, og som desuden ikke indeholder forhånds-eksisterende information om kanalens outputs. Dette udgør et første skridt i retning af en omarbejdning af ‘*quantum self-testing*’ til rent operationelle (teori-uafhængige) termer.

Kapitel 1 gennemgår en variation af en kategori-teoretisk model for fysiske teorier. Modellen inkluderer kvanteinformationsteori og klassisk informationsteori, men også mere matematiske eksempler, såsom enhver kategori med endelige produkter (f.eks. kategorierne bestående af mængder eller grupper), og ethvert partielt ordnet kommutativt monoid, passende fortolket. Nøgleegenskaben ved modellen er, at den tillader begrebet *marginalisering* (som det kendes eksempelvis fra sandsynlighedsteorien) og det duale begreb *udvidelse* (eng. ‘*dilations*’).

Udvidelser er emnet for Kapitel 2. Resultaterne, der præsenteres dér, er konceptuelt uafhængige af ‘*quantum self-testing*’, men indleder snarere en systematisk undersøgelse af udvidelser og udgør et originalt ‘*proof of concept*’ ved at demonstrere, at flere informationsteoretiske egenskaber kan udledes fra kun en håndfuld af principper, der alene refererer til strukturen af udvidelser.

Kapitel 3 indeholder nogle indlende tanker om, hvordan man kan lave en approksimativ (metrisk) teori for udvidelser, og en ny metrik for kvantekanaler, ‘*purified diamond-distance*’, introduceres. Denne generaliserer ‘*purified distance*’ fra Ref. [TCR10, Tom12].

Kapitel 4 udlægger en formalisme, hvori man kan tale om informationskanaler hvis outputs er kausalt betingede af deres inputs. Denne kan ses som et generaliseret alternativ til ‘*quantum combs*’ ([CDP09]), men kan også anskues som generalisering af abstrakte spor (eng. ‘*traces*’) i symmetriske monoidiale kategorier ([JSV96]). Formalismen giver os mulighed for at præcisere forestillingen om en *kausal udvidelse* (eng. ‘*causal dilation*’), der netop indfanger de ovennævnte kausalt strukturerede sideprocesser.

Afslutningsvist etableres forbindelsen til ‘*quantum self-testing*’ i Kapitel 5. Dette kapitel indeholder også simple beviser for et par generelle resultater om ‘*quantum self-testing*’, samt en ny karakterisering af mængden af ‘*quantum behaviours*’ i termer af ‘*non-signalling*’-egenskaber ved deres Stinespring-udvidelser.

Acknowledgements

Though a PhD study appears most of the time to be a very lonely endeavour, it goes without saying that it rests upon the support and engagement of many actors. I owe my gratitude to several people and institutions, each of whom and which have played an instrumental role.

First and foremost, I would like to thank my supervisor, professor Matthias Christandl. He dreams big, is profoundly adaptive to new ideas, and firmly believes in the resolution of problems. When he suggested to me to investigate the meaning of quantum self-testing, he likely did not expect the present dissertation as outcome. I am deeply grateful towards him for always keeping open boundaries and giving me space to pursue ways and methods I believed to be interesting and illuminating. I hope that he finds my final results and presentation to duly honour his confidence in me. When I think back at the time we spent together in addition to discussing the project, I am reminded in particular of many entertaining lunch discussions we had when I visited him at MIT in August 2018, and of our exhaustive but exciting period of planning and executing the first year mathematics course *Analyse 1* in April – June 2019. With Matthias, I learned a fundamentally different view on science and mathematics, and his vivid personality has made a lasting impression on me.

Secondly, I am highly indebted to associate professor Laura Mančinska. Despite of having no contractual obligations to my studies, she has effectively co-supervised this thesis project, taken part in the vast majority of supervision meetings with Matthias, and even filled his role enthusiastically in the 6 months of his absence staying at MIT. I am thankful for her commitment and for sharing with me her expert knowledge of the field of quantum self-testing, contributing with many suggestions and points that Matthias and I would have probably overlooked on our own.

As a PhD student, I have been employed at the QMATH Centre of Excellence at the Department of Mathematical Sciences at UCPH (the University of Copenhagen). I would like in this regard to acknowledge the funding by Sapere Aude, the European Research Council (ERC Grant Agreement no. 337603) and VILLUM FONDEN via the QMATH Centre (Grant no. 10059).

At QMATH, I have enjoyed the help and company of many colleagues, and have been astounded by the engagement and ambition with which the centre is run by its scientific founders, professors Jan Philip Solovej, Matthias Christandl and Bergfinnur Durhuus, and by centre administrator Suzanne Andersen. I wish their enterprise the best of luck in the future.

The Department of Mathematical Sciences is a unique work place, and as I was also a bachelor and master student there, it is hard not to think of it with nostalgic warmth. Many of its employees have influenced and inspired me through time, and I would like to express my gratitude in particular towards Niels Richard Hansen, Ernst Hansen, Jan Philip Solovej, Magdalena Musat and Niels Grønbaek. I would also like to acknowledge the support from people whom I have engaged with in their administrative roles, in particular PhD secretary Nina Weisse, administrative officer Mette Fulling, PhD coordinator Morten S. Risager, and head of section Henrik Laurberg Pedersen.

During my PhD studies, I have had the pleasure of spending three weeks at MIT in Boston


in August 2018 and two months the ICMAT and Universidad Complutense in Madrid in October–September 2019. Both of these experiences were very enjoyable, and I would like to thank prof. Aram Harrow at MIT and prof. David Pérez-García at Complutense for their hospitality.

The PhD committee comprises professor Roger Colbeck from the University of York, assistant professor Tobias Fritz from the University of Innsbruck, and (chair of committee) professor Nathalie Wahl from the University of Copenhagen. I am both honoured and happy about the engagement of these three, and would like to thank them for their time. I hope they will all find sentences of value within the thesis.

Last but not least, I am infinitely thankful towards my family and my friends for always supporting and encouraging me in my endeavours. They are too many to list, and I would not risk leaving any of them out.

In the past, I often smiled at dedications of highly technical academic works to people without prerequisites for understanding their content. Now, I see that such dedications serve to recognise that those people were indispensable in shaping and sustaining the person who ultimately grew capable of materialising a product of such ridiculously demanding scope.

With that in mind, I dedicate this work to my family. To my mother, who was one of the most extraordinary, ambitious and giving persons I have ever known; to my father, who ignited my interest in science and whose insights and advice continue to guide me; and to my brother, who always knows how to challenge me and whom I admire for his kindness and intellect more than he could possibly imagine.



Composition A by Piet Mondrian (1923)

Galleria Nazionale d'Arte Moderna e Contemporanea

A Note to the Reader

Everyone who has written down something for anyone to read is familiar with the trivial but crucial condition that statements must be structured *sequentially*, one sentence following the other, paragraph by paragraph, chapter after chapter.

A novelist can use this to advantage, by introducing characters and revealing plot twists according to a carefully crafted schedule. The author of an academic dissertation essentially has to do the same, but as a general rule this circumstance is hindering rather than advantageous. The reason is, of course, that abstract ideas are not connected in a linearly ordered fashion.

An additional dare is posed for academic writing because a reader cannot be counted upon to read every single sentence from the beginning to the end. Few (if any) readers of a fictional novel start by reading the first pages, then read the last, and then sporadically glance through the chapters – in contrast, the order of things which a PhD student envisions for a thesis might in the end not be the one most suitable to any given reader.

In writing this document, I have strived for the storyline to emerge clearly from the general introduction and the individual introductions to the five chapters, so as to guide you as much as possible. This is, however, my first PhD thesis, and so I hope for forgiveness in cases where I have not succeeded in coping sublimely with structural challenges.

Contents

Introduction	i
For Everyone	i
For Someone	iv
For Anyone	xi
Structure of the Thesis	xv
Preliminaries	xviii
1 Theories	1
1.1 A Mathematical Model for Physical Theories	4
1.1.A Definition of Theories – CIT and QIT	7
1.1.B States, Isomorphisms and Reversibles	11
1.2 A Reservoir of Examples	15
1.2.A Variations of CIT and QIT	16
1.2.B Cartesian Theories	18
1.2.C Thin Theories	20
1.2.D Sub-Theories	23
1.3 Pictorial Syntax	24
1.3.A Pictures for Algebra	24
1.3.B Interfaces and Channels	27
1.3.C Normal Theories	31
1.4 Summary and Outlook	32
2 Dilations	33
2.1 A Terminal Trinity: Marginals, Dilations and Non-Signalling	40
2.2 The Dilational Ordering	46
2.2.A Dilational Purity	48
2.3 Completeness and Localisability	49
2.3.A Complete Dilations	50
2.3.B Spatial and Temporal Localisability	52
2.4 Universal Dilations	56
2.5 Purification	64
2.5.A Isomorphisms and Purifiability	67
2.5.B Reversibles and Purifiability	69
2.5.C Complementarity	71
2.6 Summary and Outlook	74

3	Metric Theories	75
3.1	Topological Structure	79
3.2	Metric Structure	80
3.3	Dilationality and the Generalised Uhlmann Property	85
3.4	Purified Dilationality – the Uhlmann Property	89
3.5	The Purified Diamond-Distance P_\diamond	94
3.5.A	P_\diamond versus d_\diamond	96
3.5.B	P_\diamond versus β	97
3.6	Summary and Outlook	103
4	Contractible Theories and Causal Dilations	105
4.1	Causal Channels and Constructibility	112
4.1.A	Causal Specifications and Causal Channels	112
4.1.B	Constructible Causal Channels	115
4.1.C	The Category of Causal Channels – Relative Constructibility	119
4.2	Notions of Contraction	123
4.2.A	The Standard Notion of Contraction	124
4.2.B	General Notions of Contraction	128
4.3	Causal Dilations	134
4.3.A	Causal Dilations – Definition and Basic Properties	134
4.3.B	The Causal-Dilational Ordering	142
4.4	Density Theorems and Rigidity	147
4.4.A	... in General	149
4.4.B	... in Cartesian Theories	154
4.4.C	... in CIT	155
4.5	Summary and Outlook	162
5	Rigidity and Quantum Self-Testing	163
5.1	Density Considerations in Purifiable Theories	169
5.1.A	An Abstract Characterisation of Pure Causal Dilations	170
5.1.B	A Partial Collapse of the Causal-Dilational Ordering	174
5.2	Rigidity in QIT – Quantum Self-Testing	177
5.2.A	The Standard Definition of Quantum Self-Testing	177
5.2.B	Classically Bound Dilations	179
5.2.C	The Bridge to Quantum Self-Testing	186
5.3	Some Implications of a Reformulation	193
5.3.A	Rigidity, Security and Extremality	194
5.3.B	Extraction of the Canonical State	196
5.3.C	Extraction of Canonical Measurements	197
5.4	Summary and Outlook	199
	Conclusion	201
	References	203

Introduction

For Everyone

A friend once told me that if you ask people whether they would rather be born in 100 years than live today, the vast majority say no. Ironically, if you then ask that majority whether they would prefer having lived in the world 100 years ago, they shake their heads again. Some of them probably realise their risk aversion.

Most people assess living standards in terms of health and wealth, freedom of choice, security to education, and the like; as such, the state of humanity has indeed been on a steady rise during the last centuries, if not millennia ([R+18, Pin18]). When gauging their lives, few might think of humankind’s enterprises within mathematics and the natural sciences. Nonetheless, these too have undergone tremendous improvements during the same time.¹

Popular consensus has it that *modern natural science* began less than 500 years ago, owing to the impact of significant figures like the notorious stone-dropping Galileo Galilei (1564—1642), a main proponent and pioneer of the paradigmatic conviction that knowledge about the physical world should be acquired by experimental observation and formulated in mathematical terms ([Mac17]).² As this program unfolded over the centuries, it instilled in its practitioners the aspiration to identify a small set of valid principles, *laws of Nature*, which were not to be further explained themselves, but from which all other observed phenomena could be logically derived. (For example, Newton’s laws of motion and gravitation are simple and universal, yet allows us to derive information sufficient to safely send members of our species away from our planet and land them 380.000 km away on the Moon.)

The idea of compressing all truth to a small set of postulates is an imprint from the *mathematical science*, which itself dates back more than 2000 years as the very institutionalisation of logical inference, similarly personified by the iconic geometry-obsessed Euclid of Alexandria (ca. 300 BC). In contrast to the natural sciences, mathematics refuses external physical inputs for certification of its initial axioms and for justification of its desired conclusions; this shifts the emphasis from the actual content of statements to the logical interdependencies among statements themselves. (For example, it is known ([BT24]) that the so-called *axiom of choice*³ formally implies the absurd statement that a solid ball can be dissected into finitely many pieces which can be reassembled into two solid balls each identical to the original.)

¹Pondering the relationship between these two developments is left as an exercise to the reader.

²It goes without saying that these paragraphs represent gross simplifications of the history; proper accounts could easily fill hundreds of pages. The entry [AH16] in the Stanford Encyclopedia of Philosophy gives a decent overview of the history of *‘the scientific method’*.

³A formal version of the seemingly obvious statement that given any non-zero number of bags each of which contains at least one marble, it is possible to form a collection containing precisely one marble from each bag.

It is difficult to find a word befitting of the scale of advancement that physics and mathematics have experienced since their conceptions – the study of their evolution is a science in itself. All scientific activities are bound to progress in a trivial sense simply because knowledge is accumulative over time, at least insofar as it is recorded; as such, advancement would seem only a matter of speed. However, as articulated by the science philosopher Thomas Kuhn (1922-1996) ([Kuh12]), transitions of a much more disruptive character occasionally occur in the sciences, and they cause profoundly new mentalities to ascend.

Roughly 100 years ago, both physics and mathematics found themselves at such bewildering points of disruption, after many years of marching steadily and obliviously towards them.

In mathematics, the continued process of rigorously formalising its concepts in the language of set theory had approached a landscape inhibited by more and more intriguing entities; objects such as Peano’s space-filling curve ([Pea90]), Weierstraß’ nowhere differentiable but everywhere continuous function ([Wei72]), and Cantor’s uncountable infinities ([Can84]) were proved to formally exist by abstract arguments, though their interpretation stretched the intuition of contemporaries. This growing balloon of peculiarities was building up tensions that forced mathematicians to question the very foundations of mathematical thinking, and exhibits such as Russell’s paradox around 1903 ([Rus]) eventually became so incriminating that the balloon cracked wide open. It was exposed that mathematics ultimately did not rest on solid formal grounds, and the so-called *Foundational Crisis of Mathematics* was burning at its fullest.

In physics, a revolution of remarkably similar significance was playing out. The physicist Albert A. Michelson⁴ had barely uttered the words ([Mic]) “[...] *it seems probable that most of the grand underlying principles [in Physics] have been firmly established*” in 1894, before, as if orchestrated by the Goddess of Irony, chaos began to sprout – among other things, Maxwell’s equations for the successful theory of electromagnetism seemed to display a conflict with the principle of Galilean relativity, the so-called ‘ultraviolet catastrophe’ plagued statistical mechanics, and Nature appeared to exhibit a weird discretized behaviour with respect to the emission of light from atoms. The tendency of these beauty flaws to resist elimination and rather conspire to unite in opposition was stressing and aggravating the physical community.

Eventually, thanks to exceptional thinkers in both disciplines, these tensions were unravelled and new paradigms arose in mathematics and physics alike.

Physicists had understood that we needed to profoundly revise some of our dearest conceptions about how the world works. Albert Einstein realised that the notions of *space* and *time* behaved in surprising and malleable ways defying thousands of years of human intuition, resolving not only in 1905 the problem from Maxwell’s equations ([Ein05]), but also providing over the years 1907–1915 a new and radically different theory of gravitation ([Ein16]). Today, his *theory of general relativity* remains a landmark in physics. Similarly, a list of people too long to reproduce – but including (Einstein and) Max Planck, Niels Bohr, Werner Heisenberg, Louis de Broglie and Erwin Schrödinger – progressively and collectively grasped through the period 1900–1930 that the discrete, quantised behaviour of Nature was covering over an underlying reality inherently different from the one we experience in our daily lives. This theory, which became known as *quantum physics*, was not only

⁴A similar quote is often falsely attributed to William Thomson (Lord Kelvin).

puzzling because it seemed best phrased in unexpectedly sophisticated mathematical realms of *complex linear algebra* and *Hilbert spaces* – it also challenged the very idea that questions about the properties of a physical object are meaningful.

Mathematicians, meanwhile, came to terms with their own crisis. They managed to repair the axioms of set theory and to make precise what formal reasoning in general *is*, thus effectively making the analysis of reasoning part of mathematics itself. Two of the most striking insights were due to Kurt Gödel around 1930, who demonstrated that a formal statement can be given a finite, checkable proof provided that it is true under every possible interpretation of its content (the *Completeness Theorem*, [Göd29]), but also that any potent system of reasoning will spawn formal statements which are true under some interpretations and false under others, and thus cannot be settled by checkable mathematical proofs (the *Incompleteness Theorem*, [Göd31]).⁵ In the primeval soup of these ideas about ‘checkable’ procedures – as contemplated also by contemporaries such as Alonzo Church, Alan Turing and Emil Post – eventually emerged the formal notions of *algorithms* and *computability*, which previously had only intuitive meaning. Not long after this, Claude Shannon in 1948 ([Sha48]) conceived of a mathematical theory of *information*, and on these two pillars – the theories of computation and information – was built the field of *computer science*. Amusingly, the desire to rigorously treat an abstract mathematical universe of infinite sets had led us to create the finitistic framework of computation, to which we now owe the existence of every digital computer on Earth (and in space).⁶

It may very well have been accidental that the two crises of the sciences raged at the same time. There is, however, a poetic glow to the fact that quantum theory and computer science were conceived and born simultaneously, and, as it turns out, destined to meet again later in life. *Information is physical*, said the physicist Rolf Landauer in 1961 ([Lan61]), and he thereby ushered an era devoted to the thesis that the theory of computation and information processing cannot be separated from physics, since the processing is ultimately executed by physical entities. The specific cocktail of *quantum information theory* was given shape in the early 1980s, when various people apprehended that quantum physics may affect the efficiency of computation (Richard Feynman [Fey82] and David Deutsch [Deu85]), that it fundamentally prohibits certain standard information-theoretic tasks such as duplicating information (William K. Wootters and Wojciech H. Zurek [WZ82]), and that it provides the means for cryptographic schemes not conceivable in classical information theory (Stephen Wiesner [Wie83], Charles Bennett and Gilles Brassard [BB84]).

Over the years, the field of quantum information theory grew larger, and though it is today still relatively young, it is a well-established area of research, tri-disciplinary between physics, computer science and mathematics. Whether we will ever be able to build an operational *quantum computer* which outperforms the most powerful digital computers is a question of intense dispute, but regardless of this a vast number of insights has been gained in information theory from the influence of quantum theory, and in quantum theory from the influence of information theory ([NC02]).

⁵For example, even some statements about the natural numbers $1, 2, 3, \dots$ and the arithmetic operations $+$ and \cdot cannot be decided– they simply have different truth values under different interpretations of what these entities mean. No matter how well we try to contain them by specifying how they interact with one another (by axioms such as *for all a, b, c , it holds that $(a + b) \cdot c = a \cdot c + b \cdot c$*), we will not succeed in eliminating undecidable statements.

⁶If nothing else, let this be a testament to the fact that basic research in mathematics should always be supported.

Now, one of the subfields of quantum information theory, known as *quantum foundations*, seeks to better understand what are the core principles of quantum information theory, and how can they be phrased in general, abstract terms. Research within this subfield attempts to define a mathematical universe of *physical theories* and to understand what makes quantum theory special among them.

This PhD thesis confines to that line of thought, and aims to recast a specific phenomenon in quantum information theory, *quantum self-testing*, in general, abstract terms. In doing so, it presents a new theory of so-called *dilations*, a concept which is well known in the field but has not been studied systematically before. Intuitively, a dilation of an information channel (an information channel being for example a device which accepts an input, computes the value of a function, and then returns an output) can be thought of as encoding ‘secret computations’ which take place in the course of our interaction with the information channel. The main conclusions of this thesis are that quantum self-testing can be understood in the language of such dilations, and that in fact many features of quantum information theory itself can be derived from principles phrased exclusively in terms of dilations. In developing the formalism necessary for these conclusions, it uses the mathematical language of *category theory*, a field which arose in the 1940s ([ML13]) and today has wide applicability. As such, it is my hope that some of the ideas and results presented here may find application also in pure mathematics, or other fields outside of quantum information theory.

No one can say with certainty what the future of science is like. Physics and mathematics – and computer science, the newcomer – will probably again face critical and disruptive periods. I am thankful for having lived 100 years after the groundbreaking work that led to the exciting scientific landscape of today, and I hope that this landscape will be even more exciting to those who gaze upon it 100 years from now.

For Someone

In order to understand what quantum self-testing is, and how it came to be, we must first return to the turbulent early years of quantum physics.

Though Einstein had played a major role in establishing quantum theory,⁷ he was famously non-pleased with the philosophical inclinations it seemed to require. One of the strange features of quantum theory is that it is probabilistic: When we measure the same property in two physical systems prepared identically, we might get different results. Quantum theory predicts the *probability distributions* which the measurement results follow, but generally cannot predict the exact values obtained. When quantum theory was still young, there were (at least) two different opinions about how to interpret this circumstance.⁸

According to the *realist position*, as held by Einstein, a measurement of a physical system reveals a property which the system already possessed in advance; though we may not know e.g. what the velocity (momentum) of a particle is before we measure it, the particle surely had a velocity prior to our measurement. As such, if quantum theory predicts randomness in

⁷In 1921, he was rewarded the Nobel prize in for his discovery of the *photoelectric effect*, which posited the quantised nature of light.

⁸See e.g. the witty descriptions in Ref. [Gri05], from which I have borrowed the terms ‘realist position’ and ‘orthodox position’.

measurement outcomes, it must be because the theory itself falls short of giving a complete description of reality.

On the other hand, according to the *orthodox position*, as defended by others, the randomness of quantum theory is fundamental and exempt from ordinary intuition. It simply makes no sense to speak of a physical system having a particular property before we measure it; this was the message of quantum theory, and it needed no fix. That idea was absurd to the realists, and in 1935, Einstein and colleagues Boris Podolsky and Nathan Rosen presented a thought experiment ([EPR35]) meant to expose that it was flawed.

In high-level terms, Einstein, Podolsky and Rosen argued that in certain experimental scenarios, the outcome of one measurement seemed to be definite (i.e. non-random), yet quantum theory failed to predict its value.

More precisely, they imagined a source emitting pairs of particles going off to two different sites, A and B. At each site, an experimenter is waiting for the respective particle and can choose to measure one of two properties⁹ of it, corresponding to measurements M_A^0 or M_A^1 at site A, and M_B^0 or M_B^1 at site B. Like any other physical theory, quantum theory has a notion of *state* of a physical system. A ‘physical system’ is a somewhat abstract concept, but for example the two emitted particles considered together form a physical system; as such, quantum theory mathematically associates to this system a set of possible states, ψ .¹⁰

Einstein, Podolsky and Rosen (a trio which became known as ‘EPR’) now pointed out that according to the mathematical formalism of quantum theory, there exists a state ψ , and measurements M_A^0 , M_A^1 , M_B^0 and M_B^1 , for which the theory predicts the following: If the two particles are in the state ψ and the measurement M_A^x ($x = 0, 1$) is performed at site A and yields outcome¹¹ y_A^x , then the outcome y_B^x of the measurement M_B^x (same x) at site B can be inferred with certainty from y_A^x , i.e. there are pre-determined functions f_0 and f_1 such that $y_B^0 = f_0(y_A^0)$ and $y_B^1 = f_1(y_A^1)$.

Now, if the sites A and B are sufficiently separated, and if the measurements are performed within suitable time spans, then the principle of special relativity (that no signal can travel faster than light) ensures that the measurement at site A cannot affect the measurement at site B, and vice versa. Consequently, they argued, it must be the case that the two measurement outcomes y_B^0 and y_B^1 *were really determined all along*. Nevertheless, quantum theory also says that the state ψ does not yield definite (non-random) values for both of measurements M_B^0 and M_B^1 – in fact, the measurements M_B^0 and M_B^1 have the property that every quantum state whatsoever will give random outcomes for at least one of them. They drew from this the conclusion that the quantum states ψ simply did not model all information about the particles, and they expressed the belief that it was possible to find another theory which resolved this problem.

However, their criticism backfired spectacularly. Three decades later, in 1964, the physicist John S. Bell ([Bel64]), inspired by their paper, astounded the scientific community by demonstrating that nothing could be done to repair the alleged incompleteness of quantum theory. His insight was striking, because it ultimately meant that the ‘realist’ and ‘orthodox’ positions towards quantum theory were not a matter of philosophical taste – quantum theory was plainly incompatible with the former, and this incompatibility could moreover be subjected to an *experimental test*.

⁹In their paper [EPR35], these two properties were the *momentum* or the *position* of the particle, but this is not essential.

¹⁰In the case of two particles, the states correspond more or less to functions called *wave functions*, but again this is inessential.

¹¹For example, in the case of momentum and position, y_A^x is some real number.

Bell considered a version of the EPR-scenario in which the relevant quantum state ψ of the two particles is the so-called *singlet state*, and for which the quantum measurements were measurements of so-called *spins* of the particles, meaning in particular that the possible measurement outcomes were $+1$ or -1 . More specifically, there exist according to the formalism of quantum theory, for any unit vector $v \in \mathbb{R}^3$, a ‘spin measurement in direction v ’, $M(v)$, and for unit vectors $v_A, v_B \in \mathbb{R}^3$ the spin measurements $M_A(v_A)$ at site A and $M_B(v_B)$ at site B are such that when measuring two particles in the singlet state, the probability of obtaining measurement outcomes $y_A, y_B \in \{+1, -1\}$ is given by $\frac{1}{4} - \frac{y_A y_B}{4} v_A \cdot v_B$, where $v_A \cdot v_B$ is the scalar product of v_A and v_B . Thus, if the four measurements M_A^0, M_A^1, M_B^0 and M_B^1 in the EPR-scenario are chosen as spin measurements, with $M_A^{x_A} = M_A(v_A^{x_A})$ and $M_B^{x_B} = M_B(v_B^{x_B})$ for $x_A, x_B \in \{0, 1\}$ and some unit vectors $v_A^0, v_A^1, v_B^0, v_B^1 \in \mathbb{R}^3$, then the probability distributions predicted by quantum theory are

$$P_{\text{quant.}}^{x_A, x_B}(y_A, y_B) = \frac{1}{4} - \frac{y_A y_B}{4} v_A^{x_A} \cdot v_B^{x_B} \quad (1)$$

Now, if there is, as Einstein, Podolsky and Rosen hoped, a complete theory meeting their standards of *realism*, then the measurement outcomes merely reveal pre-existing properties which can be described by ± 1 -valued random variables $Y_A^{x_A, x_B}$ (the measurement outcome at site A) and $Y_B^{x_A, x_B}$ (the measurement outcome at site B). If moreover this assumed theory is *local*, meaning that it complies to the non-signalling principle from special relativity, then, when A and B are suitably separated, Y_A cannot depend on x_B and Y_B not on x_A . As such, what we have is really four random variables, $Y_A^{x_A}$ for $x_A \in \{0, 1\}$, and $Y_B^{x_B}$ for $x_B \in \{0, 1\}$, and their probability distributions are simply

$$P_{\text{loc. real.}}^{x_A, x_B}(y_A, y_B) = \Pr(Y_A^{x_A} = y_A, Y_B^{x_B} = y_B). \quad (2)$$

What Bell then did was to derive an inequality that the probabilities (2) are bound to obey due to the mere fact that they arise as distributions of random variables as indicated, but which the probabilities (1) as predicted by the formalism of quantum theory, do not obey (for suitable choices of the vectors $v_i^{x_i}$). As such, *Bell’s inequality* by itself is not a result about quantum theory; it is about any theory which meets the requirements of *realism* (so as to infer the existence of random variables) and *locality* (so as to conclude the independence of the outcomes at site A from the measurement chosen at site B, and vice versa). The result about quantum theory is that it violates Bell’s inequality, and hence cannot be both local and realistic.¹²

From an abstract vantage point, the collections $P = (P_{\text{loc. real.}}^{x_A, x_B})_{x_A, x_B \in \{0, 1\}}$ of probability distributions which arise from local realism (i.e. which are of the form (2)) form a convex set,¹³ and Bell’s inequality corresponds to a *half-space* which confines this convex set. (This is similar to the way in which a pyramid is confined by half-spaces, four half-spaces corresponding to its tilted sides, and one to its horizontal bottom.) In honour of Bell, we generally refer to such half-space inequalities as *Bell-inequalities*. One of the simplest derivations of

¹²It is known, incidentally, that quantum theory can be given a realistic interpretation (i.e. one in which measurable properties are described by random variables) known as *de Broglie-Bohm theory*, or simply *Bohmian mechanics* ([Boh52]), but it is, of course, non-local.

¹³In the sense that if $P_1 = (P_1^{x_A, x_B})_{x_A, x_B \in \{0, 1\}}$ and $P_2 = (P_2^{x_A, x_B})_{x_A, x_B \in \{0, 1\}}$ are two such collections, and if $\alpha \in [0, 1]$, then $P = (P^{x_A, x_B})_{x_A, x_B \in \{0, 1\}}$ is also such a collection, with $P^{x_A, x_B} := \alpha P_1^{x_A, x_B} + (1 - \alpha) P_2^{x_A, x_B}$. This is because the weight α can be encoded as the success probability of a $\{0, 1\}$ -valued random variable Z , which we may include into the random variables giving rise to P_1 and P_2 .

a Bell-inequality is not Bell's original, but was given a few years later ([CHSH69]), by J. Clauser, M. A. Horne, A. Shimony and R. A. Holt. They first observed that the random variables $Y_i^{x_i}$ must with unit probability satisfy the inequality

$$(Y_A^0 + Y_A^1) \cdot Y_B^0 + (Y_A^0 - Y_A^1) \cdot Y_B^1 \leq 2. \quad (3)$$

Indeed, since $Y_A^{x_A}$ has values ± 1 , either the sum $Y_A^0 + Y_A^1$ or the difference $Y_A^0 - Y_A^1$ is ± 2 while the other is 0, and in each case the above expression then takes one of the values ± 2 (since also $Y_B^{x_B}$ is ± 1). But now, the inequality (3) must also hold for the expectation values, that is,

$$E(Y_A^0 \cdot Y_B^0) + E(Y_A^1 \cdot Y_B^0) + E(Y_A^0 \cdot Y_B^1) - E(Y_A^1 \cdot Y_B^1) \leq 2. \quad (4)$$

Each of these four expectation values can be re-expressed using the probabilities (2), since

$$\begin{aligned} E(Y_A^{x_A} \cdot Y_B^{x_B}) &= \Pr(Y_A^{x_A} = Y_B^{x_B}) - \Pr(Y_A^{x_A} \neq Y_B^{x_B}) \\ &= P_{\text{loc. real.}}^{x_A, x_B}(1, 1) + P_{\text{loc. real.}}^{x_A, x_B}(-1, -1) - P_{\text{loc. real.}}^{x_A, x_B}(1, -1) - P_{\text{loc. real.}}^{x_A, x_B}(-1, 1), \end{aligned} \quad (5)$$

but we may equivalently keep the inequality in the form (4). This is the so-called *CHSH-inequality*. To see that it can be violated in quantum theory, note that, by Eq. (1),

$$P_{\text{quant.}}^{x_A, x_B}(1, 1) + P_{\text{quant.}}^{x_A, x_B}(-1, -1) - P_{\text{quant.}}^{x_A, x_B}(1, -1) - P_{\text{quant.}}^{x_A, x_B}(-1, 1) = -v_A^{x_A} \cdot v_B^{x_B}, \quad (6)$$

so if quantum theory were locally realistic, the CHSH-inequality would read

$$-v_A^0 \cdot v_B^0 - v_A^1 \cdot v_B^0 - v_A^0 \cdot v_B^1 + v_A^1 \cdot v_B^1 \leq 2. \quad (7)$$

However, by choosing $v_A^0 = (-1, 0, 0)$, $v_A^1 = (0, -1, 0)$, $v_B^0 = (1/\sqrt{2}, 1/\sqrt{2}, 0)$ and $v_B^1 = (1/\sqrt{2}, -1/\sqrt{2}, 0)$, we easily compute that each of the four terms attain the value $1/\sqrt{2}$, so that the entire expression equals $2\sqrt{2}$ which is evidently larger than 2.

And now we come to quantum self-testing.

Though Bell's theorem was a shock, it was not a shock that extended to comatose paralysis. On the contrary, the result stimulated a renewed interest in the set-ups from the thought experiment envisioned by Einstein, Podolsky and Rosen. An obvious question was the following: *By how much can quantum theory violate the principles of local realism?*

A precise incarnation of this question was by how much quantum theory can violate the CHSH-inequality. This problem was solved in 1980 by the mathematician Boris Cirelson ([Cir80]), who showed that the violation is at most $2\sqrt{2}$, and also coined the term *behaviour* ([Cir93]) about the collections $P = (P^{x_A, x_B})_{x_A, x_B \in \{0,1\}}$ of probability distributions producible within a given theory. To prove that there was no *quantum* behaviour which exceeded the value $2\sqrt{2}$ was not simply a matter of optimising the expression (7) over unit vectors, as the formula (1) applies only to give those quantum behaviours which result from spin

measurements on particles in the singlet state. Rather, Cirelson’s argument was rooted in the general formalism of quantum theory, in terms of linear operators on Hilbert spaces.¹⁴ The set of quantum behaviours can be shown to be convex like the smaller collection of locally realistic behaviours, and *Cirelson’s inequality* (or *Cirelson’s bound*) is thus a quantum analogue of Bell’s inequality, namely an inequality corresponding to a half-space which confines the set of possible behaviours.

One of the questions raised by his work was the following: *What are the configurations of quantum states and quantum measurements whose behaviour reach the Cirelson bound $2\sqrt{2}$?*

A number of results ([SW87, PR92, BMR92, Cir93]) soon demonstrated that the value $2\sqrt{2}$ could in fact, in a certain sense, only be obtained by measuring the singlet state using the above spin measurements. While this was curious, it was mainly considered interesting for foundational reasons.

Probably the first person to acknowledge that the scenarios considered by Bell and Cirelson could have applications in the newly emerging field of quantum information theory was the physicist Artur Ekert. In 1991, he pointed out ([Eke91]) that because the values of the CHSH-expression which exceed 2 signify the lack of local realism, such values must certify *genuine* randomness in measurement outcomes, randomness which may be used for cryptographic purposes,¹⁵ since by Bell’s argument not even a potentially untrusted manufacturer of the measurement devices could have known it in advance. Using the fact that the particular value $2\sqrt{2}$ more or less uniquely determines the configuration of state and measurements, this idea was made even more explicit at the turn of the millennium, in the papers [MY98] and [MY04] by Dominic Mayers and Andrew Yao, who gave the name *self-testing* to this phenomenon, that devices could be used to ‘test themselves’. It is important to appreciate that the idea of exploiting quantum self-testing for applications constituted an almost paradigmatic change in mindset relative to the perspective of Bell and Cirelson. Whereas they had been thinking about trustworthy experimenters who wished to establish the supremacy of quantum theory over local realism, the new ideas took the point of view that the whole experimental set-up was like a *game*, a potentially vicious scheme in which untrustworthy agents had prepared an experiment whose purpose was to fool us to believe that a certain state was being subjected to certain measurements.

The mathematical definition of self-testing (which took its modern standard form in Ref. [MYS12]) is as follows:

In quantum information theory, the *physical systems* at sites A and B are modelled by (finite-dimensional) Hilbert spaces \mathcal{H}_A and \mathcal{H}_B over the complex field \mathbb{C} . When considering the two systems as one (as we did above), the associated Hilbert space is the tensor product, $\mathcal{H}_A \otimes \mathcal{H}_B$. A *state* on this system is modelled¹⁶ by a unit vector $\psi \in \mathcal{H}_A \otimes \mathcal{H}_B$. Finally, the *measurement* $M_i^{x_i}$ ($x_i = 0, 1$) is modelled¹⁷ by a so-called *projection-valued measure* (PVM) on \mathcal{H}_i , that is, by orthogonal projections $\Pi_i^{x_i}(1)$, $\Pi_i^{x_i}(-1)$ on \mathcal{H}_i (one for each

¹⁴Though I will not reproduce it here, Cirelson’s proof was not particularly technical; what he did was basically to establish an operator inequality.

¹⁵For example, it is often of interest to generate *shared randomness* so that one may use this to establish a secret key for encryption. However, it is of course important that this randomness is the genuine randomness that comes from quantum measurements, and not randomness which was known to the potentially adversarial manufacturer of the devices in advance.

¹⁶Two comments are in place here. First of all, only the so-called *pure* states are modelled as such (we will return to this shortly). Secondly, it is more correct to say that pure states are modelled by *rank-one projections* (or, what is equivalent, one-dimensional subspaces of the Hilbert space), since for any $\alpha \in \mathbb{C}$ of unit modulus, the vectors ψ and $\alpha\psi$ correspond to the same state.

¹⁷Again, there are more general kinds of measurements than PVMs, and we shall return to this point.

possible measurement outcome $y_i = 1, -1$), which sum to the identity operator on \mathcal{H}_i , $\Pi_i^{x_i}(1) + \Pi_i^{x_i}(-1) = \mathbb{1}_{\mathcal{H}_i}$. In summary, a configuration of states and measurements is defined by a triple (ψ, Π_A, Π_B) , where $\psi \in \mathcal{H}_A \otimes \mathcal{H}_B$ is a unit vector, and where $\Pi_A = (\Pi_A^{x_A})_{x_A \in \{0,1\}}$ and $\Pi_B = (\Pi_B^{x_B})_{x_B \in \{0,1\}}$ are collections of PVMs on \mathcal{H}_A and \mathcal{H}_B , respectively. Such a triple is called a (*tensor-product*) *quantum strategy*, and the formalism of quantum theory stipulates that it gives rise to the behaviour $P = (P^{x_A, x_B})_{x_A, x_B \in \{0,1\}}$ given by the inner products

$$P^{x_A, x_B}(y_A, y_B) = \langle \psi, [\Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B)] \psi \rangle, \quad x_i \in \{0, 1\}, y_i \in \{1, -1\}. \quad (8)$$

(The quantum behaviour (1) then arises from a suitable choice of such a quantum strategy. In particular, the *singlet state* corresponds to the vector $\psi = \frac{e_0 \otimes e_1 - e_1 \otimes e_0}{\sqrt{2}} \in \mathbb{C}^2 \otimes \mathbb{C}^2$, where (e_0, e_1) is the standard basis in \mathbb{C}^2 , and the *spin measurements* correspond to projections $\Pi_i^{x_i}(\pm 1)$ which project onto various 1-dimensional subspaces of \mathbb{C}^2 .)

Now, in the case of the CHSH-inequality, quantum self-testing formally means that if (ψ, Π_A, Π_B) is any quantum strategy for which the associated behaviour reaches the Cirelson bound $2\sqrt{2}$, then this strategy is ‘reducible’, or ‘equivalent’, in a certain sense, to a fixed, *canonical* strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, namely the one described by the singlet state and the spin measurements from above. Precisely, this reducibility criterion is expressed by the existence of so-called *residual* Hilbert spaces $\mathcal{H}_A^{\text{res}}$ and $\mathcal{H}_B^{\text{res}}$, a *residual* state $\psi^{\text{res}} \in \mathcal{H}_A^{\text{res}} \otimes \mathcal{H}_B^{\text{res}}$, and isometries $W_i : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}}$ such that

$$[W_A \otimes W_B][\Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B)]\psi = [\tilde{\Pi}_A^{x_A}(y_A) \otimes \tilde{\Pi}_B^{x_B}(y_B)]\tilde{\psi} \otimes \psi^{\text{res}}, \quad x_i \in \{0, 1\}, y_i \in \{1, -1\}. \quad (9)$$

In quantum theory, the local application of an isometry W_i is like a change of coordinates, so Eq. (9) is supposed to express that, up to such local changes of coordinates, the strategy (ψ, Π_A, Π_B) is really just the canonical strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, except possibly augmented by a state ψ^{res} which is shared between the two sites A and B, but which is not acted upon by the measurements.

Of course, the above definition generalises significantly beyond the CHSH-scenario. (In fact, the scenario considered by Mayers and Yao was a different one.) In general, we use the term (*bipartite*)¹⁸ *Bell-scenario* about a quadruple of finite non-empty sets (X_A, X_B, Y_A, Y_B) , with X_i corresponding to a set of possible measurement settings (‘inputs’) at site i , and Y_i a set of possible measurement results (‘outputs’) at site i .¹⁹ The definition of a quantum strategy for this Bell-scenario generalises in the obvious way, as a triple (ψ, Π_A, Π_B) , where $\Pi_i = (\Pi_i^{x_i})_{x_i \in X_i}$ is a collection of PVMs $(\Pi_i^{x_i}(y_i))_{y_i \in Y_i}$ on \mathcal{H}_i , i.e. orthogonal projections on \mathcal{H}_i summing to $\mathbb{1}_{\mathcal{H}_i}$. The behaviour of such a strategy is given as in Eq. (8). Moreover, we no longer talk of a specific inequality being saturated, we will simply say that the quantum behaviour P *self-tests the quantum strategy* $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, if any quantum strategy (ψ, Π_A, Π_B) with behaviour P is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, by means of a residual state ψ^{res} and isometries W_A and W_B as in Eq. (9).

The traditional definition of quantum self-testing as laid out above is mathematically unambiguous. The circumstance that motivated this PhD thesis is that its *operational*

¹⁸There is also a more or less obvious generalisation from two sites A and B to more sites, but we mostly consider the bipartite scenario.

¹⁹Unfortunately, the symbol Y_i is now used for a set, whereas we previously used it for a random variable; hopefully this causes no confusion.

significance is unclear. The most convincing argument for this is by observing that the definition is intimately intertwined with the very formalism of quantum information theory: It is carved in the stones of Hilbert spaces, linear operators and vectors, and it is not at all obvious how one would formulate it independently of this, despite the fact that the narrative of self-testing – namely, ‘there is essentially only one way of realising the behaviour P ’ – suggests that a general formulation should be possible.

Not only is a reformulation desirable in order to understand the significance of the phenomenon in other theories than quantum information theory. It is also desirable in order to better understand its significance within quantum information theory.

First of all, there is consensus among many that the Hilbert space formulation of quantum information theory is mysteriously obscure. Grounded in this opinion, a number of works (see e.g. Refs. [Har01, CDP11], and the book [CS16]) have demonstrated that remarkable reformulations of the theory are possible, namely formulations which do not refer to Hilbert spaces or linear algebra, but are cast instead in a universal language pertaining to general theories of information processing. It is conceivable that quantum information will eventually be best understood and studied from such an abstract point of view, and as such it is highly relevant to have a definition of self-testing which is compatible with that mode of abstraction.

Secondly, even within the usual formalism, the significance of a ‘quantum strategy’ is somewhat unclear. For example, the most general kind of quantum states are not represented by unit vectors, but so-called *density matrices*. Similarly, the most general kinds of quantum measurements are not represented by PVMs, but *POVMs* (*positive operator-valued measures*). Whereas a number of mathematical results imply that general states and POVMs can be seen as ‘arising’ in a precise way from pure states, respectively PVMs, the meaning of these results as they apply to quantum strategies is obfuscated, at best (this point is detailed in Chapter 5). In fact, if we really take literally the assumption that the experimental set-up in a self-testing scenario is crafted by untrusted agents, then it seems presumptuous to believe in the first place that the two devices establish their outputs by the simple process of sharing a quantum state and performing measurements on it.²⁰ Though this worry might seem ludicrous to those who find it intuitively clear that we can always standardise the form of more general ‘strategies’ to triple-form (ψ, Π_A, Π_B) , it is not clear how to give a formal argument for this without having an accepted notion of ‘general strategy’, and at any rate the meaning of the components ψ , Π_A and Π_B certainly does not crystallise in the process of this standardisation.

Lastly, in order for quantum self-testing to be a practical significance, it is important that self-testing results be *robust*, such that if a strategy (ψ, Π_A, Π_B) gives rise to a behaviour which is merely close to P , then it is close to being reducible to the canonical strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, in suitable senses of the word ‘close’. (The reason for this is not only that real experiments are prone to measurement errors, but also that the probabilities $P^{x_A, x_B}(y_A, y_B)$ can never be determined precisely, but only estimated based on finitely many observations.) It seems obvious that a sensible notion of ‘closeness’ should be operational (the standard choice from Ref. [MYS12] of using the Hilbert space norm of the difference between left and right hand sides in Eq. (9) is not); the problem of defining such a distance measure is left open by this thesis, but it is certainly necessary that there first exist an operational definition in the exact case.

²⁰For example, there could be an intricate procedure by which a sequence of local operations is first executed to decide which of several shared states to use in a subsequent protocol, etc.

For Anyone

In this thesis, I present a framework which offers a fundamentally different way of looking at quantum self-testing.

Let us consider the behaviour $P = (P^{x_A, x_A})_{x_A \in X_A, x_B \in X_B}$ observed in a Bell-scenario not simply as a collection of probability distributions on the outcome set $Y_A \times Y_B$, but as a dynamic *information channel* which receives local inputs x_A, x_B and produces local outputs y_A, y_B . We make no assumptions about the constituents of this channel, but merely assume that it can indeed be ‘constructed’ from basic constituents, and that it adheres to the locality assumption so as to produce at site i the output y_i given only the input x_i .

Now, instead of asking what the individual components of the channel might be, we ask a different and purely operational question:

What are the possible side-computations that may secretly be executed in the environment during our interaction with the channel?

To intuitively understand the idea of ‘side-computations in the environment’, three simple examples of information channels are helpful. They can be pictorially displayed as

$$-A-\boxed{f}-B- \quad , \quad -\{0,1\}-\boxed{T}-\{0,1\}- \quad , \quad -\mathbb{C}^2-\boxed{\text{id}}-\mathbb{C}^2- \quad , \quad (10)$$

of which the first represents the computation of an ordinary function $f : A \rightarrow B$, the third represents the identity channel on the system \mathbb{C}^2 in quantum information theory, and the one in the middle represents the ‘*bit refreshment*’ channel in classical information theory, which accepts as input any bit and outputs a uniformly random bit, regardless of the input.

(In each case, we tacitly assume that we can interact an arbitrary number of times with independent copies of the channel, so as to establish that the input-output behaviour of the channel is really as declared.)

Suppose we interact with the first channel, f . We do so by providing an input $a \in A$ to the *input interface* of the channel, and receiving the output $b = f(a)$ at the *output interface*. (For example, this is the kind of interaction we have with an ordinary digital computer.) Now, we imagine an *environment*, consisting of additional interfaces which we do not see, but which other agents – be they untrustworthy, or simply ‘Nature’ itself – can access. (For example, when interacting with a digital computer, there might be hidden interfaces within the computer, to and from which another party can send and receive information.) Provided that we really see the behaviour f at our interfaces, what computations might be going on simultaneously between these hidden interfaces?

It is quite easy to analyse this question. Of course, the environment may, simultaneously with our use of the channel, perform a computation which is completely independent, given by some function $g : C \rightarrow D$. In this case, the total channel describing the situation is the parallel composition $f \times g : A \times C \rightarrow B \times D$. More interestingly, the environment might *copy our input* $a \in A$, and use it to compute some function $g : A \rightarrow D$, so that the total channel is given by the function $(f, g) : A \rightarrow B \times D$, $a \mapsto (f(a), g(a))$; the value $f(a)$ is returned to us, but the value $g(a)$ is kept secret in the environment, possibly to be used in other computations. Even more generally, the environment can copy our input $a \in A$ in order to decide which of several functions $g_a : C \rightarrow D$ to apply on the side. In a sense, we are describing the obvious fact that if we want someone to compute a function value $f(a)$

for us, we cannot do this without sharing with them the value of a , and thereby allowing them to keep it in memory. On the other hand, it is intuitively clear that the value of the input a is the ‘strongest’ possible information the environment can extract from our use of the channel; every other side-computation can be ‘derived’ from the one that corresponds to copying the input.

The various channels that formalise side-computations in the presence of f will be called *dilations* of f . The notion of dilation is dual to that of a *marginal*, in the sense that a dilation of f is precisely a channel whose marginal is f .

Suppose instead we interact with the third channel, $\text{id}_{\mathbb{C}^2}$. This channel is in a sense the quantum analogue of the identity function from $\{0, 1\}$ to $\{0, 1\}$; it accepts as input a quantum state on the 2-dimensional system \mathbb{C}^2 , and does nothing to it. Again, we may ask about the various possible side-computations, or, more precisely, the various possible dilations of $\text{id}_{\mathbb{C}^2}$.


Readers unfamiliar with quantum information theory might think that, once again, the environment can keep a copy of our input in memory. This, however, is not the case, due to the so-called *No-Cloning Theorem* of quantum information theory ([WZ82]). According to this result, quantum information, in contrast to classical information, cannot be copied; in fact, every dilation of $\text{id}_{\mathbb{C}^2}$ must factor in the same way as the independent side-computations for f above,²¹ with the exception that the environment may *stall* its secret computations until we feed an input to our accessible interface. Hence, there is again a ‘strongest possible’ dilation of the channel $\text{id}_{\mathbb{C}^2}$, namely the one which simply registers in the environment that an input has been provided.

Finally, suppose we interact with the ‘bit refreshment’ channel, T . As it turns out, every dilation will be derivable from one of two possible dilations, but those two should be considered genuinely different. They intuitively correspond to two different *implementations* of T , which can easily be described in words. (Here, I use the word ‘implementation’ in an intuitive sense, but a fundamental point of the work in this thesis is that this intuitive notion can be formalised by the precise notion of dilation.)

The first such implementation of T is the obvious one; our input to the channel is discarded, and as output we are given a completely fresh random bit. This seems to be merely the description of the input-output behaviour of the channel, so it may come as a surprise that it could be implemented in other ways. Indeed it can, however:

In the second implementation, our input is not discarded, but instead the environment generates a random bit and uses it to decide whether to give us back as output our original input, or to give instead the opposite of our original input. From our point of view, the input-output behaviour of the channel is still a bit refreshment.


The two corresponding dilations are given as follows. The first one can be pictorially represented as



²¹We will establish this result by an abstract argument in Proposition 2.5.5 in Chapter 2.

where r denotes a uniformly random bit, where ‘Cop’ is the copy channel, where ‘1’ denotes a trivial system which is used to stall computation, and where the wiggly lines correspond to inaccessible interfaces belonging to the environment. As such, the diagram should be read as follows: A random bit r is generated and copied. One copy is stored in the memory of the environment, while the other is saved to be eventually revealed as output to us. When we provide an input to the accessible input interface, this input is recorded in the memory of the environment, and the release of r as output at the accessible interface is triggered.

The second implementation of the channel T corresponds to the dilation represented as



where ‘XOR’ denotes the *exclusive OR*, namely the function which output 0 if its two inputs bits are identical, and 1 if they are distinct. This time, a random bit r is generated and copied, one copy stored in memory, and the other used to decide whether, when our input bit comes it (and is copied to the memory of the environment), it should be given back to us as output as it is, or first flipped.

We will ultimately see (in Chapter 4) that the two dilations (11) and (12) correspond to formally distinct situations, but to appreciate the significance of this it is important to first realise a sense in which the two dilations are *equivalent*: In equations, the channel (11) can be written as

$$\delta_b \mapsto \frac{1}{2} \tilde{\delta}_b \otimes \delta_0 \otimes \tilde{\delta}_0 + \frac{1}{2} \tilde{\delta}_b \otimes \delta_1 \otimes \tilde{\delta}_1 \quad (13)$$

where δ_z denotes the classical state which is z with certainty (i.e. the degenerate probability distribution in z), and where, somewhat intermittently, we have used the symbol $\tilde{}$ to indicate information belonging to the environment. In words, on input $b \in \{0, 1\}$, the total output of the channel is the uniform mixture of the states $\tilde{\delta}_b \otimes \delta_0 \otimes \tilde{\delta}_0$ (corresponding to the random bit being 0) and $\tilde{\delta}_b \otimes \delta_1 \otimes \tilde{\delta}_1$ (corresponding to the random bit being 1). Likewise, the channel (12) is given equationally by

$$\delta_b \mapsto \frac{1}{2} \tilde{\delta}_b \otimes \delta_b \otimes \tilde{\delta}_0 + \frac{1}{2} \tilde{\delta}_b \otimes \delta_{b \oplus 1} \otimes \tilde{\delta}_1, \quad (14)$$

where \oplus denotes addition modulo 2. Now, if *in the environment* of the channel (13) one applies the channel $\tilde{\delta}_b \otimes \tilde{\delta}_k \mapsto \tilde{\delta}_b \otimes \tilde{\delta}_{b \oplus k}$, then one effectuates the change $\tilde{\delta}_b \otimes \tilde{\delta}_k \otimes \tilde{\delta}_k \mapsto \tilde{\delta}_b \otimes \tilde{\delta}_k \otimes \tilde{\delta}_{b \oplus k}$ and thereby obtains altogether the channel (14), as can be verified by comparing the outputs for $b = 0$ and $b = 1$. Conversely, if the channel $\tilde{\delta}_b \otimes \tilde{\delta}_k \mapsto \tilde{\delta}_b \otimes \tilde{\delta}_{b \oplus k}$ is applied in the environment of (14), the channel (13) is obtained.

However, this apparent equivalence of the two dilations is deceiving, because the demonstrated ‘equivalence’ ignores *causality*: The side-information encoded by the copies of the bit r is available in the environment before we feed our input to the accessible interface – the channel needed to go from e.g. (11) to (12) needs the copy of the random bit as well as a copy of our input, and therefore does not reproduce the correct causal structure in

(12), according to which the copy of r exists before our input was presented. As it turns out, no channel which preserves the causal structure will lead us between the two dilations (11) and (12). They should be considered different, formalising the intuition that the side-information in one dilation (pre-existing knowledge of which bit will be given as output) is information about something entirely different than the side-information in the second (pre-existing knowledge of whether or not the input will be flipped).

What I will do in the thesis is to demonstrate that quantum self-testing can be viewed on the same footing as the above examples. The various ‘implementations’ of the observed quantum behaviour (i.e. the various quantum strategies) appear as *causally structured dilations* (or, as we will say, simply *causal dilations*) of the behaviour channel, formalising various possible side-computations. The self-testing phenomenon is then more or less²² the existence of a strongest possible causal dilation, which moreover has the property that it holds no pre-existing side-information about the outputs at the accessible interface (in line with Ekert’s early observation).

Even though this connection to self-testing is one of the main contributions of thesis – and certainly the unique problem which motivated the project – it is important for me to stress that the emphasis in the thesis is first and foremost on initiating an abstract and general study of *dilations*. This is not only because the structure of dilations in a given theory turns out to be very interesting in its own right, but also because the fact that quantum self-testing can be interpreted as a dilational phenomenon implies, in my opinion, that the general study of dilations is necessary and valuable by extension. A systematic study of dilations has, to the best of my knowledge, not been attempted before; I hope that the results presented in this thesis will find interest, and that the strands left open will be even as interesting as to attract the curiosity and contemplation of others.

²²There are two caveats to this equivalence, but at this point it only makes sense to describe them in high-level terms: First of all, some dilations of the behaviour channel will be very strange and not be derivable from any dilation corresponding to a quantum strategy. The root of this problem is that quantum measurements turn out to have causal dilations which go against the intuition about what a measurement is (Example 4.3.9). We will exclude the strange dilations by introducing the notion of a *classically bound* dilation. Secondly, quantum self-testing actually also implies the existence of a certain simple representative in the equivalence class of the strongest possible dilation (this representative essentially corresponds to the canonical quantum strategy), but I conjecture that such a representative can always be found (Conjecture 5.2.22).

Structure of the Thesis

It is assumed that the reader of this section has already been through the general introduction. From this point onwards, the thesis contains the following elements:

- Preliminaries
- Chapter 1 – Theories
- Chapter 2 – Dilations
- Chapter 3 – Metric Theories
- Chapter 4 – Contractible Theories and Causal Dilations
- Chapter 5 – Rigidity and Quantum Self-Testing
- Conclusion

The section **Preliminaries** collects a few non-standard mathematical facts, mostly pertaining to the formalism of quantum information theory. Some of them are used quite extensively, and it is advisable for the reader to skim them in advance.

Each of the five chapters begins with a prelude, divided into two–four subsections, more or less following the self-explanatory pattern §1. *Introduction and Motivation* – §2. *Comparison to Existing Literature* – §3. *Contributions*. Each of them moreover concludes with a summary and the mentioning of several open ends and ideas for future work.

Below, I will briefly sketch the role of each chapter – it might afterwards be beneficial for the reader to read in series the preludes to the individual chapters. This not only gives a more precise idea of their content (under §1. *Introduction and Motivation*), but also details the relations to existing literature (under §2. *Comparison to Existing Literature*) and provides overviews of the technical contributions (under §3. *Contributions*), which it would not make much sense to reproduce here before the relevant concepts have been introduced.

First, we must in **Chapter 1** agree on a mathematical framework in which to even discuss physical theories, channels and dilations. This chapter reviews a variation on the *categorical framework* for discussing operational aspects of theories ([CS16]). More precisely, a *theory* will be modelled by a *symmetric monoidal category* in which the monoidal unit is *terminal* (these concepts will be explained and heavily exemplified). This framework constitutes a natural and minimal language in which to eventually make sense of the key ingredients required for the definitions we desire.

The role of Chapter 1 is mostly that of introductory review, and it contains only few original observations. My advice for readers who believe themselves familiar with the content of Chapter 1, would be to start by skimming the introductory section and the summary (Section 1.4).

In Chapter 2, *dilations* are introduced formally, but completely disregarding causal structure. As mentioned in the general introduction, ignoring causality may effectively change the relationship among dilations – for example, the two dilations of the bit refreshment channel will be equivalent in the *dilational ordering* of Chapter 2, but not in the *causal-dilational ordering* of Chapter 4, which will eventually be the correct formalisation of ‘derivability’ among causal dilations.

However, the causality-free setting of Chapter 2 turns out to be enlightening for other reasons, namely that it allows us speak of *dilational principles* which a given theory might comply to. The power of these principles will be demonstrated in Chapter 2 by deriving from them a number of features, which previously relied on specifics of the formalism of quantum information theory, or on probabilistic concepts.

The results of Chapter 2 for the most part play no role whatsoever in establishing the connection of the framework to quantum self-testing. Rather, they are included because they are interesting in their own right, and because I believe a thorough study of dilations has to begin in the special case where causality is trivial.

Chapter 3 contains a rather general definition of *metrics* on a theory, and discusses some properties which are natural to require of such metrics, with special emphasis on compatibility with dilations. This idea leads us to introduce the *purified diamond-distance*, which is a particularly well-behaved metric in quantum information theory, generalising the purified distance of Refs. [TCR10, Tom12].

The most important thing to say about this chapter is probably that I was not sure whether to include it in the thesis or not – the observations in Chapter 3 should be considered introductory and somewhat detached from the remainder of the thesis. Nevertheless, I believe that it adds a further perspective to the theory of dilations in Chapter 2, and that the open problem of extending the metric theory to the causal setting of the two later chapters might be one of the most interesting left from the thesis.

Chapter 4 is the longest chapter of the thesis. Here, we introduce the formal apparatus which we will use to speak about *causality*, in particular the notions of *causal dilations* (which formalise causally structured side-computations) and the *causal-dilational ordering* (which formalises the idea that some causal dilations are derivable from others). In theory, Chapter 4 is a ‘causal version’ of Chapter 2, but in practice things are more subtle.

First of all, owing to the causal structure, a new operation among channels arises, namely that of *contraction*. For example, in the channel (12) which we saw a few pages ago, the wiggly output wire at the bottom can be ‘contracted’ with the straight input wire, thus creating a new circuit;²³ it is not clear that this operation can be defined solely in terms of the total input-output behaviour of the channel (12) without reference to a particular circuit-representation, but as demonstrated in Chapter 4 it often can. This is important because we have to allow such contractions to occur in the environment when defining the causal-dilational ordering (‘derivability’).²⁴ As detailed later, we can view the contraction

²³There is no reason why one would want to do so in the particular channel (12), I am merely using it as example since we have not yet seen other causal channels than (11) and (12).

²⁴The example just given is not a contraction within the environment, as the input interface involved in

operation as an instance of abstract *notions of contraction*, which are related to so-called *traces* in symmetric monoidal categories ([JSV96]).

Secondly, it is relevant to prove a number of *stability results* to consolidate the concept of a causal dilation. For example, we will see the non-obvious fact that causal dilations are actually stable under contractions in the environment as described above, and we will see (less surprisingly) that the derivability relation is ‘composable’, e.g. in the sense that derivability is preserved under serial and parallel composition of channels.

Finally, since the causal-dilational ordering is more complicated than the dilational ordering, it will not be possible to replicate the precision of Chapter 2 in its analysis. This however gives rise to the idea of *rigidity* of a causal channel, which asserts the existence of a strongest possible causal dilation. This is the concept which we will ultimately link with quantum self-testing.

That link is established finally in **Chapter 5**. Here, we essentially identify the traditional quantum strategies as causal dilations from which all other (sensible) dilations are derivable. We then establish that self-testing as ordinarily conceived implies the equivalence in the causal-dilational ordering of all causal dilations corresponding to quantum strategies, and thus in particular the existence of a causal dilation from which all others can be derived and which has no pre-existing side-information about the outputs at the accessible interface.

This chapter also contains a surprising recharacterisation of quantum behaviours as those causal channels which admit a causally structured Stinespring dilation which is non-signalling.

The thesis ends with a common **Conclusion** which is kept rather short in light of the individual chapter conclusions.

the contraction does not belong to the environment, but we will see plenty of such examples.

Preliminaries

The thesis can in principle be read by someone with little knowledge about quantum theory, whereas it requires exposition to a wide range of various elementary mathematical constructs and ideas (graphs, metric spaces, mathematical standards of formalisation and proof, etc.). The thesis can be read without previous acquaintance with category theory, though superficial or intuitive understanding of the subject is beneficial.

A few notions which are needed in the thesis, but may not be covered by standard mathematical experience, are listed below. The reader with further interest in quantum information theory may consult the standard reference [NC02], or one of many excellent lecture notes available online, e.g. [Wat].

§1. Dirac Notation.

Many practitioners of quantum physics fancy the so-called ‘Dirac notation’ ([Dir81]) for vectors, whereas mathematicians tend to dislike it, perhaps in lack of a rigorous presentation. We will not need this notation overwhelmingly, but it is used on occasion. It can easily be introduced in a precise fashion.

Let \mathcal{H} be a Hilbert space over \mathbb{C} with inner product $\langle \cdot, \cdot \rangle$, which we take to be linear in its second argument (and thus anti-linear in the first). Given a vector $\psi \in \mathcal{H}$, let us denote by $|\psi\rangle$ (‘ket ψ ’) the linear map $\mathbb{C} \rightarrow \mathcal{H}$ given by $z \mapsto z\psi$ and by $\langle\psi|$ (‘bra ψ ’) the linear map $\mathcal{H} \rightarrow \mathbb{C}$ given by $\phi \mapsto \langle\psi, \phi\rangle$. One easily checks by definition of adjoints that $|\psi\rangle^* = \langle\psi|$ and $\langle\psi|^* = |\psi\rangle$. By virtue of the Riesz representation theorem every linear functional $\mathcal{H} \rightarrow \mathbb{C}$ is of the form $\langle\psi|$ for some $\psi \in \mathcal{H}$. The merits of these bizarre-looking conventions are now threefold:

- In equations, we can replace vectors $\psi \in \mathcal{H}$ by their kets $|\psi\rangle$ without disturbing the content. For example, it is easy to check that $z_1 |\psi_1\rangle + z_2 |\psi_2\rangle = |z_1\psi_1 + z_2\psi_2\rangle$ for $z_1, z_2 \in \mathbb{C}$, and that if $A : \mathcal{H} \rightarrow \mathcal{K}$ is a linear operator with $A\psi = \phi$ then $A|\psi\rangle = |\phi\rangle$. As a result, we can ultimately forget about the vectors ψ and think of the kets $|\psi\rangle$ as fundamental and ‘belonging’ to \mathcal{H} . The corresponding bras $\langle\psi|$ can be thought of as simply alternative representations of the same underlying objects, ‘belonging’ to the dual space \mathcal{H}^* .
- The operator $\langle\phi|\psi\rangle := \langle\phi| \circ |\psi\rangle$ is the linear map $\mathbb{C} \rightarrow \mathbb{C}$ given by $z \mapsto \langle\phi, \psi\rangle z$, naturally identified with the number $\langle\phi, \psi\rangle$ itself. This justifies the suggestive identity $\langle\phi|\psi\rangle = \langle\phi, \psi\rangle$ and makes explicit mentioning of an inner product on \mathcal{H} unnecessary; it has effectively been merged with the notation for vectors.
- We have a succinct way of writing the operator $|\phi\rangle\langle\psi| := |\phi\rangle \circ \langle\psi|$ given by $\chi \mapsto \langle\psi, \chi\rangle\phi$; in particular, for $\psi \in \mathcal{H}$ a unit vector, we have a succinct notation for the projection onto the subspace spanned by ψ , namely $|\psi\rangle\langle\psi|$.

Now, once the bra-ket notation gains a life of its own, it is tempting to forget so much about the initial vectors that we insert into the symbol $|\ \rangle$ an arbitrary name for the ket rather than an actual vector; in particular, the kets in the standard basis of \mathbb{C}^n are customarily named $|0\rangle, |1\rangle, \dots, |n-1\rangle$. (As such, $|0\rangle$ denotes not, as the previous convention would dictate, the zero operator $\mathbb{C} \rightarrow \mathcal{H}$.)

In a similar spirit of inconsistency, we will actually from now on use letters ψ, ϕ, \dots from the end of the Greek alphabet to denote *rank-1 projections* (i.e. orthogonal projections onto 1-dimensional subspaces), and then write $|\psi\rangle, |\phi\rangle, \dots$ for *vector representatives*, i.e. unit vectors in the corresponding subspaces. This convention not only overwrites the above, but also abuses notation, since ‘ $|\psi\rangle$ ’ is only determined from ‘ ψ ’ up to multiplication by a complex number α of unit modulus; however, whenever we use in an equation the ‘vector representative’ $|\psi\rangle$ of the projection ψ , it will be the case that the content of the equation is insensitive to the choice of the scalar α .

§2. General CPTP Maps and Their Representations.

As we will see, systems in quantum information theory are modelled by (separable) Hilbert spaces, and the processing of quantum information between such systems by *completely positive trace-preserving (CPTP)* maps on associated operator algebras. We will mostly be interested in the case where the Hilbert spaces are finite-dimensional, but the definitions are presented generally below.

Complete Positivity (CP). Given a Hilbert space \mathcal{H} , recall that an operator A on \mathcal{H} is said to be positive, denoted $A \geq 0$, if it can be written in the form B^*B for some operator B on \mathcal{H} , with B^* denoting the adjoint (Hermitian conjugate) of B . Given Hilbert spaces \mathcal{H} and \mathcal{K} , a linear map $\Lambda : B(\mathcal{H}) \rightarrow B(\mathcal{K})$ from (bounded) operators on \mathcal{H} to (bounded) operators on \mathcal{K} is called *positive* if $\Lambda(A) \geq 0$ for all $A \geq 0$.

The map Λ is called *completely positive* if for any Hilbert space \mathcal{R} , the linear map $\Lambda \otimes \text{id}_{\mathcal{R}} : B(\mathcal{H}) \otimes B(\mathcal{R}) \rightarrow B(\mathcal{K}) \otimes B(\mathcal{R})$ is positive. (Observe the isomorphisms $B(\mathcal{H}) \otimes B(\mathcal{R}) \cong B(\mathcal{H} \otimes \mathcal{R})$ and $B(\mathcal{K}) \otimes B(\mathcal{R}) \cong B(\mathcal{K} \otimes \mathcal{R})$.)

Clearly, any completely positive map is positive, but there are positive maps which are not completely positive, for example the map $B(\mathbb{C}^2) \rightarrow B(\mathbb{C}^2)$ which maps a 2×2 matrix to its transpose.

For any linear operator $S : \mathcal{H} \rightarrow \mathcal{K}$, the map $B(\mathcal{H}) \ni A \mapsto SAS^* \in B(\mathcal{K})$ is an example of a completely positive map; it is called *conjugation by S* . We will be mostly interested in the case where S is an isometry (i.e. satisfies $S^*S = \mathbf{1}_{\mathcal{H}}$).

Trace-Preservation (TP). Let $B_1(\mathcal{H}) \subseteq B(\mathcal{H})$ denote the subspace of trace class operators on \mathcal{H} . (When \mathcal{H} is finite-dimensional, $B_1(\mathcal{H}) = B(\mathcal{H}) = \text{End}(\mathcal{H})$, the space of all linear operators on \mathcal{H} .) Let us call a linear map $\Lambda : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$ *trace-preserving* if $\text{tr}(\Lambda(A)) = \text{tr}(A)$ for all $A \in B_1(\mathcal{H})$, and if Λ is continuous w.r.t. the trace norm $\|\cdot\|_1$, given by $\|A\|_1 = \text{tr}(|A|) = \text{tr}(\sqrt{A^*A})$ (when \mathcal{H} is finite-dimensional, the continuity requirement is void).

Every isometric conjugation $A \mapsto SAS^*$ (restricted to $B_1(\mathcal{H})$) is an example of a trace-preserving map, since $\text{tr}(SAS^*) = \text{tr}(S^*SA)$ by cyclicity of the trace. Another example of a trace-preserving map is the trace itself, that is, the map $\text{tr} : B_1(\mathcal{H}) \rightarrow B_1(\mathbb{C}) \cong \mathbb{C}$.

CPTP Maps. A linear map $\Lambda : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$ is called *CPTP* if it is completely positive and trace-preserving. Both isometric conjugations and traces are examples of CPTP

maps. Moreover, the serial composition of any two CPTP maps is CPTP, and the tensor product of any two CPTP maps is also CPTP (observing again isomorphisms of the sort $B_1(\mathcal{H}_1) \otimes B_1(\mathcal{H}_2) \cong B_1(\mathcal{H}_1 \otimes \mathcal{H}_2)$).

Kraus Representations. It can be shown that $\Lambda : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$ is CPTP if and only if there exists a (countable) family $(K_i)_{i \in I}$ of linear operators $K_i : \mathcal{H} \rightarrow \mathcal{K}$ such that $\sum_{i \in I} K_i^* K_i = \mathbb{1}_{\mathcal{H}}$ and

$$\Lambda(A) = \sum_{i \in I} K_i A K_i^* \quad \text{for all } A \in B_1(\mathcal{H}). \quad (15)$$

A representation such as (15) is called a *Kraus representation* of Λ .

Stinespring Representations. It can be shown that $\Lambda : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$ is CPTP if and only if there exists a Hilbert space \mathcal{E} and an isometry $S : \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{E}$ such that

$$\Lambda(A) = [\text{id}_{B_1(\mathcal{K})} \otimes \text{tr}_{\mathcal{E}}](SAS^*) \quad \text{for all } A \in B_1(\mathcal{H}), \quad (16)$$

where $\text{tr}_{\mathcal{E}} : B_1(\mathcal{E}) \rightarrow \mathbb{C}$ denotes the trace on $B_1(\mathcal{E})$. This statement is known as *Stinespring's Dilation Theorem* ([Sti55]) and the isometric conjugation $A \mapsto SAS^*$ in the representation (16) is known as a *Stinespring dilation* of Λ (sometimes, the term ‘Stinespring dilation’ refers to the isometry S itself).

Stinespring’s theorem also contains a clause of uniqueness up to isometries, that is, if $S : \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{E}$ and $S' : \mathcal{H} \rightarrow \mathcal{K} \otimes \mathcal{E}'$ are two isometries which both define a Stinespring dilation of Λ , and if $\dim \mathcal{E} \leq \dim \mathcal{E}'$, then there exists an isometry $W : \mathcal{E} \rightarrow \mathcal{E}'$ such that $(\mathbb{1}_{\mathcal{K}} \otimes W)S = S'$.

In the special case where $\mathcal{H} \cong \mathbb{C}$, a CPTP map from $B_1(\mathcal{H}) \cong \mathbb{C}$ to $B_1(\mathcal{K})$ is called a *state*, and Stinespring dilations of the state ϱ (which are defined by isometries $\mathbb{C} \rightarrow \mathcal{K} \otimes \mathcal{E}$, i.e. by unit vectors $|\psi\rangle \in \mathcal{K} \otimes \mathcal{E}$) are typically called *purifications* of ϱ .

§3. Special CPTP Maps and Their Representations.

Classical Systems. Given a countable (often finite) set X , the associated Hilbert space of square-summable sequences $\ell^2(X)$ (which coincides with \mathbb{C}^X when X is finite) is the quantum analogue of the set X . We will call Hilbert spaces of the form $\ell^2(X)$ *classical*. Any (separable) Hilbert space \mathcal{H} is isomorphic to a classical one, but for classicality we require strict equality. In effect, this is a matter of there being chosen a preferred orthonormal basis in \mathcal{H} , indeed $\ell^2(X)$ (\mathbb{C}^X) has the canonical orthonormal basis $(|e_x\rangle)_{x \in X}$, where e_x is the sequence given by $e_x(x) = 1$ and $e_x(x') = 0$ for $x' \neq x$. By abuse of notation, we write the basis elements $|e_x\rangle$ as $|x\rangle$. In quantum information theory, the basis $(|x\rangle)_{x \in X}$ is often called the *computational basis*.

Decoherence and Classical States. Given a function²⁵ $f : X \rightarrow Y$ between countable sets, it can be naturally represented as a CPTP map $\hat{f} : B_1(\ell^2(X)) \rightarrow B_1(\ell^2(Y))$, namely the one defined by $\hat{f}(A) = \sum_{x \in X} \langle x | A | x \rangle |f(x)\rangle \langle f(x)|$, which in particular satisfies $\hat{f}(|x\rangle \langle x|) = |f(x)\rangle \langle f(x)|$.

²⁵Some mathematicians use the term ‘map’ in place of ‘function’, reserving the term ‘function’ for maps which take values in \mathbb{R} or \mathbb{C} . We do not employ this convention.

As such, the representation of the identity function $x \mapsto x$ on X is the CPTP map Δ_X given by $\Delta_X(A) = \sum_{x \in X} \langle x|A|x\rangle |x\rangle\langle x|$. We will call Δ_X the *decoherence map associated to X* , and a state ϱ on $\ell^2(X)$ is called *classical* if $\Delta_X(\varrho) = \varrho$. It is a simple exercise to verify that ϱ is classical precisely if $\varrho = \sum_{x \in X} p(x) |x\rangle\langle x|$ for some probability density $p : X \rightarrow [0, 1]$, so classical states on $B_1(\ell^2(X))$ can be identified with probability distributions on X .

Quantum Measurements. If $M : B_1(\mathcal{H}) \rightarrow B_1(\ell^2(Y))$ is a CPTP map whose domain is represented by a classical system, we say that M has *classical outcomes* if $\Delta_Y \circ M = M$. More commonly, such a CPTP map is called a *measurement on \mathcal{H} with outcomes in Y* . Using the Kraus representation of M , it is easy to verify that if M is classical then there exists a *Positive Operator-Valued Measure (POVM)* on \mathcal{H} , i.e. a family $(E_y)_{y \in Y}$ of positive operators E_y on \mathcal{H} with $\sum_{y \in Y} E_y = \mathbb{1}_{\mathcal{H}}$, such that

$$\Lambda(A) = \sum_{y \in Y} \text{tr}(E_y A) |y\rangle\langle y| \quad \text{for all } A \in B_1(\mathcal{H}); \quad (17)$$

conversely, any POVM $(E_y)_{y \in Y}$ defines a measurement. Thus, we can identify measurements with POVMs.

A measurement $M : B_1(\mathcal{H}) \rightarrow B_1(\ell^2(Y))$ is said to be *projective* if the associated POVM $(E_y)_{y \in Y}$ is a PVM (Projection-Valued Measure), i.e. if E_y is a projection on \mathcal{H} for all $y \in Y$.

Naimark's Theorem For any measurement $M : B_1(\mathcal{H}) \rightarrow B_1(\ell^2(Y))$, there exists a Hilbert space \mathcal{K}^{Nai} , a projective measurement $M^{\text{Nai}} : B_1(\mathcal{H} \otimes \mathcal{K}^{\text{Nai}}) \rightarrow B_1(\ell^2(Y))$ and a pure state ϕ^{Nai} on \mathcal{K}^{Nai} , such that

$$M = M^{\text{Nai}} \circ (\text{id}_{B_1(\mathcal{H})} \otimes \phi^{\text{Nai}}). \quad (18)$$

This statement is known as *Naimark's (Dilation) Theorem* ([Nai40]), and the representation (18) is called a *Naimark representation* of M . (Sometimes, Naimark's name is transcribed as 'Neumark'.)

Ensembles of CPTP Maps. The decoherence maps Δ_X facilitate more refined notions of classicality too. In particular, if $\Lambda : B_1(\mathcal{H}) \otimes B_1(\ell^2(X)) \rightarrow B_1(\mathcal{K})$ is a CPTP map for which a factor of the domain is a classical system, we may say that Λ is classical on this factor if $\Lambda \circ (\text{id}_{B_1(\mathcal{H})} \otimes \Delta_X) = \Lambda$. It is easy to verify that this is the case precisely if there exists a family $(\Lambda^x)_{x \in X}$ of CPTP maps $\Lambda^x : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$ such that

$$\Lambda(A \otimes B) = \sum_{x \in X} \Lambda^x(A) \langle x|B|x\rangle \quad \text{for all } A \in B_1(\mathcal{H}), B \in B_1(\ell^2(X)). \quad (19)$$

Thus, to specify a CPTP map $\Lambda : B_1(\mathcal{H}) \otimes B_1(\ell^2(X)) \rightarrow B_1(\mathcal{K})$ which is classical on $B_1(\ell^2(X))$ is precisely to specify an ensemble of CPTP maps $\Lambda^x : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$, indexed by $x \in X$. In this case we will often use the terminology that Λ 'measures' or 'reads off' the classical value x and applies the according map Λ^x .

§4. Miscellaneous.

Pre-Orders. Let P be a class of objects (e.g. a set). Recall that a relation on P is a subclass R of $P \times P$, and that we tend to write pRq rather than $(p, q) \in R$. Recall that a

relation R is reflexive if pRp for all $p \in P$, and transitive if for all $p, q, r \in P$ the conditions pRq and qRr imply the condition pRr . A relation R is called a *pre-order* if it is reflexive and transitive.²⁶ Pre-orders are typically denoted with directional symbols, such as \geq, \succeq, \supseteq etc., with the implicit convention that mirroring the symbol inverts the order (e.g. ' $p \leq q$ ' means $q \geq p$). If the conditions pRq and qRp imply $p = q$, the pre-order is commonly called a *partial order*. Most pre-orders of interest to us will not have this property, but in general the relation \sim_R defined by $p \sim_R q \Leftrightarrow pRq \wedge qRp$ is an equivalence relation on P .

Special Elements of Pre-Orders. Let \geq be a pre-order on P . An element $u \in P$ is called a *largest (greatest) element* if $u \geq p$ for all $p \in P$. An element $m \in P$ is called a *maximal element* if for all $p \in P$ with $p \geq m$ it also holds that $m \geq p$ (i.e. $m \sim_{\geq} p$).

Any largest element is a maximal element, but not necessarily conversely. For instance, in the pre-order on $\{0, 1, 2\}$ defined precisely by the reflexive conditions and the two conditions $1 \geq 0$ and $2 \geq 0$, both 1 and 2 are maximal though neither is largest. *Smallest (least) elements* and *minimal elements* are defined dually, by inverting the order.

Given a subclass $P_0 \subseteq P$ we can naturally restrict the pre-order to that subclass, and we may consequently speak of largest and maximal (respectively smallest and minimal) elements *in* P_0 by minding this restriction. For instance, in the previous example, the element 2 is a largest element in $\{0, 2\}$.

Dense Subclasses of Pre-Orders. A subclass $D \subseteq P$ is called *dense in* P , if for any $p \in P$ some $d \in D$ satisfies $d \geq p$. (As such, 'dense' means 'dense at the top'.) By extension, a class D is called dense in $P_0 \subseteq P$, if $D \subseteq P_0$ and D is dense in P_0 considered as a pre-order on its own.

This terminology has been imported from the subject of forcing in axiomatic set theory, cf. Ref. [Kun80].

The Schmidt Decomposition. If $|\psi\rangle \in \mathcal{H}_1 \otimes \mathcal{H}_2$ is any vector in a tensor-product of Hilbert spaces, then there exist a family $(p(j))_{j \in J}$ of strictly positive numbers, and orthonormal systems $(|\psi_1(j)\rangle)_{j \in J}$ in \mathcal{H}_1 and $(|\psi_2(j)\rangle)_{j \in J}$ in \mathcal{H}_2 , such that

$$|\psi\rangle = \sum_{j \in J} \sqrt{p(j)} |\psi_1(j)\rangle \otimes |\psi_2(j)\rangle, \quad (20)$$

and $\sum_{j \in J} p(j) = \|\psi\|^2$. An expression of the form (20) is called a *Schmidt decomposition* of $|\psi\rangle$. In fact, the cardinality $|J|$ is unique, as is the family $(p(j))_{j \in J}$ (up to permutation). They are referred to as the *Schmidt rank* and *Schmidt coefficients* of $|\psi\rangle$, respectively.

²⁶It is worth observing that an equivalence relation is thus a pre-order which is additionally symmetric, meaning that pRq implies qRp .

Chapter 1

Theories

§1. Introduction and Outline.

In this chapter we set up a mathematical framework for investigating general physical theories. The chapter has three sections, all of which serve mainly as review. It contains no original observations, except for a few examples in Section 1.2, the comment about functors in Remark 1.1.12, the failure of the Cantor-Schröder-Bernstein property as described in Example 1.2.7, and the definition of ‘normal’ theories (Definition 1.3.7).

General Theories. The first item on the agenda is to define mathematically what is meant by a (*physical*) *theory*. We will define a theory as a certain type of mathematical structure (like a group, or a measurable space), and as usual the concept is abstracted from a selection of prominent examples. One example with which every reader will be familiar is the theory of *sets and functions*:

We may think of a set X as a (*physical*) *system*, and think of a function $f : X \rightarrow Y$ as a (*physical*) *transformation* from the system X to the system Y . Functions can be composed *serially*, one following the other, by forming from $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ the composite $g \circ f : X \rightarrow Z$. But they can also be composed *parallelly*, one next to the other; given functions $f_1 : X_1 \rightarrow X_2$ and $f_2 : X_2 \rightarrow Y_2$, we have a function $f_1 \times f_2 : X_1 \times X_2 \rightarrow Y_1 \times Y_2$ defined by $(f_1 \times f_2)(x_1, x_2) = (f_1(x_1), f_2(x_2))$. The parallel composition of functions involves a composition of the underlying systems (sets), namely the formation of the product set $Z_1 \times Z_2$ from the individual sets Z_1 and Z_2 .

In general, a *theory* will be a structure encompassing systems, transformations, and notions of composing transformations serially and parallelly. Whereas the theory of sets and functions is undoubtedly the example known to most readers, the two most important examples for us is *Classical Information Theory*, **CIT**, and *Quantum Information Theory*, **QIT**. The systems of **CIT** are (finite) sets and its transformations are so-called *classical channels* (Markov kernels) between them, which can be thought of as probabilistic functions. The systems of **QIT** are (finite-dimensional) Hilbert spaces and its transformations are so-called *quantum channels* (CPTP maps) between them. These two theories are described in Section 1.1 (Example 1.1.10 and Example 1.1.10), where also the general definition of a theory (Definition 1.1.6) and some surrounding terminology is provided.

Specific Theories. Section 1.2 comprises a large collection of further examples of theories. Some of these will be merely curious, helping to paint a landscape, but most will serve to illustrate points later. I have categorised the examples into classes, and included

among them many mathematical ones (though none of them very technical), which admittedly stretch the boundaries of what one might call a ‘physical’ theory. In particular, the example classes include all categories with finite products (Section 1.2.B), and monoid-like structures related resource theories in the sense of Ref. [CFS16] (Section 1.2.C).

It is not necessary for the reading of the thesis to be intimately acquainted with all of the examples presented in Section 1.2, but it likely yields an elevated reading experience to familiarise oneself with one example from each class.

Pictorial Syntax. The mathematical structure that defines a theory is an algebraic entity equipped with two binary operations, serial and parallel composition of transformations. This complexity can make equations difficult to interpret and consequently obscure intuition. In the last section of the chapter, Section 1.3, we review a widely used *pictorial syntax* ([Sel10]) for arguing about transformations in a theory. This replaces algebraic expressions by pictures, and may thus tremendously clarify algebraic manipulations. We shall use the pictorial syntax in many instances throughout the thesis, and have already seen it exemplified in the general introduction when discussing dilations of the ‘bit refreshment’ channel.

In Section 1.3.B, we formally introduce the concepts of *interfaces* and *channels*, as opposed to *systems* and *transformations*. The distinction between the two (which arise from the finer points of the pictorial syntax, but which is often ignored in other presentations) might seem at this point overly formal, but it will be important later on, in particular in Chapter 4.

§2. Comparison to Existing Literature.

On the Definition of a Theory. Our model of ‘theories’ does not aim to capture every single construct that a physicist might call a theory (for example, Einstein’s theory of special relativity [Ein05] is not a theory in that sense). Rather, it aims to capture *operational* aspects of theories, in line with a ‘pragmatist’ tradition of physics (cf. Ref. [CS16]): The interest is not in the ultimate explanation about what or why Nature is, but instead in what intelligent beings can and cannot do with the physical systems and physical transformations handed to them.

Roughly speaking, there are two pillars of mathematical frameworks which intend to capture operational aspects of theories. One is the *categorical* pillar (pioneered by Refs. [AC04, Sel04, Bae06]), according to which the fundamental objects of interest are systems and transformations which can be serially and parallelly combined, as outlined above. It uses *symmetric monoidal categories* ([ML13]) as a model for theories. The second pillar is the *convex* or *probabilistic* framework, often in the incarnation of *generalised probabilistic theories* ([BW16]). In this framework, an underlying categorical structure is often implicitly present ([Bar07, BW11]), but the emphasis is on probabilities and convexity, and the study of how *state spaces* morph under the composition of systems. There has been work which quite explicitly merges the categorical and probabilistic pillars (e.g. Refs. [CDP10, Har10]), and the book [CS16] gives a fairly recent overview of various tendencies within the field.

In developing the theory presented in this thesis, I have made an effort to stay within a purely categorical framework. This is not (only) because it is more general than merged frameworks, but also because almost all defined concepts are completely independent of probabilistic structure. Precisely, the definition chosen here for a theory (Definition 1.1.6) is that of a *symmetric monoidal category in which the unit object is terminal*. As such, theories in our sense are more restricted than those modelled by arbitrary symmetric monoidal

categories ([CDKW16]), but on the other hand do not assume additional structures like *dagger compactness* or **-autonomy*, which were and still are instrumental ingredients in some works (e.g. Refs. [AC04, KU17]).¹

It is well-established that symmetric monoidal categories with terminal unit object can be interpreted as theories in which future events cannot signal to the past, and as such these are often termed *causal theories* ([CDP10, CL13, Coe14]). In fact, our notion of theories exactly coincides with that of a *causal deterministic theory* in the words of Ref. [CDP10]. However, in other treatments the terminality assumption is mostly accompanied by further standing assumptions, and in practice this renders the scope of those treatments smaller than the one presented here. Accordingly, many of the examples in Section 1.2 would be ruled out in other works (for example in Ref. [CDP10] the assumption of ‘non-determinism’ rules out our cartesian theories, and the assumption that transformations are determined by their action on states rules out our thin theories).

In the mathematical literature our notion of theories are commonly referred to as (*symmetric*) *semi-cartesian categories*, or *monoidal categories with projections* (see [Sem], and the comments between Remarks 2.3 and 2.4 in Ref. [Fri20]), but here a systematic study of the class also seems to be absent.

On the Use of the Pictorial Syntax. Ref. [Sel10] reviews a large class of pictorial syntaxes for monoidal categories, including the one for symmetric monoidal categories, attributed to [Pen71]. Nowadays, its use and interpretation are fairly standardised, with minor differences in the choice of layout (e.g. some prefer that diagrams be read from top to bottom rather than left to right). As observed pedantically in Section 1.3.B, however, the ambiguity in its representation of composite systems means that the pictures do not strictly correspond to transformations between systems, but rather to transformations between *interfaces*, that is, tuples of systems labelled by port names. We shall use the term *channels* about such transformations. The distinction is minute and for most purposes insignificant (which is probably why it has not been pointed out before), but we need it for the precise definition of marginalisation and dilations (Chapter 2), and it will become even more pressing in Chapter 4.

¹Somewhat confusingly, treatments employing dagger compactness tend to define the transformations in quantum theory as linear operators between Hilbert spaces, rather than as CPTP maps between operator algebras (see also Ref. [Coe11]). We shall use the symmetric monoidal categories exclusively as they pertain to the latter depiction.

1.1 A Mathematical Model for Physical Theories

The precise definition we choose for a theory is a *symmetric monoidal category* in which the *monoidal unit* is *terminal*. These words might intimidate certain readers, but I should like to emphasise that the concept is intuitively simple and ubiquitous, in fact intelligible to anyone who has interacted with the real world. Readers who prefer concrete rather than abstract mind-sets will not lose much by fixing ‘theory’ to mean either **CIT** (classical information theory) or **QIT** (quantum information theory).

Category theory was created in the 1940s by Samuel Eilenberg and Saunders Mac Lane ([EML45]). It was developed for applications in algebraic topology, but it soon grew wildly beyond this scope and is nowadays considered a universal language for many mathematical ideas and constructions (the original go-to reference is [ML13]; Ref. [Awo10] offers a modern and less overwhelming treatment).

In recent times, category theory has been successfully implemented also in areas outside of pure mathematics, of which Ref. [BS10] provides a very readable overview. One of these areas is the study of foundational physics, where it was realised that (symmetric monoidal) categories can be used to model physical theories.

What is a category? Formally, it is a type of mathematical structure. Poetically, it is the incarnation of the abstract idea of ‘serial composition’. More precisely, a *category* \mathbf{C} comprises a collection of *objects*, $\mathcal{X}, \mathcal{Y}, \mathcal{Z}, \dots$, and a collection of *morphisms* between these objects, T, S, R, \dots . For example, the objects could be sets and the morphisms from the set \mathcal{X} to the set \mathcal{Y} could be functions from \mathcal{X} to \mathcal{Y} . Alternatively, the objects could be groups and the morphisms from the group \mathcal{X} to the group \mathcal{Y} could be group homomorphisms from \mathcal{X} to \mathcal{Y} . We write $T : \mathcal{X} \rightarrow \mathcal{Y}$ to signify that T is a morphism from \mathcal{X} to \mathcal{Y} . A category \mathbf{C} is defined by its collection² of objects and morphisms, and by a notion of composition of morphisms: Given $T : \mathcal{X} \rightarrow \mathcal{Y}$ and $S : \mathcal{Y} \rightarrow \mathcal{Z}$, there is a morphism $S \circ T : \mathcal{X} \rightarrow \mathcal{Z}$, called the (serial) composition of T with S . In the cases of functions or group homomorphisms this composition is ordinary functional composition, but in general \circ is just an abstract binary operation. It is subject to the associativity requirement $(R \circ S) \circ T = R \circ (S \circ T)$, and it is moreover required that to every object \mathcal{Z} is associated a morphism $\text{id}_{\mathcal{Z}} : \mathcal{Z} \rightarrow \mathcal{Z}$, called identity, such that $T \circ \text{id}_{\mathcal{X}} = T = \text{id}_{\mathcal{Y}} \circ T$ for any morphism $T : \mathcal{X} \rightarrow \mathcal{Y}$. And that is it.

Like other mathematical structures, a category may be equipped with additional architecture, making it a more refined object. One such additional architecture is that of a *symmetric monoidal* structure, which adds one further mode of composition. Whereas the composition inherent in every category is *serial*, one morphism following another, a (symmetric) monoidal structure facilitates a notion of *parallel* composition. The category of sets and functions is an example of a category which allows such a structure – as discussed in the introduction, the parallel composition of $f_1 : X_1 \rightarrow Y_1$ and $f_2 : X_2 \rightarrow Y_2$ is the function $f_1 \times f_2 : X_1 \times X_2 \rightarrow Y_1 \times Y_2$, given by $(f_1 \times f_2)(x_1, x_2) = (f_1(x_1), f_2(x_2))$. Clearly, the category of groups sustains a similar construction. An example of a category with no ob-

²Readers who are used to defining a mathematical structure as a *set* equipped with certain operations or additional material, and who know something about the axioms of set theory, might worry that it is dangerous to define a mathematical structure whose underlying universe it too big to be a set (e.g. the collection of all sets, which is a proper class). There are however at least two formal escape routes: One is to use a different frame of axioms in which the notion of a (*proper*) class has formal meaning, for example the set theory of von Neumann-Bernays-Gödel. Another is to consider proper classes as entities which exist in the metalanguage, namely as predicates in first-order logic which intuitively define the class. See e.g. Ref. [ML13] for further discussions.

vious monoidal structure is that of Boolean algebras and homomorphisms between Boolean algebras.

In general, a *monoidal structure* on a category \mathbf{C} is a binary operation, \llbracket , additional to the existing serial composition. This operation has two components: For any two objects \mathcal{X} and \mathcal{Y} in \mathbf{C} , it defines an object $\mathcal{X} \llbracket \mathcal{Y}$ in \mathbf{C} , the composite of \mathcal{X} and \mathcal{Y} ; and for any two morphisms $T_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1$ and $T_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$ in \mathbf{C} it defines a parallel composite $T_1 \llbracket T_2 : \mathcal{X}_1 \llbracket \mathcal{X}_2 \rightarrow \mathcal{Y}_1 \llbracket \mathcal{Y}_2$ in \mathbf{C} . Just as the serial composition in a bare category, the parallel composition is subject to an associativity requirement,³ and also required to interplay sensibly with the serial composition (for example, one requires $(S_1 \circ T_1) \llbracket (S_2 \circ T_2) = (S_1 \llbracket S_2) \circ (T_1 \llbracket T_2)$). Moreover, one requires the existence of a special system, $\mathbf{1}$, which acts as a unit for the \llbracket -operation on systems: $\mathcal{X} \llbracket \mathbf{1} = \mathcal{X} = \mathbf{1} \llbracket \mathcal{X}$. In the example of sets and functions, we can declare as unit object any set $\{*\}$ which contains a single element.⁴ In the example of groups and group homomorphisms we can take for the object $\mathbf{1}$ any trivial group.

Finally, the word *symmetric* in ‘symmetric monoidal’ refers to the fact that the parallel composition is required to be in a certain sense symmetric. This is again well illustrated in the example of sets and functions: Whereas for functions $f : X \rightarrow X$ and $g : X \rightarrow X$ the serial compositions $g \circ f$ and $f \circ g$ are generally very different, there is a sense in which, for $f_1 : X_1 \rightarrow Y_1$ and $f_2 : X_2 \rightarrow Y_2$, the functions $f_1 \times f_2$ and $f_2 \times f_1$ are just two different ways of looking at the same function. Similarly, the sets $X \times Y$ and $Y \times X$ can be easily identified.

Readers who are interested in an accessible and more detailed introduction to symmetric monoidal categories may consult one of many well-written expositions, e.g. Refs. [BS10, Coe11]. Readers in want of more knowledge about mere categories may consult Ref. [Awo10].

For the sake of completeness – and out of respect for mathematically minded readers – I find it appropriate to reproduce below a precise definition of symmetric monoidal categories. On the other hand, any reader who feels comfortable with an intuitive impression of symmetric monoidal categories (or is creative enough to assemble a definition based on the many examples in Section 1.2), is invited to save eye power by skipping Definition 1.1.2 and going now directly to Section 1.1.A.

Remark 1.1.1. (For those Intending to Read the Definition.)

To avoid as much formalism as possible, only the definition of an especially simple kind of symmetric monoidal category is stated, namely a so-called *strict* one. This is essentially means doing away with the issues surrounding the precise relation between $(\mathcal{X} \llbracket \mathcal{Y}) \llbracket \mathcal{Z}$ and $\mathcal{X} \llbracket (\mathcal{Y} \llbracket \mathcal{Z})$, and between $\mathcal{X} \llbracket \mathbf{1}$, \mathcal{X} and $\mathbf{1} \llbracket \mathcal{X}$. This approach is standard, and it is justified by Mac Lane’s ‘Strictification Theorem’ ([ML13], Chapter XI, Section 3), according to which any monoidal category is equivalent to a strict monoidal category (via a pair of ‘strong monoidal functors’).

In practice, this means that we need never formally consider non-strict categories. Thus, we adopt the commonly held attitude that for **Theorems** and **Definitions** we assume categories to be strict, whereas for **Examples** we have no hesitations about exposing non-strict categories.

³Although the associativity requirement is cumbersome to state precisely. The reason is that in most cases of interest, the composition is only ‘almost’ associative; for example, given sets X, Y and Z , the two sets $(X \times Y) \times Z$ and $X \times (Y \times Z)$ are easily identifiable but not formally identical.

⁴Again, we do not strictly have the equalities $X \times \{*\} = X = \{*\} \times X$, but the three sets are naturally identifiable.

The strictification theorem does not go as far as to down the similar problem of the relationship between $\mathcal{X} \parallel \mathcal{Y}$ and $\mathcal{Y} \parallel \mathcal{X}$. Rather, this relationship must be formalised in terms of *swapping morphisms* $\sigma_{\mathcal{X},\mathcal{Y}} : \mathcal{X} \parallel \mathcal{Y} \rightarrow \mathcal{Y} \parallel \mathcal{X}$. Unfortunately, the conditions imposed on these morphisms take up a part of Definition 1.1.2 which in size is disproportional to their significance. \spadesuit

Definition 1.1.2. (Symmetric (Strict) Monoidal Categories ([ML13]).)

A *symmetric (strict) monoidal category* is a quadruple $(\mathbf{C}, \mathbf{1}, \parallel, \sigma)$ comprised as follows:

1. \mathbf{C} is a category.
2. $\mathbf{1}$ is an object in \mathbf{C} .
- 3a. \parallel is a map which maps pairs of objects $(\mathcal{X}, \mathcal{Y})$ to objects $\mathcal{X} \parallel \mathcal{Y}$, and pairs of morphisms $(T_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1, T_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2)$ to morphisms $T_1 \parallel T_2 : \mathcal{X}_1 \parallel \mathcal{X}_2 \rightarrow \mathcal{Y}_1 \parallel \mathcal{Y}_2$.
- 3b. \parallel is associative on objects and morphisms, with $\mathbf{1}$ and $\text{id}_{\mathbf{1}}$ as units, in the sense that
 - for any objects \mathcal{X}, \mathcal{Y} and \mathcal{Z} in \mathbf{C} ,

$$(\mathcal{X} \parallel \mathcal{Y}) \parallel \mathcal{Z} = \mathcal{X} \parallel (\mathcal{Y} \parallel \mathcal{Z}), \quad (1.1)$$

$$\mathcal{X} \parallel \mathbf{1} = \mathcal{X} = \mathbf{1} \parallel \mathcal{X}; \quad (1.2)$$

- for any morphisms $T : \mathcal{X}_1 \rightarrow \mathcal{X}_2, S : \mathcal{Y}_1 \rightarrow \mathcal{Y}_2$ and $R : \mathcal{Z}_1 \rightarrow \mathcal{Z}_2$,

$$(T \parallel S) \parallel R = T \parallel (S \parallel R), \quad (1.3)$$

$$T \parallel \text{id}_{\mathbf{1}} = T = \text{id}_{\mathbf{1}} \parallel T. \quad (1.4)$$

- 3c. \parallel is *functorial*, meaning that

- for any objects \mathcal{X} and \mathcal{Y} ,

$$\text{id}_{\mathcal{X} \parallel \mathcal{Y}} = \text{id}_{\mathcal{X}} \parallel \text{id}_{\mathcal{Y}}; \quad (1.5)$$

- for any morphisms $T_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1, T_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$ and $S_1 : \mathcal{Y}_1 \rightarrow \mathcal{Z}_1, S_2 : \mathcal{Y}_2 \rightarrow \mathcal{Z}_2$,

$$(S_1 \circ T_1) \parallel (S_2 \circ T_2) = (S_1 \parallel S_2) \circ (T_1 \parallel T_2). \quad (1.6)$$

4. σ is a collection of morphisms in \mathbf{C} called *swappings*, one morphism $\sigma_{\mathcal{X},\mathcal{Y}} : \mathcal{X} \parallel \mathcal{Y} \rightarrow \mathcal{Y} \parallel \mathcal{X}$ for each pair $(\mathcal{X}, \mathcal{Y})$ of objects in \mathbf{C} . They are subject to the conditions

$$\sigma_{\mathcal{X},\mathbf{1}} = \sigma_{\mathbf{1},\mathcal{X}} = \text{id}_{\mathcal{X}}, \quad (1.7)$$

$$\sigma_{\mathcal{Y},\mathcal{X}} \circ \sigma_{\mathcal{X},\mathcal{Y}} = \text{id}_{\mathcal{X} \parallel \mathcal{Y}}, \quad (1.8)$$

$$(\sigma_{\mathcal{Z},\mathcal{X}} \parallel \text{id}_{\mathcal{Y}}) \circ \sigma_{(\mathcal{X} \parallel \mathcal{Y}),\mathcal{Z}} = \text{id}_{\mathcal{X}} \parallel \sigma_{\mathcal{Y},\mathcal{Z}} \quad (1.9)$$

for all objects \mathcal{X}, \mathcal{Y} and \mathcal{Z} , the latter of which is to say that if in $\mathcal{X} \parallel \mathcal{Y} \parallel \mathcal{Z}$ we swap $\mathcal{X} \parallel \mathcal{Y}$ for \mathcal{Z} and then \mathcal{Z} for \mathcal{X} , this altogether amounts to swapping \mathcal{Y} for \mathcal{Z} .

Moreover, it must hold for any morphisms $T_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1$ and $T_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$, that

$$\sigma_{\mathcal{Y}_1,\mathcal{Y}_2} \circ (T_1 \parallel T_2) = (T_2 \parallel T_1) \circ \sigma_{\mathcal{X}_1,\mathcal{X}_2}. \quad (1.10)$$

■

As is customary in all mathematical disciplines, we shall often abbreviate the quadruple $(\mathbf{C}, \mathbf{1}, \parallel, \sigma)$ simply by ‘ \mathbf{C} ’, letting its family members be implicit as they are usually clear from the context.

1.1.A Definition of Theories – CIT and QIT

Having defined symmetric monoidal categories, the scariest part of the section, if not the entire chapter, is over.

We have already touched on two examples of symmetric monoidal categories, namely sets with functions and groups with group homomorphisms. Now that formalities are in order, let us baptise them properly:

Example 1.1.3. (Sets^{*}.)

The category **Sets^{*}** has non-empty⁵ sets X, Y, Z, \dots as objects and functions $f : X \rightarrow Y$ as morphisms from X to Y . Its serial composition is given by ordinary functional composition. The symmetric monoidal structure on **Sets^{*}** is given by $X \amalg Y := X \times Y$, the cartesian product of sets, and $f_1 \amalg f_2 := f_1 \times f_2$ for functions $f_1 : X_1 \rightarrow Y_1$, $f_2 : X_2 \rightarrow Y_2$, where $(f_1 \times f_2)(x_1, x_2) = (f_1(x_1), f_2(x_2))$ for $(x_1, x_2) \in X_1 \times X_2$. As unit object $\mathbf{1}$ we may take any set with a single element, say $\mathbf{1} := \{\emptyset\}$ for concreteness. It is tedious but straightforward to verify the conditions of Definition 1.1.2 (ignoring the formal difference between $(X \times Y) \times Z$ and $X \times (Y \times Z)$, and between $X \times \mathbf{1}$, X and $\mathbf{1} \times X$). The swapping functions $\sigma_{X,Y} : X \times Y \rightarrow Y \times X$ are given by $\sigma_{X,Y}(x, y) = (y, x)$. \blacklozenge

Example 1.1.4. (Groups.)

In **Groups** the objects are groups G, H, K, \dots , and the morphisms from G to H are group homomorphisms $\varphi : G \rightarrow H$. Group homomorphisms are functions after all, and so we can define the compositions exactly as in **Sets^{*}**: Serial composition of morphisms $\varphi : G \rightarrow H$ and $\psi : H \rightarrow K$ is given by the ordinary functional composition $\psi \circ \varphi$, and parallel composition of $\varphi_1 : G_1 \rightarrow H_1$ with $\varphi_2 : G_2 \rightarrow H_2$ by $\varphi_1 \times \varphi_2 : G_1 \times G_2 \rightarrow H_1 \times H_2$, where $K_1 \times K_2$ denotes the product group of K_1 and K_2 . As unit object $\mathbf{1}$ we take some fixed trivial group. \blacklozenge

Here is another example of a symmetric monoidal category, indeed historically one of the main inspirations for the very definition of the concept:

Example 1.1.5. (Vect_k.)

Let k be a field (e.g. $k = \mathbb{R}$ or $k = \mathbb{C}$), and let **Vect_k** denote the category whose objects are (finite-dimensional) vector spaces over k , and whose morphisms are k -linear maps between these spaces, with functional composition as composition. The *tensor product* \otimes defines a notion of parallel composition, making **Vect_k** a symmetric monoidal category: The composition of objects V and W is the tensor product $V \otimes W$, and the parallel composition of the linear maps $A_1 : V_1 \rightarrow W_1$ and $A_2 : V_2 \rightarrow W_2$ is the linear map $A_1 \otimes A_2 : V_1 \otimes V_2 \rightarrow W_1 \otimes W_2$ determined by $(A_1 \otimes A_2)(v_1 \otimes v_2) = A_1(v_1) \otimes A_2(v_2)$. For unit object $\mathbf{1}$ we take the 1-dimensional vector space k . \blacklozenge

Whereas we want to include **Sets^{*}** and **Groups** in our club of theories, **Vect_k** is for our purposes an imposter. (Some authors are more accommodating; see Remark 1.1.7.) The reason is that **Sets^{*}** and **Groups** admit a well-defined notion of *marginalisation*, whereas **Vect_k** does not:

We will come to think of objects as ‘systems’ in a theory, and morphisms as ‘transformations’ between those systems. The composite $\mathcal{X} \amalg \mathcal{Y}$ will represent the junction of two systems into one, and the system $\mathbf{1}$ will represent the ‘trivial system’, i.e. the system corresponding to ‘nothing’. As such, transformations $\text{tr}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbf{1}$ correspond to various ways of *discarding*,

⁵The reason for restricting to non-empty sets will become clear later (Section 1.3.C); the problem is essentially that the empty set is very destructive in its parallel composition with other sets.

or *trashing* the system \mathcal{X} , and by extension the transformations $\text{tr}_{\mathcal{X}} \parallel \text{id}_{\mathcal{Y}} : \mathcal{X} \parallel \mathcal{Y} \rightarrow \mathbf{1} \parallel \mathcal{Y} = \mathcal{Y}$ correspond to ways of discarding only the system \mathcal{X} from the composite system $\mathcal{X} \parallel \mathcal{Y}$. This is precisely the process known as marginalisation, and for it to exist and be unique, we need $\text{tr}_{\mathcal{X}}$ to exist and be unique.

Both **Sets**^{*} and **Groups** have this property; there is a unique function from X to $\mathbf{1} = \{\emptyset\}$ for any set X , and there is a unique homomorphism from G to $\mathbf{1}$ for any group G . In contrast, there are many k -linear maps from a vector space V to the vector space k (these are precisely the functionals on V).

In general, an object in a category is called *terminal* if every object admits a unique morphism to it. We thus arrive at the following definition of a theory:

Definition 1.1.6. (Theories.)

A *theory* is a symmetric (strict) monoidal category Θ , such that the monoidal unit object $\mathbf{1}$ is terminal.

The following terminology is employed:

- Objects in Θ are called *systems*, and we denote the class of all systems in Θ by Sys_{Θ} .
- Given $\mathcal{X}, \mathcal{Y} \in \text{Sys}_{\Theta}$, the system $\mathcal{X} \parallel \mathcal{Y}$ is called the *composite of \mathcal{X} and \mathcal{Y}* .
- Given $\mathcal{X}, \mathcal{Y} \in \text{Sys}_{\Theta}$, the morphisms $T : \mathcal{X} \rightarrow \mathcal{Y}$ in Θ are called *transformations from \mathcal{X} to \mathcal{Y}* , and the class of all such transformations is denoted by $\text{Trans}_{\Theta}(\mathcal{X}, \mathcal{Y})$. The class of all transformations in Θ is denoted by Trans_{Θ} .
- Given transformations $T : \mathcal{X} \rightarrow \mathcal{Y}$ and $S : \mathcal{Y} \rightarrow \mathcal{Z}$ in Θ , the transformation $S \circ T : \mathcal{X} \rightarrow \mathcal{Z}$ is called the *serial composition of T and S* .
- Given transformations $T_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1$ and $T_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$ in Θ , the transformation $T_1 \parallel T_2 : \mathcal{X}_1 \parallel \mathcal{X}_2 \rightarrow \mathcal{Y}_1 \parallel \mathcal{Y}_2$ is called the *parallel composition of T_1 and T_2* .
- The system $\mathbf{1}$ is called *the trivial system*. Given a system $\mathcal{X} \in \text{Sys}_{\Theta}$, the unique transformation from \mathcal{X} to $\mathbf{1}$ is denoted $\text{tr}_{\mathcal{X}}$ and called *the trash of \mathcal{X}* .

■

Remark 1.1.7. (All Theories are Causal.)

In some line of work, theories are simply identified with symmetric monoidal categories, and the stricter concept defined by Definition 1.1.6 is then referred to as *causal* theory, since the terminality assumption on $\mathbf{1}$ can be interpreted as an impossibility of signalling from the future to the past ([CDP10, Coe14]). Deviating from this terminology is justified on the grounds that we shall have no interest in ‘theories’ which are not causal, and that we will already use the work ‘causal’ to a near-excessive degree in other connections.

✂

Remark 1.1.8. (Typesetting.)

Generically, we typeset theories with boldface letters ($\Theta, \mathbf{Sets}, \dots$), systems of a theory with calligraphic Latin letters ($\mathcal{X}, \mathcal{Y}, \mathcal{Z}, \dots$), and transformations of a theory with ordinary capital Latin letters (T, S, R, \dots).

In specific theories (such as **Sets**^{*}, **Groups**, and **CIT** and **QIT** defined below) we may deviate from these conventions if tradition prescribes. More systematic deviations will be mentioned as introduced (for example, the special transformations to be called ‘states’ will be generically typeset with lower-case Latin letters s, t, \dots).

✂

The next two examples of theories will be our most important:

Example 1.1.9. (Classical Information Theory, **CIT**.)

The systems of **CIT** are finite, non-empty sets X, Y, Z, \dots . They compose under the cartesian product \times , as in the theory **Sets**^{*}, and the trivial system is some distinguished one-element set, say $\mathbf{1} := \{0\}$. A transformation $T : X \rightarrow Y$ can be thought of as a ‘probabilistic function’. Formally, it is a *Markov kernel*, i.e. a collection $T = (t_x)_{x \in X}$ of probability distributions on Y ; a genuine (‘deterministic’) function $f : X \rightarrow Y$ corresponds indeed to the collection $(\delta_{f(x)})_{x \in X}$, where δ_y denotes the degenerate distribution in the point $y \in Y$. (Observe in particular that a transformation from $\mathbf{1}$ to Y is simply a probability distribution on Y .) The serial and parallel composition of transformations is best described by appealing to intuition: If we think of a transformation $T : X \rightarrow Y$ as encoding a process by which on input $x \in X$ a random $y \in Y$ is produced according to the distribution t_x , then the serial composition of $T : X \rightarrow Y$ with $S : Y \rightarrow Z$ corresponds – unsurprisingly – to the process resulting from applying S after T , assuming independence of the randomness in S and T . Likewise, the parallel composition $T_1 \parallel T_2$ corresponds to the process of drawing simultaneously and independently outputs y_1 and y_2 based on the inputs x_1 and x_2 , by means of T_1 and T_2 respectively. Formally,

$$(t_{x_1}^1)_{x_1 \in X_1} \parallel (t_{x_2}^2)_{x_2 \in X_2} = (t_{x_1}^1 \otimes t_{x_2}^2)_{(x_1, x_2) \in X_1 \times X_2} \quad (1.11)$$

(with $t_{x_1}^1 \otimes t_{x_2}^2$ denoting the product distribution on $Y_1 \times Y_2$ of distributions $t_{x_1}^1$ on Y_1 and $t_{x_2}^2$ on Y_2), and

$$(s_y)_{y \in Y} \circ (t_x)_{x \in X} = (u_x)_{x \in X}, \quad u_x = \sum_{y \in Y} t_x(y) s_y. \quad (1.12)$$

As usual, it is tedious but easy to verify that **CIT** satisfies the formal conditions of Definition 1.1.6. Note that the identity transformation on X is $(\delta_x)_{x \in X}$, and that the trash $\text{tr}_X : X \rightarrow \mathbf{1}$ is the X -indexed collection of degenerate distributions on the one-element set $\mathbf{1}$.

Given a channel $T = (t_x)_{x \in X}$ we will often write $T(x)$ or $T(\delta_x)$ for the probability distribution t_x , when there is no risk of confusion. \blacklozenge

Example 1.1.10. (Quantum Information Theory, **QIT**.)

The systems of **QIT** are finite-dimensional, non-zero Hilbert spaces $\mathcal{H}, \mathcal{K}, \mathcal{L}, \dots$ over \mathbb{C} . They compose parallelly under the tensor product \otimes , and the unit object is some distinguished 1-dimensional space, say $\mathbf{1} := \mathbb{C}$. A morphism $\Lambda : \mathcal{H} \rightarrow \mathcal{K}$ is a so-called *quantum channel* from \mathcal{H} to \mathcal{K} , meaning a CPTP (completely positive trace-preserving) linear map from $\text{End}(\mathcal{H})$ to $\text{End}(\mathcal{K})$, where $\text{End}(\mathcal{L})$ denotes the space of linear operators on \mathcal{L} . The fact that a morphism $\Lambda : \mathcal{H} \rightarrow \mathcal{K}$ is a map $\Lambda : \text{End}(\mathcal{H}) \rightarrow \text{End}(\mathcal{K})$ is notationally odd-looking, but should not imply confusion in relevant instances. The serial composition of morphisms $\Lambda : \mathcal{H} \rightarrow \mathcal{K}$ and $\Phi : \mathcal{K} \rightarrow \mathcal{L}$ is given by the functional composition $\Phi \circ \Lambda$, and the parallel composition of morphisms $\Lambda_1 : \mathcal{H}_1 \rightarrow \mathcal{K}_1$ and $\Lambda_2 : \mathcal{H}_2 \rightarrow \mathcal{K}_2$ is given by the tensor product map $\Lambda_1 \otimes \Lambda_2$, determined by $(\Lambda_1 \otimes \Lambda_2)(A_1 \otimes A_2) = \Lambda_1(A_1) \otimes \Lambda_2(A_2)$ for $A_1 \in \text{End}(\mathcal{H}_1)$, $A_2 \in \text{End}(\mathcal{H}_2)$. (Observe here the isomorphisms $\text{End}(\mathcal{H}_1 \otimes \mathcal{H}_2) \cong \text{End}(\mathcal{H}_1) \otimes \text{End}(\mathcal{H}_2)$ and $\text{End}(\mathcal{K}_1 \otimes \mathcal{K}_2) \cong \text{End}(\mathcal{K}_1) \otimes \text{End}(\mathcal{K}_2)$.) The identity $\text{id}_{\mathcal{H}}$ is the identity map on $\text{End}(\mathcal{H})$ and the trash $\text{tr}_{\mathcal{H}}$ is the trace $\text{End}(\mathcal{H}) \rightarrow \text{End}(\mathbb{C}) \cong \mathbb{C}$ (this is the only trace-preserving linear map to $\text{End}(\mathbb{C})$). \blacklozenge

We will typeset, as is customary, the transformations in **QIT** with Greek letters, but often typeset linear operators which are used to define the transformations (e.g. isometries or Kraus operators) with Latin letters. Though this clashes somewhat unfortunately with the general convention of using Latin letters for the transformations themselves, this should not cause confusion.

Remark 1.1.11. (On Terminology and Notation.)

The category **CIT** is often denoted in the literature by ‘**FinStoch**’ (see e.g. Ref. [Fri20]; the terminology seems to have originated in Ref. [BF14]), and referred to as the *category of finite sets and stochastic maps between them*. The name and notation chosen in this thesis is meant to reflect the emphasis on the category as a theory of classical information, and to reinforce the physical and formal relationship with the theory **QIT**. (The category **QIT** is rarely named in the literature.) ❖

At this point, we shall not entertain any physical interpretations whatsoever of the theory **QIT**. (Realistically, most readers of these sentences will know of such an interpretation anyway.) Suffice it to say that in the same way that **CIT** models the probabilistic processing of classical information with which most of us are at least intuitively familiar, it has been determined, ultimately empirically, that **QIT** is the correct model for the processing of quantum information ([NC02]), and hence for information processing as it really is in our world (to the best of our understanding).

That said, three points about **QIT** deserve mentioning before we continue with the investigation of general theories:

Remark 1.1.12. (A Formal Relationship between **CIT** and **QIT**.)

First, it should be pointed out that the theory **CIT** is naturally contained in the theory **QIT**, by means of the following construction: To a system X in **CIT** we associate the Hilbert space $\Gamma(X) := \mathbb{C}^X$ (with its canonical inner product) in **QIT**, and to a morphism $T = (t_x)_{x \in X} : X \rightarrow Y$ in **CIT** we associate the CPTP map $\Gamma(T) : \Gamma(X) \rightarrow \Gamma(Y)$ given by $\Gamma(T)(A) = \sum_{x \in X, y \in Y} t_x(y) \langle x | A | x \rangle |y\rangle\langle y|$ for $A \in \text{End}(\mathbb{C}^X)$, such that in particular $\Gamma(T)(|x\rangle\langle x|) = \sum_{y \in Y} t_x(y) |y\rangle\langle y|$. The precise sense in which this construction gives a representation of **CIT** in **QIT** can be summarised by the observation that Γ has all the properties of a (strong) monoidal functor ([ML13]) from **CIT** to **QIT**, except that Γ does not preserve identities, i.e. $\Gamma(\text{id}_X) \neq \text{id}_{\Gamma(X)}$. Explicitly, $\Gamma(S \circ T) = \Gamma(S) \circ \Gamma(T)$, $\Gamma(X \times Y) \cong \Gamma(X) \otimes \Gamma(Y)$ and $\Gamma(T_1 \times T_2) \cong \Gamma(T_1) \otimes \Gamma(T_2)$.⁶ Moreover, Γ is injective. It is interesting to observe (and seems to have been not noted before), that the failure of Γ to preserve identities cannot be fixed by a redefinition – there simply does not exist a strong monoidal functor from **CIT** to **QIT** which is injective. In succinct terms, the reason for this, which we will come to appreciate in Chapter 2, is that the classical channel $\text{Cop} : X \rightarrow X \times X$ which deterministically copies the input (i.e. corresponds to the function $x \mapsto (x, x)$), would under such a functor Φ have to map to a quantum channel $\Phi(\text{Cop}) : \Phi(X) \rightarrow \Phi(X) \otimes \Phi(X)$, both of whose marginals are $\text{id}_{\Phi(X)}$, which is by the No Broadcasting Theorem ([BCF⁺96] – see also Corollary 2.5.7) impossible, except when $X \cong \mathbf{1}$. ❖

Remark 1.1.13. (Notions of Classicality.)

Secondly, the above-mentioned embedding Γ of **CIT** in **QIT** allows us to define notions of

⁶An earlier version of this chapter contained an entire section proposing this notion of *homomorphism* between theories, defined as maps satisfying all properties of (strong) monoidal functors, except preservation of identities. An injective such homomorphism, like Γ , can be interpreted as a *generalisation* from one theory to another.

‘classicality’ in quantum information theory. This was already reviewed in the preliminary section of the thesis, but can now be rephrased in terms of the embedding Γ .

Specifically, we call a system in **QIT** *classical* if it is of the form $\Gamma(X) = \mathbb{C}^X$ for some set X (any system in **QIT** is *isomorphic* to \mathbb{C}^X for some set X ; being classical is thus a matter of being equipped with a preferred basis). As already observed, the embedded identities $\Gamma(\text{id}_X)$ are distinct from the actual identities $\text{id}_{\Gamma(X)}$ on \mathbb{C}^X , indeed $\Gamma(\text{id}_X)$ is the *decoherence channel on X* , namely the quantum channel $\Delta_X : \mathbb{C}^X \rightarrow \mathbb{C}^X$ given by $\Delta_X(A) = \sum_{x \in X} \langle x | A | x \rangle |x\rangle\langle x|$. The fact that $\Delta_X = \Gamma(\text{id}_X)$ is distinct from $\text{id}_{\Gamma(X)}$ means that the e.g. the condition $\Gamma(\text{id}_Y) \circ M = M$ for a quantum channel $M : \mathcal{H} \rightarrow \mathbb{C}^Y$ is non-trivial, and it makes sense in this case to say that M *has classical outcomes*. The more well-known term for this concept is that M is a *measurement*. There is similarly a notion of a channel $\Lambda : \mathbb{C}^X \rightarrow \mathcal{K}$ having *classical inputs*, which corresponds to being an ensemble $(\varrho_x)_{x \in X}$ of channels $\mathbf{1} \rightarrow \mathcal{K}$ (so-called *states*). As mentioned in the preliminary section, if Λ is a quantum channel between composite systems it makes sense to speak of Λ being classical on some of the input or output systems, thus giving rise most generally to ensembles of so-called *quantum instruments*.

It is easy to verify that the embeddings of classical channels, $\Gamma(T)$, are precisely those channels $\Lambda : \mathbb{C}^X \rightarrow \mathbb{C}^Y$ which satisfy $\Delta_Y \circ \Lambda \circ \Delta_X = \Lambda$. However, certain compositions of transformations in **QIT** could result in a transformation interpretable in **CIT** even though its constituents are not. For example, for serially composable transformations Λ_1 and Λ_2 in **QIT**, the transformation $\Lambda_2 \circ \Lambda_1$ might be classical (i.e. of the form $\Gamma(T)$) even though neither Λ_1 nor Λ_2 is classical. Abstractly, this is what allows us in the first place to make statements about quantum information which are classically intelligible. By considering a more intricate combination of transformations, one can exhibit a total transformation which is classical, although no classical choice of the constituents will reproduce this transformation. This statement is essentially the famous observation of John Bell ([Bel64]) described in the introduction of the thesis, and the consequences are profound: **QIT** *is a larger theory than CIT, and this can be classically observed*.

✦

Remark 1.1.14. (On the Definition of **QIT**.)

Finally, it is very important to appreciate the fact that though **QIT** was defined in Example 1.1.10 in terms of Hilbert spaces and linear operators, there could very well be ways of defining **QIT** (up to a suitable notion isomorphism) *without making any reference to such entities*. As demonstrated by ground-breaking works such as [Har01] and later [CDP11], there are indeed completely different definitions of **QIT**, which in their formulation are much less obscure, cast in an operational language. In fact, the quest for simple and natural definitions of **QIT** is an ongoing area of research (see Ref. [CS16] for a review), and in many ways the question that motivated the present thesis – that of finding an operational definition of quantum self-testing – is very much inspired by this line of thought.

✦

1.1.B States, Isomorphisms and Reversibles

We now proceed to discuss special kinds of transformations in a given theory: *States*, *isomorphisms* and *reversibles*. The naming of states is uncontroversial, whereas there is no consensus on the naming of the latter two (see also Remark 1.1.21).

First, however, let us prove the following helpful result about trashes, which is used over and over throughout the thesis:

Lemma 1.1.15. (*Properties of Trashes.*)

The following holds of the trashes in a theory Θ :

1. For any transformation $T : \mathcal{X} \rightarrow \mathcal{Y}$ in Θ , $\text{tr}_{\mathcal{Y}} \circ T = \text{tr}_{\mathcal{X}}$.
2. For any systems \mathcal{X}, \mathcal{Y} in Θ , $\text{tr}_{\mathcal{X} \parallel \mathcal{Y}} = \text{tr}_{\mathcal{X}} \parallel \text{tr}_{\mathcal{Y}}$.
3. $\text{tr}_{\mathbf{1}} = \text{id}_{\mathbf{1}}$.

Proof. Given $T : \mathcal{X} \rightarrow \mathcal{Y}$, the transformation $\text{tr}_{\mathcal{Y}} \circ T$ is *some* transformation from \mathcal{X} to $\mathbf{1}$. Since there is only one, namely $\text{tr}_{\mathcal{X}}$, we must have $\text{tr}_{\mathcal{Y}} \circ T = \text{tr}_{\mathcal{X}}$, proving the first property. The second and third are proved similarly. \square

A theory Θ has a single system which is distinguished, namely the trivial system $\mathbf{1}$ – in fact, $\mathbf{1}$ may be the only system in Θ . The system $\mathbf{1}$ represents ‘nothing’, and whereas we have imposed that transformations to $\mathbf{1}$ are not very diverse (there is only one from each system), transformations from the system $\mathbf{1}$ are very colourful. They physically correspond to producing something from nothing:

Definition 1.1.16. (States in Θ .)

Given a system \mathcal{X} in Θ , the transformations from $\mathbf{1}$ to \mathcal{X} are called *states on \mathcal{X}* . The class of all states on \mathcal{X} is denoted by $\text{St}(\mathcal{X})$. \blacksquare

We generically denote states by lowercase Latin letters s, t, \dots

Example 1.1.17. The states in a theory are usually rather easy to understand:

- In **Sets**^{*}, a state on X is a map $s : \mathbf{1} \rightarrow X$, or, what is equivalent, an element $x \in X$. As such, some systems have many states and some have few. (If we consider instead of **Sets**^{*} the theory **Sets** in which also the empty set is included, then some systems – namely \emptyset – has *no* states.)
- In **Groups**, a state on G is a homomorphism $\sigma : \mathbf{1} \rightarrow G$. There is only one such, since it must map to the identity element in G .
- In **CIT**, a state on X is a classical channel $\mathbf{1} \rightarrow X$, or, what is equivalent, a probability distribution p on X .
- In **QIT**, a state on \mathcal{H} is a completely positive trace-preserving linear map from $\text{End}(\mathbb{C})$ to $\text{End}(\mathcal{H})$. Since $\text{End}(\mathbb{C}) \cong \mathbb{C}$ as vector spaces, such a map is characterised by a unique element $\varrho \in \text{End}(\mathcal{H})$, namely the image of $1 \in \mathbb{C}$, and by the CPTP property this element must be positive and of unit trace. Conversely, for any positive $\varrho \in \text{End}(\mathcal{H})$ with $\text{tr}(\varrho) = 1$ the map $\mathbb{C} \ni a \mapsto a\varrho \in \text{End}(\mathcal{H})$ is CPTP and hence defines a state on \mathcal{H} . In conclusion, we may identify the set of states on \mathcal{H} with the set of *density matrices on \mathcal{H}* , namely

$$\mathcal{D}(\mathcal{H}) := \{\varrho \in \text{End}(\mathcal{H}) \mid \varrho \geq 0, \text{tr}(\varrho) = 1\}. \quad (1.13)$$

- In any theory Θ , there is by assumption a unique map from $\mathbf{1}$ to $\mathbf{1}$, namely $\text{tr}_{\mathbf{1}} (= \text{id}_{\mathbf{1}})$, and this is a state on $\mathbf{1}$. It is possible that no other systems in Θ have states. \blacklozenge

Example 1.1.18. (States from States.)

The following hold in any theory:

- Since $\mathbf{1} \parallel \mathbf{1} = \mathbf{1}$, the parallel composition of states $s \in \text{St}(\mathcal{X})$ and $t \in \text{St}(\mathcal{Y})$ is a new state $s \parallel t : \mathbf{1} \rightarrow \mathcal{X} \parallel \mathcal{Y}$, on the system $\mathcal{X} \parallel \mathcal{Y}$.
Succinctly, we have a map $\text{St}(\mathcal{X}) \times \text{St}(\mathcal{Y}) \rightarrow \text{St}(\mathcal{X} \parallel \mathcal{Y})$ given by $(s, t) \mapsto s \parallel t$.
- If $s \in \text{St}(\mathcal{X})$ and $T : \mathcal{X} \rightarrow \mathcal{Y}$ is a transformation, then the serial composition $T \circ s : \mathbf{1} \rightarrow \mathcal{Y}$ is a state on \mathcal{Y} .
In other words, any transformation $T : \mathcal{X} \rightarrow \mathcal{Y}$ induces a map $\text{St}(\mathcal{X}) \rightarrow \text{St}(\mathcal{Y})$ given by $s \mapsto T \circ s$.

◆

Warning 1.1.19. (Transforming States.)

From common use of the words, it is tempting to think that a transformation is determined by its action on states, i.e. that the abstract morphism $T : \mathcal{X} \rightarrow \mathcal{Y}$ in Θ can be identified with the set-theoretic function $s \mapsto T \circ s$ from $\text{St}(\mathcal{X})$ to $\text{St}(\mathcal{Y})$. This identification can be done without harm in **CIT** and **QIT** (though in **QIT** it relies on the non-trivial fact that any operator on \mathcal{H} is a linear combination of density matrices). It is also unproblematic in the theory **Sets**^{*}, where in fact it is subtle to even distinguish the original from the impersonator, since $\text{St}(X) \cong X$. The principle that transformations be determined by their action on states is physically sound (what sense is there in two transformations being distinct if this cannot be observed on states?), and it is enforced in much existing literature, essentially by identifying transformations which act identically on states (e.g. as in Ref. [CDP10]). Nevertheless, it is not a principle we shall commit to. In fact, it may very well fail in more mathematical examples of theories, such as **Groups**, where each system has only one state, whence any two transformations from G to H act identically on states. ✖

Leaving states for now, we proceed to two other important types of transformations:

Definition 1.1.20. (Reversibles and Isomorphisms in Θ .)

- A transformation $R : \mathcal{X} \rightarrow \mathcal{Y}$ is called *reversible* if it has a left-inverse, i.e. if there exists a transformation $R^- : \mathcal{Y} \rightarrow \mathcal{X}$ such that $R^- \circ R = \text{id}_{\mathcal{X}}$.
- A transformation $\alpha : \mathcal{X} \rightarrow \mathcal{Y}$ is called an *isomorphism* if it has a two-sided inverse, i.e. if there exists a transformation $\beta : \mathcal{Y} \rightarrow \mathcal{X}$ such that $\beta \circ \alpha = \text{id}_{\mathcal{X}}$ and $\alpha \circ \beta = \text{id}_{\mathcal{Y}}$.

■

Remark 1.1.21. (Terminology.)

Refs. [CDP10, CDP11] use the term ‘reversible’ differently than we – in fact, it is used even within these references in two different ways, cf. Definitions 13 and 46 in [CDP10]. One of the uses (Def. 13) is for what we call ‘isomorphisms’. The terminology chosen in the Definition 1.1.20 is based on the grounds that (1) the term ‘isomorphism’ has been established in the mathematical literature on categories for more than half a century; (2) we will need some word for transformations with left-inverses;⁷ (3) the Latin verb *revertere* means to *turn back*. ✖

⁷In category theory proper, there is in fact a term for morphisms with left-inverses, namely *split monomorphisms*; however, I render this type of vocabulary slightly too esoteric for our purposes.

Example 1.1.22. (Isomorphisms and Reversibles.)

- In **Sets**^{*}, a map $f : X \rightarrow Y$ is an isomorphism if and only if it is bijective. It is reversible if and only if it is injective.
- In **Groups**, the isomorphisms are precisely the group isomorphisms. Any reversible transformation is an injective homomorphism (though some of these are not reversible).
- In **QIT**, every transformation $\Lambda : \mathcal{H} \rightarrow \mathcal{K}$ which is a unitary conjugation, $A \mapsto UAU^*$, is an isomorphism, with two-sided inverse given by $B \mapsto U^*BU$. It is not obvious that any isomorphism in **QIT** must take this form, but we shall later see a swift argument for this (Corollary 2.5.9) using the machinery of Chapter 2.

We shall similarly be able to characterise all reversible transformations in **QIT** (Corollary 2.5.12). For now, let us observe that if $\Sigma : \mathcal{H} \rightarrow \mathcal{K}$ is an isometric conjugation, $A \mapsto SAS^*$, then it is reversible. It is tempting to provide as left-inverse the completely positive map $B \mapsto S^*BS$, but it is not trace-preserving since $\text{tr}(S^*AS) = \text{tr}(SS^*A)$ and $SS^* \not\leq \mathbb{1}$ (unless S is unitary). Instead, pick an isometry $S' : \mathcal{K} \rightarrow \mathcal{H} \otimes \mathcal{E}$ such that $S'S = \mathbb{1}_{\mathcal{H}} \otimes |\psi\rangle$, where $|\psi\rangle$ is a unit vector in some space \mathcal{E} ; then the map $B \mapsto [\text{id}_{\text{End}(\mathcal{H})} \otimes \text{tr}_{\mathcal{E}}](S'BS'^*)$ is a completely positive trace-preserving left-inverse.

◆

It is well-known (and easy to show) that if $\alpha : \mathcal{X} \rightarrow \mathcal{Y}$ is an isomorphism, its two-sided inverse β is unique. We denote it, as is customary, by α^{-1} . Clearly, $\alpha^{-1} : \mathcal{Y} \rightarrow \mathcal{X}$ is an isomorphism with inverse $(\alpha^{-1})^{-1} = \alpha$. If $\alpha : \mathcal{X} \rightarrow \mathcal{Y}$ and $\beta : \mathcal{Y} \rightarrow \mathcal{Z}$ are isomorphisms, then their serial composition is an isomorphism too, with $(\alpha \circ \beta)^{-1} = \beta^{-1} \circ \alpha^{-1}$. Similarly, if $\alpha_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1$ and $\alpha_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$ are isomorphisms, their parallel composition is an isomorphism with $(\alpha_1 \parallel \alpha_2)^{-1} = \alpha_1^{-1} \parallel \alpha_2^{-1}$.

It is customary to call systems \mathcal{X}, \mathcal{Y} *isomorphic* if there exists an isomorphism between them. For example, systems \mathcal{H} and \mathcal{K} are isomorphic in **QIT** precisely when they have the same dimension, and systems X and Y are isomorphic in **CIT** precisely when they have the same number of elements. By the previous comments, isomorphism of systems is always an equivalence relation, well-behaved under parallel composition of systems.

Evidently, every isomorphism in a theory Θ is reversible. As Example 1.1.22 shows, there may however easily be reversible transformations in Θ which are not isomorphisms. It should also be observed that, contrary to isomorphisms, reversibles need not have unique inverses:

Remark 1.1.23. (Non-Uniqueness of Left-Inverses.)

In **Sets**^{*}, the injective inclusion map $\iota : \{0, 1\} \rightarrow \{0, 1, \dots, 8, 9\}$ given by $\iota(x) = x$ has as left-inverse any function $g : \{0, 1, \dots, 8, 9\} \rightarrow \{0, 1\}$ for which $g(0) = 0$ and $g(1) = 1$. As such, the values $g(y)$ for $y \in \{2, 3, \dots, 8, 9\}$ can be set arbitrarily as 0 or 1, so there are as many such maps g as there are subsets of the set $\{2, 3, \dots, 8, 9\}$, namely $2^8 = 256$. ✦

Let me end this section by posing a curious problem.

The reversible transformations in a theory Θ facilitate a rudimentary notion of dimension. More precisely, let us define the *dimensional ordering*, \preceq , on Sys_{Θ} by declaring that $\mathcal{X} \preceq \mathcal{Y}$ exactly if there exists a reversible transformation $R : \mathcal{X} \rightarrow \mathcal{Y}$. It is easy to see that \preceq is a pre-order, i.e. a reflexive transitive relation. (We also have $\mathcal{X}_1 \otimes \mathcal{X}_2 \preceq \mathcal{Y}_1 \otimes \mathcal{Y}_2$ if $\mathcal{X}_1 \preceq \mathcal{Y}_1$ and $\mathcal{X}_2 \preceq \mathcal{Y}_2$, and we have $\mathbf{1} \preceq \mathcal{X}$ if $\text{St}(\mathcal{X}) \neq \emptyset$.)

In **Sets*** and **CIT**, the dimensional ordering reproduces the ordering in terms of cardinality of sets ($X \preceq Y \Leftrightarrow |X| \leq |Y|$), and in **QIT** it yields the usual ordering according to dimension ($\mathcal{H} \preceq \mathcal{K} \Leftrightarrow \dim \mathcal{H} \leq \dim \mathcal{K}$). In general, however, the ordering \preceq is simply an abstract relation, not necessarily related to cardinal numbers or enjoying properties we usually expect.

One such property would be the *Cantor-Schröder-Bernstein property* (named by analogy with the Cantor-Schröder-Bernstein theorem for sets [Sch]), i.e. the principle that if $\mathcal{X} \preceq \mathcal{Y}$ and $\mathcal{Y} \preceq \mathcal{X}$ then \mathcal{X} and \mathcal{Y} are isomorphic. By Example 1.2.7, however, it is demonstrated that this property is not always satisfied.

Another expectable property would be

- *Linearity*: For all systems \mathcal{X}, \mathcal{Y} , either $\mathcal{X} \preceq \mathcal{Y}$ or $\mathcal{Y} \preceq \mathcal{X}$,

or the stronger property

- *Well-Foundedness*: For any set of systems $(\mathcal{X}_i)_{i \in I}$, there is some $i_0 \in I$ such that $\mathcal{X}_{i_0} \preceq \mathcal{X}_i$ for all $i \in I$.

Readers acquainted with the theory of ordinals and cardinals may appreciate that well-foundedness is more or less equivalent to the possibility of representing the levels of \preceq using cardinals, as we can for **Sets*** (where the equivalence class of X is represented by the cardinality $|X|$) and for **QIT** (where the equivalence class of \mathcal{H} is represented by the cardinal $\dim \mathcal{H} \in \mathbb{N}$).⁸

I do not know if it is possible to construct an example in which linearity or well-foundedness fail. In fact, I do not even know the answer to the following question:

Open Problem 1.1.24. *Is every pre-order the dimensional ordering of some theory Θ ?*

1.2 A Reservoir of Examples

Examples in mathematics serve roughly two purposes, one soft and one hard.

The soft purpose is that *examples help humans fix ideas*. For instance, a person seeing the definition of a topological space for the first time may not immediately grasp what this concept is about. Exhibiting concrete examples will help that person form a view of what a topological space is; some of these examples will fit smoothly in line with those that motivated the definition in the first place, whereas others may be surprising.

The hard purpose is that *examples uncover formal interdependencies of properties*. Some examples of topological spaces will show that certain properties cannot be derived – or are undecidable – from the axioms defining a topological space. For instance, one cannot prove that a topological space has infinitely many open sets (for this is not always true, as e.g. the trivial topology exemplifies). In a similar vein, that a set is closed does not imply that its image under a continuous map is closed (as exemplified by the map $\mathbb{R} \ni x \mapsto \frac{1}{1+x^2} \in \mathbb{R}$, which maps \mathbb{R} to $(0, 1]$).

In this section we go through a lot of examples of theories. In fact, the presented catalogue might be one the largest list of theories (in the sense of Definition 1.1.6) existing

⁸If such a representation is possible, well-foundedness of \preceq follows from well-foundedness of cardinals. If on the other hand \preceq is well-founded (and the collection of its equivalence classes is small enough to be a set) then by a standard result ([Kun80]) that set is order-isomorphic to an ordinal γ . As such, the levels of \preceq are representable as an ordering among ordinals $\beta \in \gamma$, and from this we can obtain an ordering in terms of cardinals by means of the cardinal counting map $\beta \mapsto \kappa_\beta$ from ordinals to cardinals.

at one place in the literature. Some of these examples serve the soft purpose, but most will the hard as well. Some of the examples will be so mathematical that a physicist would not call them ‘theories’ (like **Groups**, cf. Warning 1.1.19). Still, such examples may easily serve the hard purpose, demonstrating that defined concepts do not always behave as expected.

1.2.A Variations of CIT and QIT

We have already seen our two most important examples of theories, **CIT** (Example 1.1.9) and **QIT** (Example 1.1.10). Historically, much of quantum theory was conceived in a setting of infinite-dimensional Hilbert spaces, and emphasis on the finite-dimensional setting was only recently articulated ([NC02, CS16]).

There is indeed version of quantum information theory which allows (separable) infinite-dimensional Hilbert spaces as systems ([Att14, Wol19]):

Example 1.2.1. (QIT[∞].)

In the infinitary version of quantum information theory, **QIT[∞]**, systems are separable Hilbert spaces $\mathcal{H}, \mathcal{K}, \mathcal{L}, \dots$, and a transformation from \mathcal{H} to \mathcal{K} is a CPTP map $\Lambda : B_1(\mathcal{H}) \rightarrow B_1(\mathcal{K})$, as outlined in the preliminary section of the thesis. (Recall that $B_1(\mathcal{L})$ denotes the Banach space of trace-class operators on the Hilbert space \mathcal{L} .) When \mathcal{H} and \mathcal{K} are finite-dimensional, this notion of transformation restricts to that from **QIT**. The composite of systems in **QIT[∞]** is again given by the tensor product (whose construction now requires a metric completion of the algebraic tensor product), and the trivial system is $\mathbf{1} = \mathbb{C}$. Serial and parallel compositions of transformations are given, as for **QIT**, by the functional composition and tensor product of linear maps, respectively. States on a system \mathcal{X} are (by the same argument used for **QIT**) in natural bijective correspondence with linear operators $\varrho \in B_1(\mathcal{X})$, which are positive and of unit trace. \blacklozenge

The theory **QIT[∞]** does display features which **QIT** does not, but they also have a lot in common and for our purposes their differences are not profound. (One of the differences is that in **QIT[∞]** there exists a system \mathcal{H} , namely any space of infinite dimension, into which all systems admit a reversible transformation; in **QIT** there exists no such system.)

There is also a version of **CIT** which goes beyond finite sets:

Example 1.2.2. (Stoch.)

The theory **Stoch** (following the notation of Ref. [Fri20]) has measurable spaces $\mathcal{X} = (X, \mathbb{E})$ for systems, and transformations from $\mathcal{X} = (X, \mathbb{E})$ to $\mathcal{Y} = (Y, \mathbb{K})$ are *Markov kernels from \mathcal{X} to \mathcal{Y}* , that is, X -indexed collections of probability measures on \mathcal{Y} , $(\lambda_x)_{x \in X}$, for which the function $x \mapsto \lambda_x(B)$ is measurable for any fixed $B \in \mathbb{E}$. The composite of systems $\mathcal{X} = (X, \mathbb{E})$ and $\mathcal{Y} = (Y, \mathbb{K})$ is the measurable spaces $(X \times Y, \mathbb{E} \otimes \mathbb{K})$, where $\mathbb{E} \otimes \mathbb{K}$ is the product σ -algebra, and the parallel composition of transformations is defined in the obvious way by forming product measures (see any introductory book on measure theory, e.g. [Han09, Sch17, Fol99]). The serial composition is also defined in a rather obvious fashion, by *integrating* the Markov kernels, though the construction is somewhat shrouded in measurability technicalities. (For more details, see Ref. [Fri20], or the curiously historical lecture notes [Law62] which apparently constitute the first categorical presentation of Markov kernels.)

Every measurable injection $(X, \mathbb{E}) \rightarrow (Y, \mathbb{K})$ is a reversible transformation in **Stoch**, its isomorphisms are precisely the Borel-isomorphisms, and states on (X, \mathbb{E}) correspond to probability measures on (X, \mathbb{E}) . \blacklozenge

It is intuitively clear that the theory **QIT**[∞] extends **QIT**, and that **Stoch** extends **CIT** (we will be more precise about this in Section 1.2.D). However, it would seem that **Stoch** is in a sense ‘too big’ an extension of **CIT** when compared to the extension **QIT**[∞] of **QIT**. Indeed, as mentioned earlier, the theory **CIT** embeds into **QIT**, but there is no obvious sense in which **Stoch** embeds into **QIT**[∞], since there is no canonical way of associating a Hilbert space to a measurable space (X, \mathbb{E}) . One could speculate that by taking measure spaces (X, \mathbb{E}, μ) in place of measurable spaces (X, \mathbb{E}) as systems (and by requiring a sufficient compatibility of the transformations, e.g. having appropriate densities w.r.t. the ground measures), one could massage **Stoch** into a theory for which such an embedding would be possible, but I do not know of any such construction.

So far, all of our theories have had plenty of states, with the exception of **Groups**. It is possible to devise an ‘information’ theory which also does not have many states:

Example 1.2.3. (Oblivious Information Theory.)

In **QIT**, every system \mathcal{H} has a unique *invariant state*, $\tau_{\mathcal{H}}$, described by the density matrix $\frac{1}{\dim \mathcal{H}} \mathbb{1}_{\mathcal{H}}$, and defined by the property that $\alpha \circ \tau_{\mathcal{H}} = \tau_{\mathcal{H}}$ for all isomorphisms (unitary conjugations) $\alpha : \mathcal{H} \rightarrow \mathcal{H}$. The states $\tau_{\mathcal{H}}$ are also called *fully mixed*, and they can be interpreted as representing complete obliviousness about the system \mathcal{H} . Let us define *oblivious quantum information theory*, **OblQIT**, as the theory in which this obliviousness is preserved: The systems of **OblQIT** are finite-dimensional non-zero Hilbert spaces, and the transformations from \mathcal{H} to \mathcal{K} are the CPTP maps $\Lambda : \text{End}(\mathcal{H}) \rightarrow \text{End}(\mathcal{K})$ for which $\Lambda(\tau_{\mathcal{H}}) = \tau_{\mathcal{K}}$. The composition of systems and transformations is given as in **QIT**, and the trivial system is again $\mathbf{1} := \mathbb{C}$. Importantly, the unique state $\text{tr}_{\mathbf{1}} = \text{id}_{\mathbf{1}}$ on the system $\mathbf{1}$ is invariant, and therefore all trashes $\text{tr}_{\mathcal{H}}$ in **QIT** are valid transformations in **OblQIT**, so $\mathbf{1}$ is indeed terminal. Note also that any state $\sigma : \mathbf{1} \rightarrow \mathcal{H}$ in **OblQIT** must be invariant, i.e. the system \mathcal{H} admits a unique state, namely $\tau_{\mathcal{H}}$.

Of course, there is nothing quantum about this idea: Classical information theory, **CIT**, also has unique invariant states (namely the uniform distributions) and by restricting to transformations that map uniform distributions to uniform distributions we similarly obtain an oblivious classical information theory, **OblCIT**. An interesting exercise for the reader is to verify that, contrary to the what is the case in **CIT**, every reversible transformation in **OblCIT** is an isomorphism. \blacklozenge

Rather than restricting the class of transformations, we can enlarge it:

Example 1.2.4. (Negative Information Theory.)

Let us define **NCIT** (‘negative classical information theory’) as the theory whose systems and composition of systems is the same as in **CIT**, but whose transformations from X to Y are collections $t = (t_x)_{x \in X}$ of ‘not necessarily positive probability distributions on Y ’, that is, of functions $t_x : Y \rightarrow \mathbb{R}$ such that

$$\sum_{y \in Y} t_x(y) = 1. \tag{1.14}$$

The potential usefulness of negative probabilities has been discussed in e.g. Ref. [Fey87]. We define serial and parallel composition in **NCIT** by the same equations as for **CIT**. Importantly, the system $\mathbf{1} = \{0\}$ remains terminal in **NCIT** because δ_0 is the only \mathbb{R} -valued function on $\{0\}$ satisfying the normalisation (1.14).

Certainly, the theory **NCIT** seems quite distinct from **CIT**, e.g. in admitting on the system $\{0, 1\}$ the ‘unbounded’ set of states $\{t\delta_0 + (1-t)\delta_1 \mid t \in \mathbb{R}\}$, and in admitting

‘convex combinations’ such as $\delta_0 = \frac{1}{2}(3\delta_0 - 2\delta_1) + \frac{1}{2}(-\delta_0 + 2\delta_1)$. But to prove a categorical distinction from the theory **CIT** we cannot refer to notions of convexity or boundedness; we have to point out a distinction visible in terms of the serial and parallel composition. To this end, consider in **NCIT** the transformations $\alpha^w = (\alpha_j^w)_{j \in \{0,1\}} : \{0,1\} \rightarrow \{0,1\}$ given for $w \in \mathbb{R}$ by

$$\alpha_0^w = (1-w)\delta_0 + w\delta_1, \quad \alpha_1^w = -w\delta_0 + (1+w)\delta_1. \quad (1.15)$$

It is easily verified that $\alpha^0 = \text{id}_{\{0,1\}}$ and that $\alpha^u \circ \alpha^v = \alpha^{u+v}$ for all $u, v \in \mathbb{R}$, so the map $w \mapsto \alpha_w$ is an injective group homomorphism from the additive group of reals to the group of automorphisms of $\{0,1\}$ (i.e. isomorphisms $\{0,1\} \rightarrow \{0,1\}$) in **NCIT**. In particular, any automorphism α^w with $w \neq 0$ has infinite multiplicative group order; in **CIT**, on the other hand, every automorphism of every system is a bijection on a finite set and hence has finite multiplicative order. \blacklozenge

1.2.B Cartesian Theories

The theory **Sets**^{*} (Example 1.1.3) has the following feature: Any transformation into a composite system, say $f : X \rightarrow Y_1 \times Y_2$, is given by two components, $f = (f_1, f_2)$ with $f_1 : X \rightarrow Y_1$ and $f_2 : X \rightarrow Y_2$. These two component functions are the *marginals* of f , and they completely determine f . Similarly, a transformation $\varphi : G \rightarrow H_1 \times H_2$ in **Groups** is determined by its marginals.

On the other hand, transformations in **CIT** (or **QIT**) are not determined by marginals; for examples, the states $k := \frac{1}{2}\delta_0 \otimes \delta_0 + \frac{1}{2}\delta_1 \otimes \delta_1$ and $p := (\frac{1}{2}\delta_0 + \frac{1}{2}\delta_1) \otimes (\frac{1}{2}\delta_0 + \frac{1}{2}\delta_1)$ on the system $\{0,1\} \times \{0,1\}$ have identical marginals, but they are not the same; the state k represents two copies of a uniformly random bit, whereas p represents two independent uniformly random bits.

Let us be slightly more precise:

Definition 1.2.5. (Marginal-Determined.)

Let \mathcal{Y}_1 and \mathcal{Y}_2 be systems in a theory Θ . Let $\pi_1 : \mathcal{Y}_1 \parallel \mathcal{Y}_2 \rightarrow \mathcal{Y}_1$ and $\pi_2 : \mathcal{Y}_1 \parallel \mathcal{Y}_2 \rightarrow \mathcal{Y}_2$ denote the *factor projections*, $\text{id}_{\mathcal{Y}_1} \parallel \text{tr}_{\mathcal{Y}_2}$ and $\text{tr}_{\mathcal{Y}_1} \parallel \text{id}_{\mathcal{Y}_2}$, respectively. We say that the system pair $(\mathcal{Y}_1, \mathcal{Y}_2)$ is *marginal-determined*, if for any two transformations $T_1 : \mathcal{X} \rightarrow \mathcal{Y}_1$ and $T_2 : \mathcal{X} \rightarrow \mathcal{Y}_2$, there exists a unique transformation $T : \mathcal{X} \rightarrow \mathcal{Y}_1 \parallel \mathcal{Y}_2$ such that $\pi_1 \circ T = T_1$ and $\pi_2 \circ T = T_2$. \blacksquare

Definition 1.2.6. (Cartesian Theories.)

A theory Θ is called *cartesian* if any pair of systems in Θ is marginal-determined. \blacksquare

The introductory lines serve to illustrate that **Sets**^{*} and **Groups** are cartesian theories, whereas **CIT** and **QIT** are not. Readers acquainted with category theory will realise that cartesian theories are precisely *categories with finite products* ([ML13, Awo10]), sometimes referred to as *cartesian categories* (hence the name). This realisation immediately gives a true bombardment of theory examples, including all sorts of categories whose morphisms are functions on ‘structured sets’ which admit a notion of product:

- **Top**^{*}, in which the systems are non-empty topological spaces and the transformations continuous maps;
- **Rings**, in which the systems are algebraic rings and the transformations are ring homomorphisms;

- \mathbf{Man}^∞ , in which the systems are differentiable manifolds and the transformations are smooth maps;
- \vdots


In all these cases, the serial composition of transformations is given by ordinary functional composition and the parallel composition is given by means of the products that these categories facilitate (products of topological spaces, of rings, of manifolds, ...).

Another example of this is the theory **Graphs** of graphs and homomorphisms. However, in the same way that the empty set in **Sets** introduces some pathological features making the theory **Sets*** nicer in the end, so do graphs with un-looped vertices cause problems in **Graphs** (see Section 1.3.C). Thus, we shall consider instead the theory **Graphs***, whose objects are (non-empty) graphs in which every vertex has a loop:


Example 1.2.7. (Graphs*.)

The category **Graphs*** has non-empty looped graphs for objects and graph homomorphisms for morphisms, with functional composition as serial composition. Any graph with one vertex and its loop is terminal; we fix one and call it **1**. One can consider several ‘products’ of graphs G and H (see [Gra]), but only one kind will make **Graphs*** a cartesian theory, and this is the so-called *direct product*, $G \times H$. It has as vertices pairs of vertices in G and H , and it has an edge between (u_1, v_1) and (u_2, v_2) precisely if u_1 and u_2 are adjacent in G and v_1 and v_2 in H .⁹ Transformations compose parallelly as one would expect. The isomorphisms in **Graphs** are precisely the graph isomorphisms, and a state on the graph G is simply a vertex in G (for this it matters that every vertex has a loop).

The collection of reversible transformations allows us to exhibit a feature which we have not encountered earlier. It can be phrased as the failure in **Graphs*** of the Cantor-Schröder-Bernstein property discussed in Section 1.1.B, and explicitly it is the following: There are systems G, H with the property that we can find reversible transformations $g : G \rightarrow H$ and $h : H \rightarrow G$, though there is no isomorphism between G and H . It is an easy exercise to see that this can only be the case if G and H are infinite, but we may actually take them rather simple. Indeed, let G be the graph



and H the graph



both extending infinitely to the right in a periodic fashion. For reversible transformation $g : G \rightarrow H$ we take the unique injective homomorphism that maps the leftmost ‘head’,

⁹There is a distinct product, $G \square H$, which quite confusingly is called the *cartesian product* of G and H ; it can also be seen as a *funny tensor product* of categories ([Gra]). The operation \square makes **Graphs*** a symmetric monoidal category in a different way. Thus, **Graphs*** can be considered a theory in two distinct ways: Either by equipping it with the direct product \times , or by equipping it with the product \square .

$a_0b_0c_0$, in G to the leftmost head, $a'_1b'_1c'_1$, in H . (The image of G under g is all of H except for the edges r'_0 and d'_1 and their common vertex.) Now, g has as left-inverse the map $g^- : H \rightarrow G$ which acts as the inverse of g on its image, and by collapsing r'_0, d'_1 and their common vertex to the left ‘ear’ of the leftmost head in G . A reversible $h : H \rightarrow G$ is constructed completely analogously, by mapping the structure $d'_1a'_1b'_1c'_1$ to $d_1a_1b_1c_1$. The graphs G and H are not isomorphic, however, since H has a loop (namely l_1) whose vertex only has one other edge. \blacklozenge

Another example of a cartesian theory whose transformations are functions on ‘structured sets’ can be obtained from linear algebra. We saw in Example 1.1.5 that \mathbf{Vect}_k , the category of vector spaces over k and k -linear maps between them, can be augmented to a symmetric monoidal category by means of the tensor product, \otimes . We also discussed, however, that this does not constitute a theory in the sense of Definition 1.1.6, since the monoidal unit k fails to be terminal. It turns out that we can augment \mathbf{Vect}_k with another notion of parallel composition that does make it into a theory:

Example 1.2.8. (\mathbf{Vect}_k .)

Let k be a field and consider on the category \mathbf{Vect}_k the symmetric monoidal structure defined by the *direct sum*, \oplus ; the composite of systems V and W is $V \oplus W$, and the parallel composition of $A_1 : V_1 \rightarrow W_1$ with $A_2 : V_2 \rightarrow W_2$ is $A_1 \oplus A_2 : V_1 \oplus V_2 \rightarrow W_1 \oplus W_2$. For terminal object and \oplus -unit we fix a zero-dimensional space, $\mathbf{1} := 0$. The trashes $\text{tr}_V : V \rightarrow 0$ must then be the zero-maps, and the factor projections $\pi_j : V_1 \oplus V_2 \rightarrow V_j$ are consequently the ordinary projections P_j onto the subspaces V_j . The theory is cartesian because any linear map $A : W \rightarrow V_1 \oplus V_2$ is specified by the projected maps $A_j := P_j A$ and because any such pair of maps A_1, A_2 defines a map to $V_1 \oplus V_2$ by $x \mapsto A_1 x \oplus A_2 x$. \blacklozenge

We end with an example demonstrating that the systems in a cartesian theory need not be ‘sets with structure’:

Example 1.2.9. (The Interval Theory.)

Consider the real unit interval $[0, 1]$, and define a theory Θ as follows:

Systems of Θ are numbers $x, y, z, \dots \in [0, 1]$. For any $x, y \in [0, 1]$ there is at most one transformation from x to y , and there is one precisely if $x \geq y$. (It does not matter what the transformation actually *is*, but for concreteness we may choose it to be the pair (x, y) .) Serial composition of transformations can be defined uniquely, since $x \geq y$ and $y \geq z$ implies $x \geq z$, and each system has an identity transformation since $x \geq x$. The associative and symmetric composition of systems in Θ is given by $x \parallel y := \max\{x, y\}$, and the system $0 \in [0, 1]$ is a unit for this operation which is terminal in Θ since $x \geq 0$ always. Parallel composition of transformations is also uniquely defined, by the observation that $x_1 \geq y_1, x_2 \geq y_2$ imply $\max\{x_1, x_2\} \geq \max\{y_1, y_2\}$. Finally, the theory is cartesian since $z \geq x_1$ and $z \geq x_2$ if and only if $z \geq \max\{x_1, x_2\}$. \blacklozenge

1.2.C Thin Theories

We just saw in Example 1.2.9 that there exist theories with at most one transformation from one system to another. We can make an entire example class out of such theories, and they turn out to have a fairly graspable characterisation. As actual physical theories tend to have many transformations, this class of theories is mostly interesting for purely mathematical purposes, or for finding counterexamples.

In category theory, categories with at most one morphism from one object to another are called *thin* ([[Thi](#)]), so we adopt the same terminology:

Definition 1.2.10. (Thin Theories.)

A theory Θ is called *thin* if for any systems $\mathcal{X}, \mathcal{Y} \in \text{Sys}_\Theta$ there is at most one transformation from \mathcal{X} to \mathcal{Y} . ■

Obviously, the composition of transformations – serial and parallel – in a thin theory is unexciting. Really, it is the composition of its systems which is interesting.

In any theory Θ , the composition \parallel on systems gives Sys_Θ the structure of a *monoid*, $(\text{Sys}_\Theta, \parallel, \mathbf{1})$, with unit object $\mathbf{1}$. And in a thin theory, the transformation structure can be compactly summarised as follows: Let us define on Sys_Θ a relation \succeq by $\mathcal{X} \succeq \mathcal{Y}$ if and only if there is a transformation from \mathcal{X} to \mathcal{Y} . By the axioms of identities and serial composition, this relation is reflexive and transitive, i.e. it is a pre-order on Sys_Θ . Terminality of $\mathbf{1}$ means $\mathcal{X} \succeq \mathbf{1}$ for all $\mathcal{X} \in \text{Sys}_\Theta$, and the parallel composition of transformations means that $\mathcal{X}_1 \parallel \mathcal{X}_2 \succeq \mathcal{Y}_1 \parallel \mathcal{Y}_2$ when $\mathcal{X}_1 \succeq \mathcal{Y}_1$ and $\mathcal{X}_2 \succeq \mathcal{Y}_2$. Finally, the symmetry condition of the theory implies that $\mathcal{X} \parallel \mathcal{Y}$ and $\mathcal{Y} \parallel \mathcal{X}$ are equivalent under the pre-order, i.e. $\mathcal{X} \parallel \mathcal{Y} \succeq \mathcal{Y} \parallel \mathcal{X}$ and $\mathcal{Y} \parallel \mathcal{X} \succeq \mathcal{X} \parallel \mathcal{Y}$.

Conversely, it is easy to see that any monoid equipped with a pre-order subject to these conditions defines a thin theory.

In summary, we have proved the following:

A thin (strict) theory is the same thing as a pre-ordered quasi-commutative monoid,

where by this horrifying sequence of words I mean a quadruple $(M, \star, 1, \succeq)$, such that

- $(M, \star, 1)$ is a monoid;
- \succeq is a pre-order on M ;
- $x \succeq 1$ for all $x \in M$;
- $x_1 \succeq y_1, x_2 \succeq y_2 \Rightarrow x_1 \star x_2 \succeq y_1 \star y_2$ for all $x_1, x_2, y_1, y_2 \in M$;
- $x \star y \simeq y \star x$ for all $x, y \in M$.¹⁰

The identity transformations in Θ are the relationships $x \succeq x$, and the trashes are the relationships $x \succeq 1$.

Remark 1.2.11. (Strictness.)

If we drop the strictness assumption for the theory, the strict associativity of the operation \star and strict unitality of the element 1 are replaced by \simeq -equivalences in the monoid $(M, \star, 1)$ (e.g. $x \star 1 \simeq x \simeq 1 \star x$ rather than $x \star 1 = x = 1 \star x$). As such, the characterisation of thin theories by means of pre-ordered monoid-like structures is not contingent on the strictness. ✂

Some of the simplest examples of thin theories are partially ordered commutative monoids:

Example 1.2.12. (One Ordering, two Compositions.)

The non-negative integers $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ form a commutative monoid with unit 0 , both when equipped with addition, $(n, m) \mapsto n + m$, and when equipped with the max-function, $(n, m) \mapsto \max\{n, m\}$. The usual ordering \geq on \mathbb{N}_0 satisfies the required compatibility conditions with these binary operations, so we have two thin theories $(\mathbb{N}_0, +, 0, \geq)$ and $(\mathbb{N}_0, \max, 0, \geq)$. ◆

¹⁰Here, $z \simeq w$ means $z \succeq w$ and $w \succeq z$.

Example 1.2.13. (One Composition, two Orderings.)

The natural numbers, $\mathbb{N} = \{1, 2, 3, \dots\}$ form a commutative monoid with unit 1 when equipped with multiplication, $(n, m) \mapsto n \cdot m$. The usual ordering \geq on \mathbb{N}_0 satisfies the required compatibility conditions, and so does the *divisibility ordering*, \geq_{div} , according to which $n \geq_{\text{div}} m$ precisely if m divides n . Thus we have two thin theories, $(\mathbb{N}, \cdot, 1, \geq)$ and $(\mathbb{N}, \cdot, 1, \geq_{\text{div}})$. \blacklozenge

Example 1.2.14. (Powersets.)

For any set A , the powerset $\mathcal{P}(A)$ is a commutative monoid with unit \emptyset when equipped with the union-operation \cup . Set-theoretic inclusion \supseteq is a compatible partial order, so we have a thin theory $(\mathcal{P}(A), \cup, \emptyset, \supseteq)$. \blacklozenge

There are also examples of thin theories in which the pre-order does not meet the condition to be a partial order. The condition $x \simeq y \Rightarrow x = y$ can break down violently, or subtly:

Example 1.2.15. (Any Monoid is a Thin Theory.)

Let $(M, \star, 1)$ be any monoid. By putting the trivial relation \succeq on M , which renders $x \succeq y$ for all $x, y \in M$, we see that $(M, \star, 1)$ is augmented to a thin theory. (This is even independent of whether or not the monoid is commutative.) \blacklozenge

Example 1.2.16. (Logic.)

Let P_1, P_2, P_3, \dots be infinitely many symbols, and consider the set of all *well-formed formulas* that can be generated from these symbols along with the logical connectives $\wedge, \vee, \neg, \top, \perp, \rightarrow, \leftrightarrow$ and the parentheses $)$ and $($. For example, $\neg(P_1 \wedge P_3)$ is a well-formed formula, whereas $) \rightarrow P_2 \neg$ is not. The set of well-formed formulas is pre-ordered by the relation $\psi \succeq \phi$ asserting that formula ϕ is provable from formula ψ (using some standard inference system, see e.g. Ref. [End01]), and by equipping it with the binary operation that maps the pair ψ_1, ψ_2 to $(\psi_1 \wedge \psi_2)$, it becomes a thin theory **Logic**, with unit \top (though it is neither strictly commutative nor strictly associative). \blacklozenge

Remark 1.2.17. (Relation to Resource Theories [CFS16].)

In Ref. [CFS16] the authors propose ordered monoids $(M, \star, 1, \succeq)$ as a model of *resource convertibility*; more precisely, a thin theory in our language is in their language a *waste-free theory of resource convertibility*. The elements $x, y, z, \dots \in M$ correspond to resources in some universe, and the pre-order assignment $x \succeq y$ reflects that resource x can be converted into resource y at no cost. The binary operation \star simply represents the junction of two resources, and $1 \in M$ represents a void resource. (The adjective ‘waste-free’ refers to the relations $x \succeq 1$.) The scope of this interpretation is large, ranging from economy to chemistry. It also includes the theory of resource convertibility associated to (*bipartite*) *quantum entanglement*. \blacklozenge

States, isomorphisms and reversible transformations are strange notions in thin theories:

Remark 1.2.18. (States, Isomorphisms and Reversibles in Thin Theories.)

An isomorphism from x to y in a thin theory Θ is a pair of relationships $x \succeq y$ and $y \succeq x$, whose two compositions yield the identities on x and y , respectively. However, since the serial composition of the transformation $x \succeq y$ with $y \succeq x$ must yield some transformation from x to x and since there is by assumption only one, namely the identity $x \succeq x$, any pair of relationships between two systems x, y witnesses an isomorphism. In short, for x and y to be isomorphic is precisely the condition $x \simeq y$. The same argument implies that any reversible transformation is necessarily an isomorphism.

Not all systems in a thin theory have states. Actually, for x to have a state means precisely that $1 \succeq x$, which is to say that $x \simeq 1$. \spadesuit

In Example 1.2.9 we saw a cartesian thin theory. We end this subsection by classifying those thin theories which are also cartesian:

Proposition 1.2.19. (*Thin Cartesian Theories.*)

A thin theory described by the pre-ordered quasi-commutative monoid $(M, \star, 1, \succeq)$ is cartesian if and only if for every $x, y \in M$, the element $x \star y \in M$ is a least upper bound for x and y in the pre-order (M, \succeq) .

Proof. The thin theory described by $(M, \star, 1, \succeq)$ is cartesian if and only if for any $x, y, z \in M$ it holds that $(z \succeq x) \wedge (z \succeq y) \Leftrightarrow z \succeq x \star y$. This is precisely to say that $x \star y$ is an upper bound for x and y which is least among all upper bounds. \square

It follows that for instance the theory **Logic** from Example 1.2.16 is also cartesian.

1.2.D Sub-Theories

Some of the above examples of theories were *nested*, one inside the other. For example, the oblivious version of quantum information, **OblQIT**, was a *sub-theory* of **QIT**, in the sense that all of its systems and transformations, along with their serial and parallel composition, came from **QIT**. In a similar way, **QIT** itself was a sub-theory of **QIT** $^\infty$, as was **CIT** of **Stoch**.

In fact, **CIT** is also a sub-theory of **QIT**, or, more precisely, **QIT** has a sub-theory which is ‘isomorphic’ to **CIT**, by the construction mentioned earlier, according to which we associate the Hilbert space $\hat{X} := \mathbb{C}^X$ to the finite set X and the CPTP map $\hat{T} : \text{End}(\mathbb{C}^X) \rightarrow \text{End}(\mathbb{C}^Y)$ given by $\hat{T}(A) = \sum_{x \in X, y \in Y} t_x(y) \langle x | A | x \rangle | y \rangle \langle y |$ to the classical channel $T = (t_x)_{x \in X} : X \rightarrow Y$. This last example, however, is different than the others in a significant way: The identity transformations id_X in the smaller theory do not coincide with the identity transformations $\text{id}_{\hat{X}}$ in the larger theory. We have already noticed in Remark 1.1.12 that this is not an artefact of this specific embedding of **CIT** in **QIT**, but a living condition of any such embedding. Similarly, the correct definition of ‘sub-theory’ should not require identities to agree.

Definition 1.2.20. (Sub-Theories.)

Let Θ be a theory. A *sub-theory* of Θ is a theory Θ_0 for which

- $\text{Sys}_{\Theta_0} \subseteq \text{Sys}_\Theta$, and the trivial system and composition of systems in Θ_0 are the same as in Θ ;
- for any systems $\mathcal{X}, \mathcal{Y} \in \text{Sys}_{\Theta_0}$, $\text{Trans}_{\Theta_0}(\mathcal{X}, \mathcal{Y}) \subseteq \text{Trans}_\Theta(\mathcal{X}, \mathcal{Y})$ and the serial and parallel compositions in Θ_0 are the same as in Θ .

■

Remark 1.2.21. Since identity transformations (and swapping transformations) in Θ_0 are not required to be the same as in Θ , a sub-theory is a weaker notion than that of a (symmetric monoidal) sub-category. \spadesuit

The concept of sub-theory is relevant in the context of example appropriation, since it paves the way for an explosion: Whenever we have a theory Θ , we can choose from

it any collection of systems and transformations and consider as new example theory the sub-theory Θ_0 of Θ that this collection generates.¹¹

Example 1.2.22. (**FinSets**^{*}.)

FinSets^{*} is the the sub-theory of **Sets**^{*} generated by the finite sets and all the functions between them. That is, **FinSets**^{*} has as systems finite sets X, Y, Z, \dots , composing under the cartesian product, and the transformations from X to Y are functions $f : X \rightarrow Y$, composing serially and parallelly as in **Sets**^{*}. Observe that **FinSets**^{*} can also be regarded as a sub-theory of **CIT**, generated this time by all the systems, but only the transformations $(t_x)_{x \in X} : X \rightarrow Y$ corresponding to deterministic functions. \blacklozenge

Example 1.2.23. (**T – REX**.)

T – REX is the sub-theory of **FinSets**^{*} with all the same systems, but consisting only of the surjective functions. Note that the serial and parallel compositions of surjective functions are surjective, and that the trashes and identities are surjective. Though it has the ‘cartesian product’ for parallel composition, the theory **T – REX** is not a cartesian theory (e.g. there is no transformation $D : X \rightarrow X \times X$ with marginals equal to id_X when $|X| \geq 2$). In general, the theory **T – REX** is extremely strange, and it will provide us with many counterexamples throughout. One bizarre feature is that, though it is much more intricate than a thin theory, it retains the property of having states only on those systems which are isomorphic to **1**. \blacklozenge

Readers who know about the theory of computation and algorithms may also define a sub-theory of **Sets**^{*} whose systems are collections of strings and whose transformations are algorithms (computable functions).

1.3 Pictorial Syntax

So far, we have used an algebraic syntax in terms of the symbols ‘ \parallel ’ and ‘ \circ ’ to represent composite transformations in a theory. Whereas this is in principle unproblematic, it is more or less undebatable that, to the human eye, the nature of already rather simple compositions can be obfuscated by the algebraic notation. For example, if s is a state on $\mathcal{Z}_1 \parallel \mathcal{Z}_2$ and if $T_1 : \mathcal{X}_1 \parallel \mathcal{Z}_1 \rightarrow \mathcal{Y}_1$ and $T_2 : \mathcal{X}_2 \parallel \mathcal{Z}_2 \rightarrow \mathcal{Y}_2$ are transformations, then what is the appropriate intuition about the transformation $(T_1 \parallel T_2) \circ (\text{id}_{\mathcal{X}_1} \parallel s \parallel \text{id}_{\mathcal{X}_2})$?

1.3.A Pictures for Algebra

A viable and very effective solution is to introduce a pictorial syntax for systems, transformations and the two modes of composition. The basic idea is to pictorially denote a transformation $T : \mathcal{X} \rightarrow \mathcal{Y}$ as a box with incoming and outgoing wires, as such:

$$-\mathcal{X}-\boxed{T}-\mathcal{Y}- \tag{1.16}$$

This reinforces the interpretation of T as a ‘process’ which transforms input from the system \mathcal{X} to outputs on the system \mathcal{Y} . If \mathcal{X} and \mathcal{Y} are composite systems, say $\mathcal{X} = \mathcal{X}_1 \parallel \mathcal{X}_2$ and $\mathcal{Y} = \mathcal{Y}_1 \parallel \mathcal{Y}_2 \parallel \mathcal{Y}_3$, we may detail the representation by drawing one wire for each factor:

¹¹To the extent that the collections are not so bizarre that this procedure cannot be formalised in the formal language which we use, cf. the earlier comments on sets versus proper classes.

$$\begin{array}{c}
-\mathcal{X}_1- \\
-\mathcal{X}_2- \\
\boxed{T} \\
-\mathcal{Y}_1- \\
-\mathcal{Y}_2- \\
-\mathcal{Y}_3-
\end{array}
\quad (1.17)$$

Note that the associativity $(\mathcal{Y}_1 \parallel \mathcal{Y}_2) \parallel \mathcal{Y}_3 = \mathcal{Y}_1 \parallel (\mathcal{Y}_2 \parallel \mathcal{Y}_3)$ is built into this notation, and if we moreover agree that the trivial system $\mathbf{1}$ may be represented by empty space (no wire at all), then relations as $\mathcal{Z} \parallel \mathbf{1} = \mathcal{Z} = \mathcal{Z} \parallel \mathbf{1}$ are also automatic. As such, we may represent a state $s : \mathbf{1} \rightarrow \mathcal{X}$ by a box with no incoming wires, and a trash $\text{tr}_{\mathcal{Z}} : \mathcal{Z} \rightarrow \mathbf{1}$ by a box with no outgoing wires, as

$$\boxed{s} - \mathcal{X} - \quad , \quad \text{respectively} \quad - \mathcal{Z} - \boxed{\text{tr}} \quad . \quad (1.18)$$

The serial composition of transformations $-\mathcal{X}-\boxed{T}-\mathcal{Y}-$ and $-\mathcal{Y}-\boxed{S}-\mathcal{Z}-$ is represented by indeed connecting them serially, as

$$-\mathcal{X}-\boxed{T}-\mathcal{Y}-\boxed{S}-\mathcal{Z}- \quad (1.19)$$

(with suitable modifications when the systems are represented as composites with several wires). This visual representation agrees with the Western reading direction from left to right, but disagrees with the unfortunate direction of functional composition mimicked in the notation ' $S \circ T$ '.

The parallel composition of $-\mathcal{X}_1-\boxed{T_1}-\mathcal{Y}_1-$ and $-\mathcal{X}_2-\boxed{T_2}-\mathcal{Y}_2-$ is represented by vertical juxtaposition, as

$$\begin{array}{c}
-\mathcal{X}_1-\boxed{T_1}-\mathcal{Y}_1- \\
-\mathcal{X}_2-\boxed{T_2}-\mathcal{Y}_2-
\end{array}
\quad (1.20)$$

(again modified if there are more incoming and outgoing wires to each box). Importantly, this notation is consistent with the convention of representing composite systems as a stack of wires as in Eq. (1.17): The parallel composition $T_1 \parallel T_2$ indeed has domain $\mathcal{X}_1 \parallel \mathcal{X}_2$ and codomain $\mathcal{Y}_1 \parallel \mathcal{Y}_2$.

An identity transformation $\text{id}_{\mathcal{Z}} : \mathcal{Z} \rightarrow \mathcal{Z}$ can be represented simply as the wire

$$-\mathcal{Z}- \quad , \quad (1.21)$$

and when combined with the convention on representing serial composition by serial connection, this consistently suggests the facts that $\text{id}_{\mathcal{Z}} \circ T = T$ and $S \circ \text{id}_{\mathcal{Z}} = S$ for transformations $T : \mathcal{X} \rightarrow \mathcal{Z}$ and $S : \mathcal{Z} \rightarrow \mathcal{Y}$.

Within this pictorial syntax, the transformation $T := (T_1 \parallel T_2) \circ (\text{id}_{\mathcal{X}_1} \parallel s \parallel \text{id}_{\mathcal{X}_2})$ from above is now drawn as

$$\begin{array}{c}
-\mathcal{X}_1- \\
-\mathcal{Z}_1- \\
-\mathcal{Z}_1- \\
-\mathcal{X}_2- \\
\boxed{s} \\
\boxed{T_1} \\
\boxed{T_2} \\
-\mathcal{Y}_1- \\
-\mathcal{Y}_2-
\end{array}
\quad (1.22)$$

This picture aids the intuition about the transformation T , by providing the interpretation that a state s is shared across two different sites, at each of which a transformation is then applied to form locally at site i a connection from \mathcal{X}_i to \mathcal{Y}_i , using the part of the state s which is present at site i .

We can use the pictorial syntax not only to better display the nature of composite transformations, but also to manipulate them more transparently. For instance, we can apply to T the trash $\text{tr}_{\mathcal{Y}_1} : \mathcal{Y}_1 \rightarrow \mathbf{1}$ to the upper wire in (1.22), and in parallel apply $\text{id}_{\mathcal{Y}_2}$ to the lower wire, thus computing that

$$\begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \boxed{s} \\
 \text{---}\mathcal{X}_2\text{---} \\
 \text{---}\mathcal{Y}_1\text{---} \\
 \boxed{\text{tr}} \\
 \text{---}\mathcal{Y}_2\text{---}
 \end{array}
 \begin{array}{c}
 \boxed{T_1} \\
 \boxed{T_2}
 \end{array}
 =
 \begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \boxed{s} \\
 \text{---}\mathcal{X}_2\text{---} \\
 \text{---}\mathcal{Y}_2\text{---}
 \end{array}
 \begin{array}{c}
 \boxed{\text{tr}} \\
 \boxed{T_2}
 \end{array}
 =
 \begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \boxed{s} \\
 \text{---}\mathcal{X}_2\text{---} \\
 \text{---}\mathcal{Y}_2\text{---}
 \end{array}
 \begin{array}{c}
 \boxed{\text{tr}} \\
 \boxed{\text{tr}} \\
 \boxed{T_2}
 \end{array},
 \quad (1.23)$$

resting for the equalities on Lemma 1.1.15. (For clarity, the labels ‘ \mathcal{X}_1 ’ and ‘ \mathcal{X}_2 ’ on the internal wires have been omitted, and we shall often omit wire labels when they are irrelevant or clear from context.)

The transformation $\begin{array}{c} \boxed{s} \\ \boxed{\text{tr}} \\ \text{---}\mathcal{X}_2\text{---} \\ \boxed{T_2} \\ \text{---}\mathcal{Y}_2\text{---} \end{array}$ can now be renamed as $\text{---}\mathcal{X}_2\text{---}\boxed{T'_2}\text{---}\mathcal{Y}_2\text{---}$, and the above computation then altogether suggests that by trashing the system \mathcal{Y}_1 from T we obtain something of the form $\text{tr}_{\mathcal{X}_1} \llbracket T'_2$ for some transformation $T'_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$. (Physically this means that the output on \mathcal{Y}_2 alone is unaffected by the input to \mathcal{X}_1 ; we shall consider such *non-signalling* properties systematically in the next section.)

The pictorial syntax laid out above is ubiquitous in the literature, and has been since the introduction of symmetric monoidal categories. Ref. [Sel10] gives a detailed and formal survey of general *graphical calculi* for monoidal categories, and therein the author essentially attributes the boxes-and-wires representation to Roger Penrose, dating back almost 50 years ([Pen71]).

Remark 1.3.1. (On the Validity of Pictorial Reasoning.)

It is only fair for the reader to question the exact relationship between the algebraic and pictorial syntaxes. Is it really the case that one can deduce the algebraic identity

$$(\text{tr}_{\mathcal{Y}_1} \llbracket \text{id}_{\mathcal{Y}_2}) \circ ((T_1 \llbracket T_2) \circ (\text{id}_{\mathcal{X}_1} \llbracket s \llbracket \text{id}_{\mathcal{X}_2})) = \text{tr}_{\mathcal{X}_1} \llbracket T'_2 \quad (1.24)$$

for some $T'_2 : \mathcal{X}_2 \rightarrow \mathcal{Y}_2$ just by reference to the graphical manipulation in Eq. (1.23)? It is instructive to consider for each step of the manipulation the translation from pictures to algebraic symbolism, and to verify that this is indeed the case.

One might worry that in general care must be taken when translating between the conclusions of pictorial and algebraic manipulations. It would be dangerous if graphical manipulations suggested algebraically invalid derivations, and conversely regrettable if some algebraic derivations had no graphical counterpart. Fortunately, it is a mathematical fact that this does not happen (see Thm. 2.1 in [Sel10] for a formal statement, and Ref. [JS91] for an even more precise treatment). This fact is the ultimate power – and justification – of the pictorial syntax.

I have only defined the pictorial syntax by examples, and I shall take the attitude of not being uptight about the formal correspondence between derivations in the two syntaxes. Rather, pictures will be used where they enlighten, and with the implicit understanding that they really represent underlying algebraic arguments which can be distilled upon desire. ✖

1.3.B Interfaces and Channels

Though the pictorial approach to notation is intuitively superior, and equivalent to the algebraic with regards to deductive power, there is a sense in which it is distinct from the algebra it intends to represent.

The problem has to do with multiplicity of wires. According to the pictorial syntax, a composite system $\mathcal{X} = \mathcal{X}_1 \parallel \mathcal{X}_2$ can, when it appears as domain or codomain of a transformation, be represented equally well as $\text{---}\mathcal{X}\text{---}$ and as $\begin{array}{c} \text{---}\mathcal{X}_1\text{---} \\ \text{---}\mathcal{X}_2\text{---} \end{array}$. Similarly, though we made the convention that we can choose to represent the system $\mathbf{1}$ by empty space, we made no convention that we must do so. Thus, in the thin theory $(\mathbb{N}, \cdot, 1, \geq)$ (Example 1.2.13), for instance, the three pictures

$$\begin{array}{c} \text{---}2\text{---} \\ \text{---}3\text{---} \\ \text{---}5\text{---} \end{array}, \quad \begin{array}{c} \text{---}6\text{---} \\ \text{---}5\text{---} \end{array}, \quad \begin{array}{c} \text{---}1\text{---} \\ \text{---}1\text{---} \\ \text{---}30\text{---} \\ \text{---}1\text{---} \end{array} \quad (1.25)$$

are all valid pictorial representations of the system 30, and in the theory **QIT** the pictures

$$\begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{\text{id}_{\mathbb{C}^2 \otimes \mathbb{C}^2}} \begin{array}{c} \text{---}\mathbb{C}^2 \otimes \mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2 \otimes \mathbb{C}^2\text{---} \end{array}, \quad \begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{\text{id}_{\mathbb{C}^2 \otimes \mathbb{C}^2}} \begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array}, \quad \begin{array}{c} \text{---}\mathbb{C}^2 \otimes \mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2 \otimes \mathbb{C}^2\text{---} \end{array} \boxed{\text{id}_{\mathbb{C}^2 \otimes \mathbb{C}^2}} \begin{array}{c} \text{---}\mathbb{C}^2 \otimes \mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2 \otimes \mathbb{C}^2\text{---} \end{array} \quad (1.26)$$

are all valid representations of the transformation $\text{id}_{\mathbb{C}^2 \otimes \mathbb{C}^2}$ (there is one more). Now, this ambiguity is not in conflict with the equivalence of derivations in the algebraic and pictorial syntaxes mentioned earlier,¹² but it does raise the following question:

If it is not transformations between systems, then what exactly are we drawing when we draw a picture in the pictorial syntax?

We will need the viewpoint that the various pictorial representations of the same formal object correspond to various ways of thinking about that object. For example, in **CIT**, the difference between drawing a state $p : \mathbf{1} \rightarrow X$ as

$$\boxed{p}\text{---}X\text{---} \quad \text{and as} \quad \text{---}\mathbf{1}\text{---}\boxed{p}\text{---}X\text{---} \quad (1.27)$$

will be that the latter represents the viewpoint that the process under consideration accepts an input to the trivial system, whereas the former does not. When we define

¹²This is essentially because the various representations of the identity transformations are isomorphisms, and so can act as regrouping devices which collect a bunch of wires into one, much like cable collectors used in offices.

causal specifications in Chapter 4, the utility of this is amplified, since it will allow the interpretation that the trivial system can act as a *button* that has to be pushed before the state p materialises, whereas absence of the trivial system as input signifies that the state p was always present.

Similarly, if X is a composite system, say $X = X_1 \times X_2$, the two drawings

$$\boxed{p} \begin{array}{l} \text{---}X_1\text{---} \\ \text{---}X_2\text{---} \end{array} \quad \text{and} \quad \boxed{p} \text{---}X\text{---} \quad (1.28)$$

will signify that the state p is presented across two different ports (first drawing), or that all of p is given through one single port (second drawing). Generally, we want wires to represent *ports* which can be labelled by distinct names and associated with (possibly different) systems. None of this intuition is reflected a priori in the formal algebraic structure of symmetric monoidal categories, so if we want we an algebraic counterpart to the pictures we have to encode the additional structure explicitly.

We do this by introducing the concept of an *interface*. Intuitively, an interface in a theory Θ is specified by a finite collection of ports distinguishable from each other by unique names, and each of which able to transmit information of a certain type as defined by an associated system in Θ . Formally, we can pack this information compactly by defining an interface as a *map* from its set of port names to the associated systems:

Definition 1.3.2. (Interfaces in Θ .)

Let Θ be a theory. An *interface in Θ* is a map, \mathbb{X} , whose domain, denoted $\text{ports}(\mathbb{X})$, is a finite set of so-called *ports* (or *port names*), and which assigns to each port $p \in \text{ports}(\mathbb{X})$ a system $\mathbb{X}(p) \in \text{Sys}_\Theta$. ■

Example 1.3.3. (Simple Interfaces.)

Let us call an interface *simple* if it has just a single port. A simple interface \mathbb{X} is completely specified by the name $p \in \text{ports}(\mathbb{X})$ and the associated system $\mathbb{X}(p) \in \text{Sys}_\Theta$. ◆

Example 1.3.4. (Pairs of Qubits.)

Let \mathbb{X} be an interface in **QIT** with a single port, say $\text{ports}(\mathbb{X}) = \{1\}$, and with $\mathbb{X}(1) = \mathbb{C}^2 \otimes \mathbb{C}^2$. Let \mathbb{X}' be an interface with two ports, say $\text{ports}(\mathbb{X}') = \{a, b\}$ and with $\mathbb{X}'(a) = \mathbb{X}'(b) = \mathbb{C}^2$. The interfaces \mathbb{X} and \mathbb{X}' are distinct, though they both represent the same total system $\mathbb{C}^2 \otimes \mathbb{C}^2$ (i.e. a pair of qubits). The interface \mathbb{X} represents a scenario in which this system is thought of as a single entity, while \mathbb{X}' represents a bipartite situation where each factor is separate and clearly labelled. The pictorial depictions of \mathbb{X} and \mathbb{X}' are illustrated by Eq. (1.26). ◆

Example 1.3.5. (Interfaces in Thin Theories.)

In a thin theory Θ , identified with the pre-ordered quasi-commutative monoid $(M, \star, \succeq, 1)$, the distinction between a system and an interface is not foggy at all; indeed, a system in M is an element of M whereas an interface \mathbb{X} corresponds to a tuple of elements in M , indexed by the port names $\text{ports}(\mathbb{X})$. ◆

Defining interfaces as maps allows for smooth notation in many regards:

Since maps are set-theoretically identified with their graphs, the notation $\mathbb{X}_0 \subseteq \mathbb{X}$ makes sense for interfaces \mathbb{X}_0 and \mathbb{X} , and it means that $\text{ports}(\mathbb{X}_0) \subseteq \text{ports}(\mathbb{X})$ with $\mathbb{X}_0(p) = \mathbb{X}(p)$ for all $p \in \text{ports}(\mathbb{X}_0)$. We shall say in this case that \mathbb{X}_0 is a *sub-interface* of \mathbb{X} . If \mathbb{X}_0 is

a sub-interface of \mathbb{X} , we can form the *complementary interface* $\mathbb{X} \setminus \mathbb{X}_0$, namely the unique sub-interface of \mathbb{X} with $\text{ports}(\mathbb{X} \setminus \mathbb{X}_0) = \text{ports}(\mathbb{X}) \setminus \text{ports}(\mathbb{X}_0)$.

Two interfaces \mathbb{X}_1 and \mathbb{X}_2 are called *parallelly composable* if $\text{ports}(\mathbb{X}_1) \cap \text{ports}(\mathbb{X}_2) = \emptyset$, i.e. if they share no port names. In this case the union $\mathbb{X}_1 \cup \mathbb{X}_2$ is an interface, namely the one given by $\text{ports}(\mathbb{X}_1 \cup \mathbb{X}_2) = \text{ports}(\mathbb{X}_1) \cup \text{ports}(\mathbb{X}_2)$ and $(\mathbb{X}_1 \cup \mathbb{X}_2)(p_j) = \mathbb{X}_j(p_j)$ for $p_j \in \text{ports}(\mathbb{X}_j)$. The interface $\mathbb{X}_1 \cup \mathbb{X}_2$ is called the *composite* of \mathbb{X}_1 and \mathbb{X}_2 . **Importantly, we do not allow the parallel composition of interfaces which have port names in common.** This has the positive consequence that all input wires in a drawing are identified by a unique port name, as are all output systems. (We do allow, however, that an input and output wire are labelled by the same name, as long as these correspond to the same system, cf. the definition of a channel below.)

We define *the trivial interface*, denoted \mathbb{I} , as the empty map, i.e. as the unique interface with no port names. The difference between the two drawings in Eq. (1.27) is exactly that the input interface of the left one is the trivial interface, whereas the input interface of the right one has a single port with associated system $\mathbf{1}$. The trivial interface \mathbb{I} is a sub-interface of every interface and it is parallelly composable with every interface \mathbb{X} , with $\mathbb{I} \cup \mathbb{X} = \mathbb{X}$.

Finally, it is obvious that every non-trivial interface \mathbb{X} factors into simple interfaces (i.e. interfaces with a single port) as $\mathbb{X}_1 \cup \dots \cup \mathbb{X}_{|\mathbb{X}|}$ where $|\mathbb{X}| := |\text{ports}(\mathbb{X})|$ is the *size of \mathbb{X}* , and where the interfaces $\mathbb{X}_1, \dots, \mathbb{X}_{|\mathbb{X}|}$ are the simple sub-interfaces of \mathbb{X} .

Having defined interfaces, we must replace transformations with entities whose domains and codomains are interfaces rather than systems. We will call them *channels*.¹³ A channel T from an interface \mathbb{X} to an interface \mathbb{Y} is more or less a transformation from \mathcal{X} , the total system corresponding to the interface \mathbb{X} , to \mathcal{Y} , the total system corresponding to the interface \mathbb{Y} . However, channels need to be defined in such a way that they retain the information about how the various ports of the interfaces relate to these total systems. For example, if $\text{ports}(\mathbb{X}) = \{1, 2\}$ and $\text{ports}(\mathbb{Y}) = \{a, b\}$, and if $T : \mathbb{X}(1) \parallel \mathbb{X}(2) \rightarrow \mathbb{Y}(a) \parallel \mathbb{Y}(b)$ is a transformation, then we want to think of this transformation as identical to the transformation $T' : \mathbb{X}(1) \parallel \mathbb{X}(2) \rightarrow \mathbb{Y}(b) \parallel \mathbb{Y}(a)$ given by $T' = \sigma_{\mathbb{Y}(a), \mathbb{Y}(b)} \circ T$, provided that we know the port sequence in each case (a-b versus b-a). Pictorially,

$$\begin{array}{c} 1 - \mathbb{X}(1) - \boxed{T} - \mathbb{Y}(a) - a \\ 2 - \mathbb{X}(2) - \boxed{T} - \mathbb{Y}(b) - b \end{array} \quad \text{and} \quad \begin{array}{c} 1 - \mathbb{X}(1) - \boxed{T} - \mathbb{Y}(b) - b \\ 2 - \mathbb{X}(2) - \boxed{T} - \mathbb{Y}(a) - a \end{array} \quad (1.29)$$

should be merely different drawings of the same thing. In general, we want to define a channel as an equivalence class of transformations related by swappings on their input and output systems.

Given an interface \mathbb{Z} , a bijection $\ell : \{1, \dots, |\mathbb{Z}|\} \rightarrow \text{ports}(\mathbb{Z})$ is a choice of enumeration of the ports in \mathbb{Z} . Given such an enumeration ℓ , we define *the system (in \mathbb{Z}) corresponding to ℓ* as the composite

$$\mathbb{Z}[\ell] := \mathbb{Z}(\ell(1)) \parallel \mathbb{Z}(\ell(2)) \parallel \dots \parallel \mathbb{Z}(\ell(|\mathbb{Z}|)), \quad (1.30)$$

with the natural convention that if $\mathbb{Z} = \mathbb{I}$ and ℓ is the empty enumeration, then $\mathbb{Z}[\ell] = \mathbf{1}$. We formalise channels as follows:

¹³Somewhat confusingly, this conflicts with the use of the term ‘quantum [classical] channel’ for the transformations in **QIT [CIT]**, but the confusion should not cause any serious harm.

Definition 1.3.6. (Channels in Θ .)

Let \mathbb{X} and \mathbb{Y} be interfaces in a theory Θ , subject to the condition that if $\mathfrak{p} \in \text{ports}(\mathbb{X}) \cap \text{ports}(\mathbb{Y})$, then $\mathbb{X}(\mathfrak{p}) = \mathbb{Y}(\mathfrak{p})$. A *channel from \mathbb{X} to \mathbb{Y}* is a triple $(T, \mathbb{X}, \mathbb{Y})$, where T is family of transformations in Θ ,

$$T = (\ell_Y T_{\ell_X} : \mathbb{X}[\ell_X] \rightarrow \mathbb{Y}[\ell_Y])_{\ell_X, \ell_Y}, \quad (1.31)$$

indexed by the collection of enumerations ℓ_X of $\text{ports}(\mathbb{X})$ and ℓ_Y of $\text{ports}(\mathbb{Y})$, and subject to the condition that for any permutations π_X of $\text{ports}(\mathbb{X})$ and π_Y of $\text{ports}(\mathbb{Y})$, we have

$$\pi_Y \circ \ell_Y T_{\pi_X \circ \ell_X} = \sigma_{\pi_Y} \circ \ell_Y T_{\ell_X} \circ \sigma_{\pi_X}^{-1}, \quad (1.32)$$

where $\sigma_{\pi_X} : \mathbb{X}[\ell_X] \rightarrow \mathbb{X}[\pi_X \circ \ell_X]$ and $\sigma_{\pi_Y} : \mathbb{Y}[\ell_Y] \rightarrow \mathbb{Y}[\pi_Y \circ \ell_Y]$ denote the transformations which swap the individual factors in accordance with the permutations π_X and π_Y , respectively. \blacksquare

We will write $T : \mathbb{X} \rightarrow \mathbb{Y}$ to indicate that T is a channel from \mathbb{X} to \mathbb{Y} , and by abuse of notation we often write simply ‘ T ’ for all of the individual components $\ell_Y T_{\ell_X}$.

What we have obtained in summary is the following:

- In a diagram, every *wire* corresponds to a simple interface; all incoming wires are formally given distinct port names (whereas their systems may coincide), and likewise so are all outgoing wires.
- Every *box* corresponds to a channel between the interfaces that connect to it; as such, the order in which these are drawn from top to bottom is insignificant, as the port names formally keep track of the matching.
- A port name may occur both as input and output, but in that case the associated systems are identical (and we will really think of it as the same physical port).

In effect, we have replaced the theory Θ by another category, $\text{IC}(\Theta)$, namely the *category of interfaces and channels in Θ* , where the objects are interfaces and the morphisms are channels, which compose serially in an obvious (but tedious) way. This category is moreover ‘partially’ symmetric monoidal, in the sense that some interfaces (namely those with no overlapping port names) are deemed parallely composable, and some channels are deemed parallely composable (namely those for which the interfaces are parallely composable, and for which the parallel composition does not result in a violation of the condition that a repeated port name can correspond to different systems). Though this might all seem very formal, we will almost never explicitly flesh out the formalities and so an intuitive understanding of these concepts suffices.

As a last convention, it is only natural to introduce an abbreviation which confuses the reader by overwriting the above correspondence – namely, that in order to keep diagrams visually simple we will often write

$$-\mathbb{X}-\boxed{T}-\mathbb{Y}- \quad \text{as abbreviation for} \quad \mathbb{X} \begin{array}{c} \text{---} \\ \vdots \\ \text{---} \end{array} \boxed{T} \begin{array}{c} \text{---} \\ \vdots \\ \text{---} \end{array} \mathbb{Y}, \quad (1.33)$$

with T a channel from \mathbb{X} to \mathbb{Y} . In this way, we do not need to draw a specific number of wires (or to draw dotted lines) when arguing about general channels.

1.3.C Normal Theories

The pictorial syntax raises another problem which can be viewed as a discrepancy between pictures and algebra, though in a more subtle way than before. Consider the equation

$$\begin{array}{ccc} \mathcal{X} \text{---} \boxed{T} \text{---} \mathcal{Y} & = & \mathcal{X} \text{---} \boxed{T'} \text{---} \mathcal{Y} \\ \mathcal{Z} \text{---} \boxed{S} \text{---} \mathcal{W} & = & \mathcal{Z} \text{---} \boxed{S'} \text{---} \mathcal{W} \end{array} . \quad (1.34)$$

Is it necessarily the case that $T = T'$ and $S = S'$? The graphical language somehow invites the presumption that this holds, whereas the algebraic equation $T \parallel S = T' \parallel S'$ does not enforce the same suspicion. To understand whether the conclusion is legal, it clearly suffices (by symmetry) to know whether the identity $T = T'$ is implied by $T \parallel S = T' \parallel S'$. Moreover, since the identity $T \parallel S = T' \parallel S'$ evidently implies $T \parallel \text{tr}_{\mathcal{Z}} = T' \parallel \text{tr}_{\mathcal{Z}}$, it actually suffices to know the answer in the case $S = S' = \text{tr}_{\mathcal{Z}}$. This motivates the following definition:

Definition 1.3.7. (Normality.)

A theory Θ is called *normal* if for any systems $\mathcal{X}, \mathcal{Y}, \mathcal{Z} \in \text{Sys}_{\Theta}$, and for any transformations $T, T' : \mathcal{X} \rightarrow \mathcal{Y}$, it holds that

$$T \parallel \text{tr}_{\mathcal{Z}} = T' \parallel \text{tr}_{\mathcal{Z}} \Rightarrow T = T'. \quad (1.35)$$

■

Our reasons for regarding **Sets** and **Graphs** as poorer theories than **Sets*** and **Graphs*** is precisely that the latter are not normal:

Example 1.3.8. (**Sets** is Not Normal.)

For any sets X and Y , and for any functions $f, g : X \rightarrow Y$, we have $f \times \text{tr}_{\emptyset} = g \times \text{tr}_{\emptyset}$, as both functions are empty. (Here, tr_{\emptyset} is the unique empty map $\emptyset \rightarrow \mathbf{1}$.) ♦

Example 1.3.9. (**Graphs** is Not Normal.)

The (direct) product of graphs G and H , $G \times H$, renders (u_1, v_1) adjacent to (u_2, v_2) if and only if both u_1 is adjacent to u_2 in G and v_1 is adjacent to v_2 in H . In particular, if \bullet denotes a graph with precisely one vertex and no edges, then $G \times \bullet$ is the graph with the same vertex set as G , but with no edges at all. Now, if $f : G \rightarrow H$ is any graph homomorphism from G to some graph H , then $f \times \text{tr}_{\bullet} : G \times \bullet \rightarrow H$ is the homomorphism which acts as f on the vertex set. Since $G \times \bullet$ has no edges, however, we cannot tell from $f \times \text{tr}_{\bullet}$ how f acts on edges, so in particular if we choose H to be a graph for which some pair of vertices has two distinct edges between them, we can find $f, f' : G \rightarrow H$ distinct with $f \times \text{tr}_{\bullet} = f' \times \text{tr}_{\bullet}$. ♦

Luckily, we have the following:

Proposition 1.3.10. (*Virtually All Theories are Normal.*)

The following theories are normal:

1. *All theories in which every system has at least one state.*
2. *All thin theories.*
3. *All sub-theories of a normal theory.*

In fact, every single theory presented as an example in Section 1.2 is normal.

Proof. 1. If s is a state on \mathcal{Z} , then for any transformation $T : \mathcal{X} \rightarrow \mathcal{Y}$ we have $T = T \parallel \text{id}_1 = (T \parallel \text{tr}_{\mathcal{Z}}) \circ (\text{id}_{\mathcal{X}} \parallel s)$, so T can be recovered uniquely from $T \parallel \text{tr}_{\mathcal{Z}}$. Pictorially,

$$\boxed{T} = \begin{array}{c} \boxed{T} \\ \boxed{s} \text{---} \boxed{\text{tr}} \end{array} . \quad (1.36)$$

2. A thin theory is normal for the simple reason that any two transformations from \mathcal{X} to \mathcal{Y} are identical, regardless of whether they satisfy additional constraints or not.

3. If Θ is a sub-theory of a theory Θ , then the normality condition for Θ concerns a smaller class of systems and transformations than that corresponding to a normality condition for Θ .

The final statement is left for the reader to ponder. □

We make the following convention:

From now on, we always use the term ‘theory’ to mean normal theory.

1.4 Summary and Outlook

In this chapter, we have seen a mathematical definition of ‘physical’ theories (Definition 1.1.6), and many examples of this concept (Section 1.2). We have also defined the special kinds of transformations called *states* (Definition 1.1.16), *reversibles* and *isomorphisms* (Definition 1.1.20). Finally, we have discussed a pictorial syntax for representing the elements of a theory, and we have defined in this context the notions of *interfaces* (Definition 1.3.2) and *channels* (Definition 1.3.6). We also observed the notion of a theory being *normal* (Definition 1.3.7), and we will employ that assumption implicitly from now on.

The reign of category theory in mathematics has hatched a practice which is nowadays standard in most of its disciplines: When defining a class of mathematical structures (for examples: physical theories), one should define along with them the appropriate notion of *structure-preserving maps* (or, *homomorphisms*) between them – mathematical extremists would say that one has not even understood what is integral to a structure before one has defined what are the structure-preserving maps.

As such, it would be natural to develop a theory of homomorphisms between theories, as alluded to in Remark 1.1.12. It is not a priori clear that anything substantial can be gained by this exercise, but from initial thoughts it seems to me that the notion of a theory Θ which allows an embedding $\Gamma : \mathbf{CIT} \rightarrow \Theta$ is highly interesting: By virtue of such an embedding, we can basically talk about *probabilities* and *classicality* and construct in Θ a *convex structure*, in the same way as we do in \mathbf{QIT} .¹⁴ In existing treatments (such as Ref. [CDP10]), probabilistic and convex structure have to be separately planted on top of the compositional structure in the theory.

¹⁴Though one apparently has to assume by axiom that there exist for any finite tuple (T_1, \dots, T_n) of transformations $T_k : \mathcal{X} \rightarrow \mathcal{Y}$ in Θ a unique transformation $T : \mathcal{X} \parallel \Gamma(\{1, \dots, n\}) \rightarrow \mathcal{Y}$ which reads off the value k in the classical system $\Gamma(\{1, \dots, n\})$ and chooses the respective transformation T_k .

Chapter 2

Dilations

§1. Introduction and Outline.

Consider a channel $\begin{array}{c} \mathcal{X}_1 \\ \boxed{L} \\ \mathcal{X}_2 \end{array} \begin{array}{c} \mathcal{Y}_1 \\ \mathcal{Y}_2 \\ \mathcal{Y}_3 \end{array}$ and suppose that, for some reason or another, we only have access to the third output port, corresponding to the system \mathcal{Y}_3 . By the interpretation of trashes as discarding, the channel we really ‘see’ is the *marginal*

$$\begin{array}{c} \mathcal{X}_1 \\ \boxed{T} \\ \mathcal{X}_2 \end{array} \mathcal{Y}_3 \quad := \quad \begin{array}{c} \mathcal{X}_1 \\ \boxed{L} \\ \mathcal{X}_2 \end{array} \begin{array}{c} \mathcal{Y}_1 \text{ tr} \\ \mathcal{Y}_2 \text{ tr} \\ \mathcal{Y}_3 \end{array} . \quad (2.1)$$

For example, in the theory **Sets**^{*}, L would be a function, and T would be just the third component of that function, with the two other components discarded. (In **CIT** or **QIT**, the act of marginalisation may affect correlations between outputs as well.) We shall now develop the first symptoms of an obsession, namely the quest for answering the following question:

Knowing T , what are the possible L that T could have come from?

In fact, we will include among the possible L also some which have inputs additional to those of \mathcal{X}_1 and \mathcal{X}_2 , as long as these inputs do not *signal* to the interface accessible to us. Precisely, given a channel $\begin{array}{c} \mathbb{X} \\ \boxed{T} \\ \mathbb{Y} \end{array}$ we define a *dilation of T* to be a channel $\begin{array}{c} \mathbb{X} \\ \boxed{L} \\ \mathbb{D} \end{array} \begin{array}{c} \mathbb{Y} \\ \mathbb{E} \\ \text{tr} \end{array}$ such that

$$\begin{array}{c} \mathbb{X} \\ \boxed{L} \\ \mathbb{D} \end{array} \begin{array}{c} \mathbb{Y} \\ \mathbb{E} \\ \text{tr} \end{array} = \begin{array}{c} \mathbb{X} \\ \boxed{T} \\ \mathbb{D} \end{array} \begin{array}{c} \mathbb{Y} \\ \text{tr} \end{array} , \quad (2.2)$$

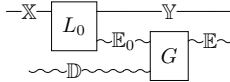
and we thus desire to understand all possible dilations of T . In some theories, this question will be a purely mathematical curiosity (for example in thin theories), but mostly it has a clear physical interpretation, as in the information theories **CIT** and **QIT**: What we will imagine is a dichotomy between interfaces \mathbb{X} and \mathbb{Y} which are *accessible* or *open* to us, and interfaces \mathbb{D} and \mathbb{E} which are *inaccessible* or *hidden*. (This dichotomy is visually represented using wiggly wires above.) The hidden interfaces may be controlled by untrustworthy

agents, or may represent simply parts of Nature which is not within our reach; we will often use the word *the environment* as umbrella term. By interacting with the open interfaces, we may establish that we are interacting with the channel T ,¹ but the point is that T could be ‘implemented’ in many different ways across the hidden interfaces, as formalised by its various dilations.

In this chapter, I introduce a very general theory of dilations, applicable to all of the theories we have seen in Chapter 1. Indeed, the singular reason for demanding in Definition 1.1.6 that the trivial system $\mathbf{1}$ be terminal, is that this requirement minimally facilitates the notion of dilations. Whereas this is well-known ([CDP10, Chi14a]), the structure of dilations has, to the best of my knowledge, never been studied systematically in the literature before, neither for specific theories nor for theories in general. As such, all the theory and results laid out in this chapter are new.

Dilations, Non-Signalling and DiVincenzo’s Property. Naturally, we begin in Section 2.1 by defining *marginalisation* and *dilations*. Due to the unconventional choice of defining dilations as two-sided (i.e. with the interface \mathbb{D} possibly non-trivial), the concept of dilations will be closely related to that of *non-signalling*, which is also defined and exemplified.

Some 20 years ago, D. DiVincenzo proposed a later confirmed conjecture ([BGNP01, ESW02]) about the structure of non-signalling channels in **QIT**, which is here promoted to a general property that a theory might have. In the context of dilations, the DiVincenzo

property entails that any dilation is of the form  for some channel G

and some one-sided dilation L_0 , thus effectively reducing the study of dilations to that of one-sided dilations. In later sections it is proved that many theories enjoy this property (but it fails for example in some thin theories and in the theory **T – REX**).

A Hierarchy of Dilations. The most important observation about the dilations of a given channel $\text{---}X\text{---}[T]\text{---}Y\text{---}$ is that they do not constitute just a messy class of channels; they are naturally *ordered*. For example, consider in the theory **CIT** the state $\text{---}p\text{---}X\text{---}$, with p the uniform distribution on the set $X = \{\square, \square, \square, \square, \square, \square\}$, representing the throw of a fair die. Among its one-sided dilations $\text{---}p\text{---}X\text{---}$ are the one with $E = X$ and $\ell(x, x') = p(x)\delta_{x,x'}$ (diagonal distribution), the one with $X = \{\text{even}, \text{odd}\}$ and $\ell(x, k) = p(x)\delta_{\text{par}(x),k}$ (where $\text{par}(x)$ denotes the parity of x), and the ones with E arbitrary and $\ell = p \otimes r$ for some state r on E . These dilations have clear meanings in terms of *side-information*: The first corresponds to the environment keeping record of the precise outcome of die throw, the second corresponds to the environment knowing only the parity, and the third kinds correspond to the environment possessing information completely independent of the die throw.

The *dilational ordering* introduced in Section 2.2 formalises the intuition that the second dilation may be derived from the first, and the third in turn from the second. In general, the dilational ordering (Definition 2.2.1) will describe how one dilation $\text{---}X\text{---}[L']\text{---}Y\text{---}$ might

¹An operationally minded reader might inquire as to how we would determine that the channel we have access to is T ; even in the case where T is a function between finite sets, to confirm its identity we need to check its value on every possible input. Whereas this problem is very real, we shall completely ignore it. One can either regard it as an assumption of mathematical character, or imagine iterated use of identical and independent copies of the channel.

be *derivable* from another dilation $\begin{array}{c} \text{---X---} \\ \boxed{L} \\ \sim_{\mathbb{E}} \end{array} \text{---Y---}$ by constructions taking place in the environment, specifically, if there exists a channel G such that

$$\begin{array}{c} \text{---X---} \\ \boxed{L'} \\ \sim_{\mathbb{E}'} \end{array} \text{---Y---} = \begin{array}{c} \text{---Y---} \\ \boxed{L} \\ \sim_{\mathbb{E}} \end{array} \begin{array}{c} \text{---Y---} \\ \boxed{G} \\ \sim_{\mathbb{E}'} \end{array} . \quad (2.3)$$

At this early point in our story, we have to require L to be one-sided for the definition of the relation to be well-tempered (though L' can be two-sided, cf. Definition 2.2.1); in fact, the generalisation to arbitrary dilations cannot be executed in a satisfactory manner without treating causality, which we postpone to Chapter 4.² As such, there is a sense in which the dilational ordering introduced here is provisional. Nevertheless, its study is by no means futile – it uncovers over-arching principles which govern the structure of dilations in many theories, and whose distillation and ramifications are the topic of the remainder of the chapter.

Dilational Axioms. In the extreme case, the dilational ordering collapses to a single level, namely when the channel $\begin{array}{c} \text{---X---} \\ \boxed{T} \\ \text{---Y---} \end{array}$ has only the *trivial* dilations $\begin{array}{c} \text{---X---} \\ \boxed{T} \\ \text{---Y---} \\ \sim_{\mathbb{D}} \end{array} \begin{array}{c} \text{---Y---} \\ \boxed{S} \\ \sim_{\mathbb{E}} \end{array}$ obtained by parallel composition. This phenomenon will be called *dilational purity* of T (Section 2.2.A), and though it is significant we cannot (and want not) expect such astounding simplicity in the structure of dilations of general channels.

However, in Section 2.3 we shall see two milder forms of collapse. They will be presented as possible *dilational axioms* which a theory might or might not comply to, but as we will see they reign in virtually all theories of physical interest:

Completeness. The first such axiom is the existence of *complete dilations*, which are simply largest (greatest) elements in the dilational ordering, i.e. dilations from which any other dilation can be derived. In the above example with the fair die, the dilation which represented a copy of the outcome was in fact complete. In general, completeness in **CIT** is obtained by copying (Theorem 2.3.21), and likewise in cartesian theories completeness is obtained by copying (Theorem 2.3.6). In **QIT**, on the other hand, Stinespring dilations serve as complete (Theorem 2.3.22).

The existence of complete dilations in a theory is ultimately an information-theoretic principle: *There is a largest amount of side-information to be had.* It provides conceptual clarity in the dilational ordering, in addition to being of technical importance in many later considerations.

Localisability of Side-Information. The other such axiom governs not the structure of dilations of a single channel, but how the dilational structure behaves under the two modes of composition in the theory. As such, it actually comprises two separate axioms:

²The essential problem can be appreciated in a concrete example: If $\begin{array}{c} \text{---X---} \\ \boxed{L} \\ \sim_{\mathbb{E}} \end{array} \text{---Y---}$ is a dilation of $\begin{array}{c} \text{---X---} \\ \boxed{T} \\ \text{---Y---} \end{array}$, and $\begin{array}{c} \text{---Y---} \\ \boxed{M} \\ \sim_{\mathbb{K}} \end{array} \text{---Z---}$ a dilation of $\begin{array}{c} \text{---Y---} \\ \boxed{S} \\ \text{---Z---} \end{array}$ whose hidden input interface matches the hidden output interface of L , then the channel $\begin{array}{c} \text{---X---} \\ \boxed{L} \\ \sim_{\mathbb{D}} \end{array} \begin{array}{c} \text{---Y---} \\ \boxed{M} \\ \sim_{\mathbb{K}} \end{array} \text{---Z---}$ is a dilation of $\begin{array}{c} \text{---X---} \\ \boxed{T} \\ \text{---Z---} \end{array}$; it would seem reasonable that in the dilational ordering, this dilation should be greater than $\begin{array}{c} \text{---X---} \\ \boxed{L} \\ \sim_{\mathbb{D}} \end{array} \begin{array}{c} \text{---Y---} \\ \boxed{M} \\ \sim_{\mathbb{K}} \end{array} \text{---Z---}$ (which is also a dilation), but the operation required to realise this is that of *contracting* the interface \mathbb{E} , an operation which at this point is not feasible.

Two channels $\text{--}\mathbb{X}\text{--}\boxed{T}\text{--}\mathbb{Y}\text{--}$ and $\text{--}\mathbb{Y}\text{--}\boxed{S}\text{--}\mathbb{Z}\text{--}$ in succession define a third, serially composed channel $\text{--}\mathbb{X}\text{--}\boxed{T}\text{--}\boxed{S}\text{--}\mathbb{Z}\text{--}$. *Temporal localisability of side-information* will refer to the principle that any (one-sided) dilation of the serial composition can be derived from the composition of (one-sided) dilations of the individual channels T and S , that is, can be ‘temporally localised’ in the composition. Specific examples will demonstrate that this axiom is not automatic, but its information-theoretic interpretation certainly suggests that it should hold in physically sensible theories: It expresses that any side-information about a composite channel must stem from side-information about the first and second channels. Ref. [CDP11] considers a related principle (‘Atomicity of Composition’ – see details below), and in the words of one of the authors ([CS16], p. 194), “*Although [the failure of this principle] is logically conceivable, it raises a puzzling questions: What is the extra information about?*”

The obvious sibling to temporal localisability is *spatial localisability of side-information*, which asserts that any (one-sided) dilation of a parallel composition $\begin{array}{c} \text{--}\mathbb{X}_1\text{--}\boxed{T_1}\text{--}\mathbb{Y}_1\text{--} \\ \text{--}\mathbb{X}_2\text{--}\boxed{T_2}\text{--}\mathbb{Y}_2\text{--} \end{array}$ must derive from a parallel composition of (one-sided) dilations. Its interpretation and motivation is analogous to that of temporal localisability: Side-information about the independent execution of two channels should be ‘spatially localisable’, as side-information about the individual channels. Again, this principle is not automatic, as demonstrated by examples.

The two localisability principles hold in many theories, in particular in all cartesian theories (Theorem 2.3.16), and in the two information theories **CIT** and **QIT** (Theorem 2.3.17 and Theorem 2.3.18). The principles have interesting consequences, most of which factor through two specific consequences: Firstly, temporal localisability implies that any channel has a reversible dilation, a fact that will be used in later sections over and over. Secondly, spatial localisability implies the DiVincenzo property, demonstrating that this property is not in any way a quantum feature but owes its validity to a different information-theoretic principle entirely, a circumstance which has apparently not been noted before.

Universal Dilations. For some applications, it will be important that the channel G which derives in the environment a given dilation from a complete dilation (cf. Eq. (2.3)) is unique. This phenomenon is embodied in the notion of a *universal* dilation (Definition 2.4.1) presented in Section 2.4. In the present chapter, the significance of universal dilations over complete dilations is mainly expressed by Proposition 2.4.8. In Chapter 4, however, it will be instrumental in the form of Lemma 4.2.5, which ultimately implies that we can introduce a notion of *contraction* for channels (Theorem 4.2.6).

Cartesian theories are easily proved to have universal dilations (Theorem 2.4.5), and it is also not hard to prove the existence of universal dilations for **CIT**, by cutting down the hidden system of a complete dilation (Theorem 2.4.6). The theory **QIT** has universal dilations too, namely minimal Stinespring dilations (Theorem 2.4.11), but this fact is non-trivial and can be seen as a generalisation of the injectivity of the Choi-Jamiołkowski isomorphism (Lemma 2.4.12), a result which might be of independent interest.

Purification. The two information theories **CIT** and **QIT** share the properties of completeness (even universality), and of spatial and temporal localisability. They are distinguished, however, by *purifiability* (Definition 2.5.1), which most elegantly can be described by saying that every quantum channel has a dilation which is *a complete dilation of itself*, or, equivalently, which is dilationally pure. This phenomenon very effectively separates the natures of **QIT** and **CIT**, and Section 2.5 serves to demonstrate this separation in a sequence of results about purifiable theories. Purifiable theories include other creatures than

QIT, though, for example we shall see that a thin theory $(M, \star, 1, \succeq)$ is purifiable precisely if it satisfies the cancellation law $x \star y \succeq x \star z \Rightarrow y \succeq z$.

Section 2.5.A is about isomorphisms in purifiable theories, and proves among other things an extremely general version (Corollary 2.5.7) of the *No Broadcasting Theorem* ([BCF+96]).

In Section 2.5.B, we prove a structure-theorem for reversible channels (Theorem 2.5.11), which in the case of **QIT** specialises to the result that a quantum channel is reversible precisely if it is the tensoring with a state followed by an isometric conjugation.

Finally, in Section 2.5.C, I introduce a general notion of *complementarity* between channels, generalising the concept of complementary for quantum channels ([DS05]), and made possible by a combination of purifiability with the other dilational principles. In particular, this entails an abstract version of the complementarity between reversible and ‘completely forgetful’ channels (Theorem 2.5.18). This complementarity comprises a long list of impossibility (‘no go’) theorems, since it implies not only a strong version of the No Broadcasting Theorem (and hence the *No Cloning Theorem* ([WZ82]) and the *No Deletion Theorem* ([PB00])), but also the *No Hiding* ([BP07]) and *No Masking* ([MPSS18]) *Theorems*.

§2. Comparison to Existing Literature.

Purity from Dilations. The most well-known notion of ‘purity’ in information theory is arguably that of *probabilistic purity*, which refers simply to convex extremality. This purity notion in general differs from that of dilational purity (Definition 2.2.9), cf. Remark 2.2.13. The difference is qualitative too, in the sense that convex purity requires an ambient convex structure to make sense, whereas dilational purity is well-defined in general.

The idea of defining purity categorically in terms of dilations is not new, but has been considered also in Ref. [Chi14b]; in lack of the DiVincenzo property, the definition given there is weaker than the one here (the latter producing the intended notion), but the ideas behind the two definitions are identical. Other authors ([CH17]) have considered a different categorically definable purity notion (in terms of so-called *factorisation systems*), but that notion is distinct from dilational purity as it reproduces convex purity in the case of **CIT**. See Remark 2.2.13 for further details.

Purification Principles. ‘Purifiability’ in the abstract has been considered in the literature before, in the form of the *Purification Postulate* of Refs. [CDP10, CDP11]. The seminal result of that work is that, within a large ground class of theories, five axioms along with the Purification Postulate characterise the theory **QIT**. Whereas the Purification Postulate is intimately related with *purifiability* in the sense proposed here (Definition 2.5.1), there are differences, as detailed in Remark 2.5.4. The most important difference is that the purifiability notion proposed in this chapter is not only less restrictive than that of Refs. [CDP10, CDP11] within their ground class of theories – its effective scope is also larger outside of this ground class. Indeed, in Refs. [CDP10, CDP11] the Purification Postulate is often used in combination with a number of ‘standing assumptions’ (some of which pertain to probabilistic structure), and this has the subtle side-effect that some theories which accidentally satisfy the Purification Postulate, but violate the standing assumptions, are nothing like quantum theory.³ For example, every cartesian theory (e.g. the theory **Sets**^{*}) would comply the the Purification Principle of Refs. [CDP10, CDP11], which pertains only to purification of states. In contrast, this does not happen for the notion of purifiability presented here, essentially because it concerns purification of all kinds of channels.

³The authors are of course aware of this circumstance, cf. the remark following Cor. 6 on p. 16 in Ref. [CDP10].

It should also be mentioned that the results derived in this chapter based on the purifiability principle (most importantly those of Theorem 2.5.11, Theorem 2.5.17 and Theorem 2.5.18) do not have counterparts in the work of Refs. [CDP10, CDP11]. Furthermore, they are all derived in the simplest possible framework, from few principles, which are principles exclusively about dilations. Finally, it is interesting to note that (as explained in Remark 2.5.4), the Purification Postulate of Refs. [CDP10, CDP11] implies the existence of complete dilations, whereas the framework here separates these two phenomena.

Dilations in General and Axioms about Them. It is well-known in the literature that general theories in the sense of Definition 1.1.6 facilitate the concepts of marginalisation and of (one-sided) dilations ([CDP10, Chi14a]). Be that as it may, the structure of dilations has to the best of my knowledge never been studied systematically before. (Refs. [CDP10, CDP11] has a dilational axiom in focus – the ‘Purification Principle’ mentioned just above – but this axiom is easily stated without referring to any dilational ordering.) Whereas a principle vaguely related to that of temporal localisability is mentioned in Ref. [CDP11] (see details below), the dilational ordering in general, and the completeness, (spatial) localisability and universality axioms in particular have no counterparts in the existing literature.

Temporal Localisability and ‘Atomicity of Composition’. Ref. [CDP11] considers a principle under the name *Atomicity of Composition*, according to which the serial composition of two *atomic* transformations is atomic; the notion of an ‘atomic’ transformation has, as far as I can see, no exact counterpart in the framework of this thesis, because its definition ultimately relies on convex structure (it can be seen as a broadening of convex extremality), but it is meant to capture the impossibility of extracting further information from a transformation. As such, atomicity is in spirit similar to dilational purity (formally it subsumes it), and a thoughtless identification of the two notions would thus translate to the serial composition of pure transformations being pure. That statement follows rigorously from temporal localisability (Definition 2.3.11), and though it is evidently less general it would seem that the two principles are really cut from the same stone.

The DiVincenzo Property and Spatial Localisability. The conjecture of DiVincenzo about the structure of non-signalling channels in **QIT** has been studied extensively. It was first mentioned publicly in Ref. [BGNP01], where it was shown to be true for a particular class of measurement channels. In Ref. [ESW02], its validity was extended to all channels, with a proof based on the uniqueness of Stinespring dilations; the proof was then shortened in Ref. [PHHH06].⁴ Finally, in Ref. [CDP10], it was shown that the structure-theorem holds not only in **QIT**, but more generally in the presence of the Purification Postulate, using a uniqueness of dilations similar to that of Stinespring’s.

Proposition 2.3.20 in this chapter shows that, contrary to what the preceding works suggest (see also p. 201 in [CS16]), the structure-theorem has nothing to do with purification. Rather, its root is the principle of spatial localisability (Definition 2.3.10), from which the DiVincenzo property can be derived in one sentence. It is true that in **QIT** the validity of spatial localisability is best proved by an argument involving Stinespring dilations (this provides the link to the existing proofs), but the principle itself is conceptually independent of purification – it governs many non-quantum theories, including classical information theory.

⁴When viewed correctly, however, the proof in Ref. [PHHH06] is conceptually equivalent to that in Ref. [ESW02], as it too uses uniqueness of Stinespring dilations albeit in the form of purification of states as mediated by the Choi-Jamiołkowski isomorphism.

§3. Contributions.

The original contributions of this chapter are the following:

1. Defining the *dilational ordering* among dilations of a channel T in a general theory Θ (as modelled by a symmetric monoidal category with terminal unit).
2. Identifying *completeness* (Definition 2.3.2) and *localisability* (Definition 2.3.12) as axioms about the dilational ordering which are simple to state, easy to interpret, true in the majority of example theories, and potent in deriving consequences. Among these consequences are the existence of reversible dilations (Proposition 2.3.19), and the DiVincenzo property (Proposition 2.3.20), which has often before been presented as a quantum feature resting on Stinespring’s dilation theorem.
3. Identifying the concept of *universal dilations*, which reveal the precise structure of the dilational ordering (Proposition 2.4.8), have occasional technical significance (e.g. in Theorem 2.5.11), and will be instrumental in defining in Chapter 4 a new operation (*contraction*) generalising and connecting the works of Refs. [JSV96] and [CDP09]. In proving the existence of universal dilations in **QIT**, the injectivity of the Choi-Jamiołkowski isomorphism ([JLF13, dP67, Cho75, Jam72]) is generalised in a non-trivial way from pure states of full rank to isometric channels of full rank (Lemma 2.4.12). This is the most technical result of the chapter.
4. Presenting a notion of *purifiability* (Definition 2.5.1) which broadens the scope of the ‘Purification Postulate’ of Refs. [CDP10, CDP11], and in conjunction with other dilational principles allows the derivation of new general results, including a structure-theorem for reversible transformations (Theorem 2.5.11), and an abstract notion of *complementarity* between channels (Theorem 2.5.17), which places the notion of complementary quantum channels ([DS05]) in a general setting.
5. Proving for the first time a general version of the complementarity between reversible and completely forgetful⁵ channels (Theorem 2.5.18) known from quantum information theory. This theorem moreover comprises many of the impossibility theorems known in the case of **QIT**, including the fairly recent *No Hiding* and *No Masking Theorems* ([BP07, MPSS18]), which to my best knowledge have not been proved in general frameworks before.

Most importantly, however, all of the principles which we entertain in this chapter are *principles about dilations and nothing else*. As such, the results of the chapter constitute a proof of concept, demonstrating that a remarkable amount of features, including many features of quantum information theory, can be derived from a small set of principles which pertain exclusively to the nature of dilations. In particular, neither the statement of these principles nor the proofs of their consequences depend in any way on probabilistic structure or on ‘dagger-compact structure’ ([AC04]). In short, the framework is the most general possible in which it even makes sense to speak of dilations.

⁵A *completely forgetful* channel is a channel which is the serial composition of a trash and a state.

2.1 A Terminal Trinity: Marginals, Dilations and Non-Signalling

In Definition 1.1.6 we imposed on a theory the requirement that the trivial system $\mathbf{1}$ be terminal. This requirement supports the concepts of *marginalisation*, *dilation* and *non-signalling*. These three concepts are reviewed and exemplified in this section.

Marginalisation we already saw in the form of (somewhat intermittently named) *factor projections* in connection with cartesian theories in Section 1.2.B. To treat general marginalisation succinctly and precisely, however, we need to work in the realm of channels and interfaces (cf. Definition 1.3.2 and Definition 1.3.6):

Definition 2.1.1. (Projections and Marginals.)

Let \mathbb{Y} be an interface in a theory Θ , and let $\mathbb{Y}_0 \subseteq \mathbb{Y}$ be a sub-interface. The *projection from \mathbb{Y} to \mathbb{Y}_0* is the channel $\pi_{\mathbb{Y} \rightarrow \mathbb{Y}_0} : \mathbb{Y} \rightarrow \mathbb{Y}_0$ given by $\text{id}_{\mathbb{Y}_0} \parallel \text{tr}_{\mathbb{Y} \setminus \mathbb{Y}_0}$. Pictorially, $\pi_{\mathbb{Y} \rightarrow \mathbb{Y}_0}$ is

$$\begin{array}{c} \text{---}\mathbb{Y}_0\text{---} \\ \text{---}\mathbb{Y} \setminus \mathbb{Y}_0\text{---} \boxed{\text{tr}} \end{array} \left(\text{abbreviating} \begin{array}{c} \text{---}\mathbb{Y}_0\text{---} \\ \vdots \\ \text{---}\boxed{\text{tr}}\text{---} \\ \mathbb{Y} \setminus \mathbb{Y}_0 : \\ \text{---}\boxed{\text{tr}}\text{---} \end{array} \right). \quad (2.4)$$

For a channel $T : \mathbb{X} \rightarrow \mathbb{Y}$, the \mathbb{Y}_0 -marginal of T is the serial composition $\pi_{\mathbb{Y} \rightarrow \mathbb{Y}_0} \circ T : \mathbb{X} \rightarrow \mathbb{Y}_0$. Pictorially, the \mathbb{Y}_0 -marginal of T is

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \boxed{T} \text{---}\mathbb{Y}_0\text{---} \\ \text{---}\mathbb{Y} \setminus \mathbb{Y}_0\text{---} \boxed{\text{tr}} \end{array} \left(\text{abbreviating} \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \vdots \\ \text{---}\mathbb{X} : \\ \text{---}\boxed{T}\text{---} \\ \vdots \\ \text{---}\mathbb{Y} \setminus \mathbb{Y}_0\text{---} \\ \text{---}\boxed{\text{tr}}\text{---} \end{array} \right). \quad (2.5)$$

■

A given channel $T : \mathbb{X} \rightarrow \mathbb{Y}$ has $2^{|\mathbb{Y}|}$ marginals, one for each sub-interface of \mathbb{Y} . The \mathbb{Y} -marginal is the full channel T , whereas the \mathbb{I} -marginal is the trash $\text{tr}_{\mathbb{X}}$. The intermediate marginals are typically more interesting.

Example 2.1.2. (Marginals in \mathbf{Sets}^* .)

Let $\begin{array}{c} \text{---}X\text{---} \\ \boxed{f} \\ \text{---}Z\text{---} \end{array} \begin{array}{c} \text{---}Y\text{---} \\ \text{---} \end{array}$ be a channel in \mathbf{Sets}^* . Since \mathbf{Sets}^* is cartesian, f must take the form $x \mapsto (f_1(x), f_2(x))$ for some functions $f_1 : X \rightarrow Y$ and $f_2 : X \rightarrow Z$, and these uniquely determine f . These two functions are precisely the marginals of f corresponding to the two simple interfaces for Y and Z . (This example generalises to any cartesian theory.) ◆

Example 2.1.3. (Marginals in \mathbf{CIT} .)

In \mathbf{CIT} , marginalisation specialises to the usual notion in terms of summing over elements of the discarded systems. In particular, the marginals of a state $\begin{array}{c} \text{---}X\text{---} \\ \boxed{p} \\ \text{---}Y\text{---} \end{array}$ are the states $\begin{array}{c} \text{---}X\text{---} \\ \boxed{p_x} \end{array}$ and $\begin{array}{c} \text{---}Y\text{---} \\ \boxed{p_y} \end{array}$ with densities given by $p_x(x) = \sum_{y \in Y} p(x, y)$ and $p_y(y) = \sum_{x \in X} p(x, y)$, respectively. Marginals of a general channel in \mathbf{CIT} is given by marginalising each component. ◆

Example 2.1.4. (Marginals in QIT.)

Marginalisation in QIT also specialises to the usual notion, since projections are given by the *partial trace*. For example, both non-trivial marginals of the maximally entangled qubit

state $\begin{array}{c} \boxed{\phi} \\ \text{---} \mathbb{C}^2 \text{---} \\ \text{---} \mathbb{C}^2 \text{---} \\ \text{---} \tau_{\mathbb{C}^2} \text{---} \end{array}$ with vector representative $|\phi\rangle = \frac{|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle}{\sqrt{2}}$ are given by the fully mixed qubit state $\begin{array}{c} \boxed{\tau_{\mathbb{C}^2}} \\ \text{---} \mathbb{C}^2 \text{---} \end{array}$. \blacklozenge

Example 2.1.5. (Marginals in a Thin Theory.)

Let Θ be a thin theory, identified with the pre-ordered quasi-commutative monoid $(M, \star, 1, \succeq$

$\left. \begin{array}{c} \text{---} y_1 \text{---} \\ \text{---} y_2 \text{---} \\ \text{---} y_2 \text{---} \end{array} \right\} \text{---} x \text{---}$ is simply the relationship $x \succeq y_1 \star y_2 \star y_3$. In addition the two

trivial marginals $x \succeq y_1 \star y_2 \star y_3$ and $x \succeq 1$, there are six marginal relationships: $x \succeq y_2 \star y_3$, $x \succeq y_1 \star y_3$, $x \succeq y_1 \star y_2$, $x \succeq y_1$, $x \succeq y_2$ and $x \succeq y_3$. \blacklozenge

Now, *dilations* are usually defined as dual to marginals, but we shall employ a slightly more general definition, according to which dilations of a channel $T : \mathbb{X} \rightarrow \mathbb{Y}$ may not only have outputs additional to those of \mathbb{Y} , but also inputs additional to those of \mathbb{X} . The reason for this is both one of more robust interpretation (two-sided dilations model general *side-computations*, whereas one-sided dilations model only *escaping side-information*), and one of mathematical robustness.

Definition 2.1.6. (Dilations.)

Let $T : \mathbb{X} \rightarrow \mathbb{Y}$ be a channel in Θ . A *dilation* of T is a channel $L : \mathbb{X} \cup \mathbb{D} \rightarrow \mathbb{Y} \cup \mathbb{E}$ whose \mathbb{Y} -marginal is $T \parallel \text{tr}_{\mathbb{D}}$, i.e. for which

$$\begin{array}{c} \text{---} \mathbb{X} \text{---} \\ \sim \mathbb{D} \sim \\ \boxed{L} \\ \text{---} \mathbb{Y} \text{---} \\ \sim \mathbb{E} \sim \\ \boxed{\text{tr}} \end{array} = \begin{array}{c} \text{---} \mathbb{X} \text{---} \\ \sim \mathbb{D} \sim \\ \boxed{T} \\ \text{---} \mathbb{Y} \text{---} \\ \sim \mathbb{D} \sim \\ \boxed{\text{tr}} \end{array} . \quad (2.6)$$

The interfaces \mathbb{D} and \mathbb{E} are called *hidden* or *inaccessible (relative to T)*, and by convention we use wiggly lines to pictorially represent the wires corresponding to the hidden interfaces. A dilation is called *one-sided* if $\mathbb{D} = \mathbb{I}$. \blacksquare

Remark 2.1.7. (On the Importance of Normality for Dilations.)

The normality condition on Θ , cf. Definition 1.3.7 ($T \parallel \text{tr}_{\mathbb{Z}} = T' \parallel \text{tr}_{\mathbb{Z}} \Rightarrow T = T'$), is crucial for the concept of dilations not to be pathological. Indeed, if Θ is not normal and if $T, T' : \mathbb{X} \rightarrow \mathbb{Y}$ are distinct channels with $T \parallel \text{tr}_{\mathbb{Z}} = T' \parallel \text{tr}_{\mathbb{Z}}$ for some interface \mathbb{Z} , then the channel $L := T \parallel \text{tr}_{\mathbb{Z}} = T' \parallel \text{tr}_{\mathbb{Z}}$ is a dilation of both T and T' . In other words, one cannot tell from the dilation L which is the channel that it dilates, although we know which are the hidden interfaces. \blacklozenge

Remark 2.1.8. (On the Importance of Interfaces for Dilations.)

Consider in the thin theory $(\mathbb{N}_0, \max, 0, \geq)$ from Example 1.2.12 a channel $\text{---} x \text{---} \boxed{} \text{---} y \text{---}$ with $10 \geq x \geq y$. The channel $\begin{array}{c} \text{---} x \text{---} \boxed{} \text{---} y \text{---} \\ \sim 10 \sim \boxed{} \sim 10 \sim \end{array}$ is a dilation. As a channel, it determines the channel $\text{---} x \text{---} \boxed{} \text{---} y \text{---}$ that it dilates, because it formally determines its input and output interfaces which determine x and y . As a transformation, however, it is simply the relationship $\max\{10, x\} \geq \max\{10, y\}$, i.e. $10 \geq 10$, with no recognition of the values of x and y , that is, of the channel that it dilates. \blacklozenge

To characterise the dilations of a given channel, it obviously suffices to characterise those whose hidden interfaces are simple. Nevertheless, it is generally a challenging task to do this (indeed, this is why it is possible to write an entire chapter about it). At this point, we have only a few such characterisations within reach:

Example 2.1.9. (Dilations in Cartesian Theories.)

Let \boxed{T}_{-X-Y-} be a channel in **Sets**^{*} (or indeed in any other cartesian theory). We can determine all of its dilations. Indeed, given a dilation $\begin{array}{c} \sim D \sim \\ \boxed{L} \\ \sim E \sim \\ -X- \end{array} \begin{array}{c} \boxed{L} \\ \boxed{\text{tr}} \\ -Y- \end{array}$, consider the channel

$$\begin{array}{c} \sim D \sim \\ \boxed{L} \\ \sim E \sim \\ -X- \end{array} \begin{array}{c} \boxed{L} \\ \boxed{\text{tr}} \\ -Y- \end{array}, \quad (2.7)$$

where (following Ref. [Fri20]) the channel $\boxed{\text{tr}}_{-X-}$ denotes the diagonal function $x \mapsto (x, x)$ (or, in a general cartesian theory, the unique channel with both marginals equal to id_X). The channel (2.7) is easily seen to have precisely the same marginals as L , so by the cartesian property it must in fact be equal to L . (In the absence of a hidden input, this is easy to understand intuitively in the theory **Sets**^{*}: A one-sided dilation L is determined by its marginals; one of those is T , and the other is $\begin{array}{c} \sim E \sim \\ \boxed{L} \\ \boxed{\text{tr}} \\ -X- \end{array}$.) Now, this shows that any dilation of T is of the form

$$\begin{array}{c} \sim D \sim \\ \boxed{G} \\ \sim E \sim \\ -X- \end{array} \begin{array}{c} \boxed{G} \\ \boxed{\text{tr}} \\ -Y- \end{array}, \quad (2.8)$$

for some channel G . Conversely, it is obvious that the channel (2.8) is a dilation for any choice of G , so these are precisely the dilations of T . ◆

Example 2.1.10. (Dilations in Thin Theories.)

Let Θ be a thin theory, described the pre-ordered quasi-commutative monoid $(M, \star, 1, \succeq)$. Interfaces \mathbb{X} and \mathbb{Y} in Θ are simply tuples of elements $(x_1, \dots, x_n) \in M^n$ and $(y_1, \dots, y_m) \in M^m$, and a channel $\mathbb{X} \succeq \mathbb{Y}$ is nothing but the relationship $x_1 \star \dots \star x_n \succeq y_1 \star \dots \star y_m$. A dilation of $\mathbb{X} \succeq \mathbb{Y}$ is determined by a pair of interfaces $\mathbb{D} = (d_1, \dots, d_k)$ and $\mathbb{E} = (e_1, \dots, e_\ell)$ such that $\mathbb{X} \cup \mathbb{D} \succeq \mathbb{Y} \cup \mathbb{E}$, i.e. $x_1 \star \dots \star x_n \star d_1 \star \dots \star d_k \succeq y_1 \star \dots \star y_m \star e_1 \star \dots \star e_\ell$. Of course, these can be understood by understanding the dilations with simple hidden interfaces, namely the pairs of elements $d, e \in M$ which satisfy $(x_1 \star \dots \star x_n) \star d \succeq (y_1 \star \dots \star y_m) \star e$. (The general dilations then arise by \star -factorisations of d and e .) ◆

In other theories, even determining all the one-sided dilations of states is non-trivial:

Example 2.1.11. (One-Sided Dilations of States in CIT.)

Let \boxed{p}_{-X-} be a state in **CIT**. What are its one-sided dilations? Taking the hidden interface \mathbb{E} to be simple, we are looking for states $\begin{array}{c} \sim E \sim \\ \boxed{\ell} \\ -X- \end{array}$ with $\begin{array}{c} \sim E \sim \\ \boxed{\ell} \\ \boxed{\text{tr}} \\ -X- \end{array} = \boxed{p}_{-X-}$.

Some obvious dilations are the *trivial* dilations of the form $\begin{array}{c} \sim E \sim \\ \boxed{r} \\ \boxed{p} \\ -X- \end{array}$, corresponding to

independent side-information; but we also have e.g. the dilation $\boxed{\hat{p}} \begin{array}{c} \sim X \sim \\ \sim X \sim \end{array} := \boxed{p} \begin{array}{c} \sim X \sim \\ \sim X \sim \end{array}$ given by copying, i.e. with density given by $\hat{p}(x, x') = \delta_{x, x'} p(x)$. Intuitively, a copy seems like the strongest sort of side-information we can have, and this intuition can be formalised: Every one-sided dilation ℓ of p is in fact of the form

$$\boxed{\hat{p}} \begin{array}{c} \sim X \sim \\ \sim X \sim \end{array} \begin{array}{c} \boxed{G} \\ \sim E \sim \end{array} \quad (2.9)$$

for some channel G . It is easy to see that any G yields a dilation. The converse owes to the existence of conditional distributions: Because the X -marginal of a dilation ℓ is p , the density $\ell : X \times E \rightarrow [0, 1]$ can be expressed as $\ell(x, e) = g_x(e)p(x)$ for probability densities $g_x : E \rightarrow [0, 1]$, which we may interpret as the *conditional distribution* (w.r.t. ℓ) of the E -output given the X -output. However, this precisely means that if we take $G = (g_x)_{x \in X}$, then (2.9) is satisfied. (Observe that g_x is uniquely determined if $p(x) > 0$, whereas it can be chosen arbitrarily if $p(x) = 0$.) \blacklozenge

Example 2.1.12. (One-Sided Dilations of States in **QIT**.)

Let $\boxed{\varrho} \begin{array}{c} \sim \mathcal{H} \sim \\ \sim \mathcal{H} \sim \end{array}$ be a state in **QIT**, and let us determine its one-sided dilations $\boxed{\xi} \begin{array}{c} \sim \mathcal{E} \sim \\ \sim \mathcal{H} \sim \end{array}$.

By a well-known result, the state ϱ has a *purification*, that is, a dilation $\boxed{\pi} \begin{array}{c} \sim \mathcal{H} \sim \\ \sim \mathcal{H} \sim \end{array}$ with π (probabilistically) pure, and this purification is unique up to isometric conjugations on the hidden system, cf. the paragraph on Stinespring Representations in the preliminary section on the thesis. Like the classical copy in Example 2.1.11, the fixed purification π can help us

determine all possible dilations of ϱ : Indeed, if $\boxed{\xi} \begin{array}{c} \sim \mathcal{E} \sim \\ \sim \mathcal{H} \sim \end{array}$ is any dilation, let $\boxed{\pi'} \begin{array}{c} \sim \mathcal{E}' \sim \\ \sim \mathcal{E} \sim \\ \sim \mathcal{H} \sim \end{array}$ be

a purification of ξ ; then π and π' are both purifications of ϱ , and by the uniqueness clause we find an isometric channel Σ such that $\boxed{\pi} \begin{array}{c} \sim \mathcal{H} \sim \\ \sim \mathcal{H} \sim \end{array} \begin{array}{c} \boxed{\Sigma} \\ \sim \mathcal{E}' \sim \\ \sim \mathcal{E} \sim \end{array} = \boxed{\pi'} \begin{array}{c} \sim \mathcal{E}' \sim \\ \sim \mathcal{E} \sim \\ \sim \mathcal{H} \sim \end{array}$. By trashing \mathcal{E}' ,

we find that $\boxed{\xi} \begin{array}{c} \sim \mathcal{E} \sim \\ \sim \mathcal{H} \sim \end{array}$ equals

$$\boxed{\pi} \begin{array}{c} \sim \mathcal{H} \sim \\ \sim \mathcal{H} \sim \end{array} \begin{array}{c} \boxed{\Gamma} \\ \sim \mathcal{E} \sim \end{array}, \quad (2.10)$$

with Γ denoting the marginal of Σ . Hence, the one-sided dilations of $\boxed{\varrho} \begin{array}{c} \sim \mathcal{H} \sim \\ \sim \mathcal{H} \sim \end{array}$ are precisely those states of the form (2.10) for some channel Γ . \blacklozenge

Based on Example 2.1.9, Example 2.1.11 and Example 2.1.12, it is very natural to suspect that the ‘one-dilation-to-rule-them all’-phenomenon is more than a mere coincidence. This will be the topic of Section 2.3.A when we introduce and study *complete* dilations.

For now, however, we shall raise a different question, namely how we might hope to systematically understand the two-sided dilations in **CIT** and **QIT** (and for that matter in other theories) based on the one-sided dilations. To have a better perspective on this question, it is beneficial to introduce an additional concept, namely that of *non-signalling* ([CS16]):

Definition 2.1.13. (Non-Signalling.)

Let $T : \mathbb{X} \rightarrow \mathbb{Y}$ be a channel in Θ , and let \mathbb{X}_0 and \mathbb{Y}_0 be sub-interfaces of \mathbb{X} and \mathbb{Y} , respectively. We say that T is *non-signalling from \mathbb{X}_0 to \mathbb{Y}_0* if there exists a channel $T' : \mathbb{X} \setminus \mathbb{X}_0 \rightarrow \mathbb{Y}_0$ such that

$$\begin{array}{c} \text{---} \mathbb{X} \setminus \mathbb{X}_0 \text{---} \\ \text{---} \mathbb{X}_0 \text{---} \end{array} \boxed{T} \begin{array}{c} \text{---} \mathbb{Y}_0 \text{---} \\ \text{---} \mathbb{Y} \setminus \mathbb{Y}_0 \text{---} \end{array} \boxed{\text{tr}} = \begin{array}{c} \text{---} \mathbb{X} \setminus \mathbb{X}_0 \text{---} \\ \text{---} \mathbb{X}_0 \text{---} \end{array} \boxed{T'} \begin{array}{c} \text{---} \mathbb{Y}_0 \text{---} \\ \text{---} \text{tr} \end{array} . \quad (2.11)$$

■

The interpretation of non-signalling is disclosed by its very name: Non-signalling from \mathbb{X}_0 to \mathbb{Y}_0 means that no inputs in the interface \mathbb{X}_0 affect the outputs in the interface \mathbb{Y}_0 .

Example 2.1.14. (Bell-Channels and Non-Signalling.)

Consider in a theory Θ a channel of form

$$\begin{array}{c} \text{---} \mathcal{X}_1 \text{---} \\ \text{---} \mathcal{Z}_1 \text{---} \\ \text{---} \mathcal{Z}_1 \text{---} \\ \text{---} \mathcal{X}_2 \text{---} \end{array} \begin{array}{c} \boxed{T_1} \\ \boxed{T_2} \end{array} \begin{array}{c} \text{---} \mathcal{Y}_1 \text{---} \\ \text{---} \mathcal{Y}_2 \text{---} \end{array} , \quad (2.12)$$

as in our earlier discussion of the pictorial syntax. Let us call such a channel a (*bipartite*) *Bell-channel*, as it corresponds to the experimental set-up considered by J. Bell ([Bel64]): A state s is shared across two sites, and locally at each site a channel T_i can be applied to part of the state and a local input. The ground-breaking demonstration of [Bel64] (mentioned in the general introduction, and alluded to in connection with Example 1.1.10) was that the components T_1, T_2 and s can be chosen from **QIT** in such a way that the total channel (2.12) is interpretable in **CIT**, though no choice of T_1, T_2 and s from **CIT** will reproduce it. A key point in this argument is that it is quite easy to understand the Bell-channels in **CIT**, because any randomness from the channels T_i can be extracted and moved to s , thus effectively realising the channel (2.12) as a convex combination of products of deterministic channels (functions).

It is not clear that the Bell-channels in **QIT** are similarly easy to characterise, but one thing can be said, in fact about Bell-channels in any theory: They must be non-signalling from \mathcal{X}_1 to \mathcal{Y}_2 and from \mathcal{X}_2 to \mathcal{Y}_1 . Indeed, as demonstrated by our earlier computation,

$$\begin{array}{c} \text{---} \mathcal{X}_1 \text{---} \\ \text{---} \mathcal{X}_2 \text{---} \end{array} \begin{array}{c} \boxed{T_1} \\ \boxed{T_2} \end{array} \begin{array}{c} \text{---} \mathcal{Y}_1 \text{---} \\ \text{---} \mathcal{Y}_2 \text{---} \end{array} \boxed{\text{tr}} = \begin{array}{c} \text{---} \mathcal{X}_1 \text{---} \\ \text{---} \mathcal{X}_2 \text{---} \end{array} \begin{array}{c} \boxed{\text{tr}} \\ \boxed{T_2} \end{array} \begin{array}{c} \text{---} \mathcal{Y}_1 \text{---} \\ \text{---} \mathcal{Y}_2 \text{---} \end{array} = \begin{array}{c} \text{---} \mathcal{X}_1 \text{---} \\ \text{---} \mathcal{X}_2 \text{---} \end{array} \begin{array}{c} \boxed{\text{tr}} \\ \boxed{T_2} \end{array} \begin{array}{c} \text{---} \mathcal{Y}_1 \text{---} \\ \text{---} \mathcal{Y}_2 \text{---} \end{array} \boxed{\text{tr}} = \begin{array}{c} \text{---} \mathcal{X}_1 \text{---} \\ \text{---} \mathcal{X}_2 \text{---} \end{array} \begin{array}{c} \boxed{\text{tr}} \\ \boxed{T_2'} \end{array} \begin{array}{c} \text{---} \mathcal{Y}_1 \text{---} \\ \text{---} \mathcal{Y}_2 \text{---} \end{array} , \quad (2.13)$$

and similarly with \mathcal{Y}_2 trashed in place of \mathcal{Y}_1 .

For some time, it was unknown whether the class of Bell-channels in **QIT** (with classical inputs and outputs) was restrained only by these two non-signalling conditions. The work of Refs. [Cir80, PR94] settled this in the negative, and we now know that in fact the class of Bell-channels in **QIT** is highly complex ([BCP+14, GKW+18]).

◆

Example 2.1.15. (The PR Box.)

The concrete example given by S. Popescu and D. Rohrlich in Ref. [PR94] of a non-signalling channel in **CIT**, which cannot be realised as a Bell-channel in **QIT**, was the classical channel

$\begin{array}{c} \{0,1\} \\ \square \\ \{0,1\} \end{array} \begin{array}{c} \{0,1\} \\ P \\ \{0,1\} \end{array}$, determined by the probability distributions $(P^{x_1, x_2})_{x_1, x_2 \in \{0,1\}}$ given by

$$P^{x_1, x_2}(y_1, y_2) = \begin{cases} \frac{1}{2} & \text{for } y_1 \oplus y_2 = x_1 \cdot x_2 \\ 0 & \text{for } y_1 \oplus y_2 \neq x_1 \cdot x_2 \end{cases}, \quad (2.14)$$

with \oplus denoting addition modulo 2. It is easy to verify that each of the marginals equals the channel $\begin{array}{c} \{0,1\} \\ \square \\ \{0,1\} \end{array} \begin{array}{c} \{0,1\} \\ F \\ \{0,1\} \end{array}$, with F the channel given by $F(k) = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1$ for $k = 0, 1$, so in particular P is non-signalling. The reader is not expected to see why P cannot be realised as a Bell-channel, but this is the result of Refs. [Cir80, PR94]. The channel P has since become known in the folklore as the *Popescu-Rohrlich Box*, or simply *PR Box* (though it had in fact been identified earlier, e.g. by Cirelson himself ([Cir93])).

◆

Example 2.1.16. (Non-Signalling in **Logic**.)

Consider the thin theory of propositional logic, **Logic**, from Example 1.2.16. The channel $\begin{array}{c} P_0 \\ \square \\ P_1 \end{array} \begin{array}{c} P_0 \\ Q_0 \\ P_1 \\ Q_1 \end{array}$ asserts the relation $P_0 \wedge Q_0 \succeq P_1 \wedge Q_1$ (i.e. " $P_0 \wedge Q_0$ implies $P_1 \wedge Q_1$ "). To say that the channel is non-signalling from the P_0 -port to the Q_1 -port would be to say that the marginal relation $\begin{array}{c} P_0 \\ \square \\ P_1 \end{array} \begin{array}{c} P_0 \\ Q_0 \\ P_1 \\ \text{tr} \end{array}$ is of the form $\begin{array}{c} P_0 \\ \square \\ P_1 \end{array} \begin{array}{c} P_0 \\ \text{tr} \\ P_1 \end{array}$, or, what is actually equivalent, that " P_0 implies P_1 ". This of course does not generally follow from " $P_0 \wedge Q_0$ implies $P_1 \wedge Q_1$ ", so not all channels in the theory **Logic** are non-signalling.

◆

Non-signalling as such shall not occupy us before Chapter 4, but it is relevant in the context of dilations because, due to our choice of defining them as two-sided, dilations and non-signalling are intimately related – indeed, $T : \mathbb{X} \rightarrow \mathbb{Y}$ is non-signalling from \mathbb{X}_0 to \mathbb{Y}_0 if and only if it is the dilation of some channel $T' : \mathbb{X} \setminus \mathbb{X}_0 \rightarrow \mathbb{Y}_0$. This also means that any understanding we have of non-signalling channels can be transferred to an understanding about two-sided dilations. We end this section by discussing one such situation.

It is obvious that a channel of the form

$$\begin{array}{c} \mathbb{X} \setminus \mathbb{X}_0 \\ \square \\ \mathbb{X}_0 \end{array} \begin{array}{c} T_1 \\ \text{H} \\ T_2 \end{array} \begin{array}{c} \mathbb{Y}_0 \\ \mathbb{Y} \setminus \mathbb{Y}_0 \end{array} \quad (2.15)$$

is non-signalling from \mathbb{X}_0 to \mathbb{Y}_0 ; trashing $\mathbb{Y} \setminus \mathbb{Y}_0$ results in trashing \mathbb{X}_0 . More than 20 years ago, it was conjectured by DiVincenzo ([BGNP01, ESW02]) that in **QIT** channels of the form (2.15) are in fact the only non-signalling channels. This conjecture was proven in Ref. [BGNP01] in a special case, and later in Ref. [ESW02] for arbitrary channels. Phrased in terms of dilations, this means that in **QIT** every dilation of a channel $\begin{array}{c} \mathbb{X} \\ \square \\ \mathbb{Y} \end{array}$ must take the form

$$\begin{array}{c} \mathbb{X} \\ \square \\ \mathbb{Y} \end{array} \begin{array}{c} L_0 \\ \text{E}_0 \\ G \end{array} \begin{array}{c} \mathbb{Y} \\ \text{E} \end{array} \quad (2.16)$$

for some channels L_0 and G . Moreover, by trashing \mathbb{E} we easily read off that L_0 is in this case a one-sided dilation of T . Conversely, the channel (2.16) defines a dilation of T for any one-sided dilation L_0 and any channel G . **In other words, to determine the dilations of a channel in QIT, we need only determine the one-sided dilations.** (In particular, this immediately hands us all dilations of states in QIT, cf. Example 2.1.12.)

Definition 2.1.17. (The DiVincenzo Property.)

We say that a theory Θ has the *DiVincenzo property* if every channel $-X-T-Y-$ which is non-signalling from X_0 to Y_0 is of the form

$$\begin{array}{ccc}
 -X \setminus X_0- & \boxed{T_1} & -Y_0- \\
 & | & \\
 & \text{---} \mathbb{H} \text{---} & \\
 & | & \\
 - & \boxed{T_2} & -Y \setminus Y_0- \\
 & / & \\
 & X_0 &
 \end{array} \tag{2.17}$$

for some channels T_1 and T_2 . ■

Remark 2.1.18. When checking for the DiVincenzo Property, it obviously suffices to consider the case where X and Y are bipartite interfaces and X_0 and Y_0 are simple (i.e. single-port) sub-interfaces. ✦

Example 2.1.19. (Failure of the DiVincenzo Property.)

Consider the thin theory $(\mathbb{N}, \cdot, 1, \geq)$ from Example 1.2.13. The channel $-7- \boxed{\quad} -4-$ is non-signalling from the 2-port to the 4-port, since $7 \geq 4$. However, it cannot be written as

$$\begin{array}{ccc}
 -7- & \boxed{\quad} & -4- \\
 & | & \\
 & z & \\
 & | & \\
 - & \boxed{\quad} & -3- \\
 & / & \\
 & 2 &
 \end{array} \quad \text{for any } z \in \mathbb{N}, \text{ since } 7 \geq 4 \cdot z \text{ forces } z = 1, \text{ which violates } 2 \cdot z \geq 3.$$

◆

In the next section, we will see that the DiVincenzo property is very generic, and we will understand that it is a consequence of an information-theoretic principle governing dilations. As observed just above, this will greatly simplify the study of dilations.

We must begin, however, by giving shape to the intuition that some dilations can be *constructed* from other dilations, a phenomenon which we also observed in connection with Example 2.1.9, Example 2.1.11, and Example 2.1.12.

2.2 The Dilational Ordering

Dilations are not isolated islands – some of them are connected by bridges. Ideally, these bridges should connect general two-sided dilations to general two-sided dilations, but it turns out that we will not be successful in attempting such a definition before introducing the theory of *causal* channels in Chapter 4. For now, we must restrict our attention to the case where at least the initial dilation is one-sided:

Definition 2.2.1. (The Dilational Ordering.)

Let $-X-T-Y-$ be a channel in Θ , and denote by $\text{Dil}(T)$ the class of all its dilations. The *dilational ordering on $\text{Dil}(T)$* is the relation \supseteq_T on $\text{Dil}(T)$ given by

$$\begin{array}{c}
\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \end{array} \triangleright_T \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L'} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}'\sim \end{array} \Leftrightarrow \exists G : \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{E}\text{---} \\ \sim\mathbb{D}'\sim \end{array} \begin{array}{c} \text{---}\mathbb{Y}\text{---} \\ \boxed{G} \\ \text{---}\mathbb{E}'\sim \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L'} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}'\sim \end{array}
\end{array} \quad (2.18)$$

We say that L' is derivable from L if $L \triangleright_T L'$. ■

Remark 2.2.2. (On the Nature of the Relation \triangleright_T .)

The domain of the relation \triangleright_T is somewhat awkward (since one dilation may be two-sided and the other not). If we restrict both dilations to be one-sided, however, \triangleright_T restricts to a reflexive and transitive relation, i.e. a pre-order. The reason that we do not confine ourselves to one-sided dilations in the definition is that we soon wish to give meaning to the idea that every dilation (even if two-sided), is derivable from a single fixed dilation. ✂

Remark 2.2.3. (Notation – On the Dependence of \triangleright_T on T .)

We shall often drop the subscript from \triangleright_T and write simply \triangleright . Note, however, that it does matter what the base interfaces \mathbb{X} and \mathbb{Y} are; for example, dilations $L : \mathbb{X} \cup \mathbb{D} \rightarrow \mathbb{Y} \cup \mathbb{E}$ of T (with hidden interfaces \mathbb{D} and \mathbb{E}) can also be seen as dilations of $\text{tr}_{\mathbb{I}}$ (with hidden interfaces $\mathbb{X} \cup \mathbb{D}$ and $\mathbb{Y} \cup \mathbb{E}$), but then of course the dilational ordering trivialises. ✂

Example 2.2.4. (Large Dilations.)

Example 2.1.11 shows that for any state $\boxed{p}\text{---}\mathbb{X}\text{---}$ in **CIT**, the class of one-sided dilations of $\boxed{p}\text{---}\mathbb{X}\text{---}$ has a \triangleright -largest element, namely the dilation given by copying. Example 2.1.12 demonstrates that this is also the case for states in **QIT**, with a \triangleright -largest element given by purification. ◆

Example 2.2.5. (DiVincenzo and Derivability.)

In theories which have the DiVincenzo Property, any dilation L of T is derivable from a one-sided dilation L_0 . In fact, this statement is equivalent to the DiVincenzo Property. ◆

Example 2.2.6. (Dilational Ordering in Cartesian Theories.)

By Example 2.1.9, every dilation of a channel $\text{---}\mathbb{X}\text{---}\boxed{T}\text{---}\mathbb{Y}\text{---}$ is derivable from the dilation given by copying the input. ◆

Example 2.2.7. (Dilational Ordering in a Thin Theory.)

Consider the channel $\text{---}7\text{---}\boxed{}\text{---}3\text{---}$ in the thin theory $(\mathbb{N}, \cdot, 1, \geq)$. Every one-sided dilation of this channel is derivable from the one-sided dilation $\text{---}7\text{---}\boxed{}\text{---}3\text{---}$. Some dilations of $\text{---}7\text{---}\boxed{}\text{---}3\text{---}$, however, are not derivable from this dilation, for example the dilation $\begin{array}{c} \sim 3 \sim \\ \boxed{} \\ \sim 7 \sim \\ \text{---}7\text{---}\boxed{}\text{---}3\text{---} \end{array}$. ◆

Remark 2.2.8. (One-Sided Dilations as a Category.)

The class of one-sided dilations of $\text{---}\mathbb{X}\text{---}\boxed{T}\text{---}\mathbb{Y}\text{---}$ does not merely have the structure of a pre-order; in fact, the very definition of the dilational ordering reveals that the pre-order is the shadow of a category whose objects are the one-sided dilations of T and whose morphisms from $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{E}\sim \end{array}$ to $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L'} \\ \text{---}\mathbb{E}'\sim \end{array}$ are the channels $\sim\mathbb{E}\sim\boxed{G}\sim\mathbb{E}'\sim$ such that $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{E}\sim \end{array} \begin{array}{c} \text{---}\mathbb{Y}\text{---} \\ \boxed{G} \\ \text{---}\mathbb{E}'\sim \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L'} \\ \text{---}\mathbb{E}'\sim \end{array}$. We shall not have occasion to study this category in any detail. (Also, it seems to vanish in the elaboration of Chapter 4.) ✂

It turns out that the dilational ordering is rather well-behaved in the theories of interest to us, confined by simple principles with a clear interpretation. Before we study those principles in Section 2.3.A and Section 2.3.B, we will examine the possibility of a rather brutal collapse of the dilational ordering.

2.2.A Dilational Purity

Any channel $-X-[T]-Y-$ has some dilations, namely the *trivial dilations* of the form $-X-[T]-Y-$ $\sim_{\mathbb{D}} \sim_{\mathbb{D}} [S] \sim_{\mathbb{E}} \sim_{\mathbb{D}}$.

In the dilational ordering, the trivial dilations are precisely those that can be derived from the (very trivial) dilation $-X-[T]-Y-$ itself. It may happen that there are no other dilations:

Definition 2.2.9. (Dilational Purity.)

A channel $-X-[T]-Y-$ is called *dilationally pure* if every dilation of T is of the form $-X-[T]-Y-$ $\sim_{\mathbb{D}} \sim_{\mathbb{D}} [S] \sim_{\mathbb{E}} \sim_{\mathbb{D}}$ for some channel S . ■

From an operational point of view, dilational purity of T signifies that any side-computation in the environment of T must be independent from it. We will see in the next section that all isometric channels in **QIT** are dilationally pure, and since identity channels are isometric this will immediately imply a rather strong version of the *No Broadcasting Theorem* ([BCF⁺96]).

Example 2.2.10. (Dilational Purity in Thin Theories.)

In a thin theory, consider the identity channel $-x-\square-x-$. Its dilations are $\sim_y \sim_x \square \sim_x \sim_z$, with $x \star y \succeq x \star z$. If the monoid satisfies the cancellation law $x \star y \succeq x \star z \Leftarrow y \succeq z$, then these dilations factor as $\sim_y \sim_x \square \sim_x \sim_z$, so the identity $-x-\square-x-$ is dilationally pure. ◆

In **CIT** we have the following characterisation of the dilationally pure channels:

Proposition 2.2.11. (Dilationally Pure Channels in **CIT**.)

A channel in **CIT** is dilationally pure if and only if it is a probabilistically pure state.

Proof. According to Example 2.1.11, a probabilistically pure state has only trivial one-sided dilations. The fact that a two-sided dilation must also be trivial can either be verified by a direct consideration or seen as a consequence of the DiVincenzo property for **CIT**, which we shall prove in Section 2.3.B.

Assume conversely that $-X-[T]-Y-$ is dilationally pure. Consider the copy channel $-X-\square-$ from earlier. Clearly, $-X-\square-[T]-Y-$ is a dilation of T , so by dilational purity it must be of the form $-X-[T]-Y-$ $\sim_{\mathbb{D}} \sim_{\mathbb{D}} [s] \sim_{\mathbb{E}} \sim_{\mathbb{D}}$ for some state s . However, we now have

$$-X-\square-\square-\sim_X \sim_X = -X-\square-[T]-\text{tr} = -X-[T]-\text{tr} [s] \sim_X \sim_X, \quad (2.19)$$

which is only possible if the system X corresponding to X has size $|X| = 1$, so T must be a state.⁶ It is then an easy exercise to verify that a state which is dilationally pure must be probabilistically pure; if it is not, then we obtain a non-trivial dilation by copying the output. □

⁶Ignoring that we strictly speaking defined states as having the particular domain **1**.

In the next section we shall be able to also give a complete characterisation of the dilationally pure channels in **QIT**. For now, we content ourselves with a confinement:

Lemma 2.2.12. (*Necessary Condition for Dilational Purity in QIT.*)
If a channel in QIT is dilationally pure, then it is isometric.

Proof. It suffices to consider channels with simple input and output interfaces. Assume that $\text{-}\mathcal{X}\text{-}\boxed{\Lambda}\text{-}\mathcal{Y}\text{-}$ is a dilationally pure quantum channel, and let $(K_i)_{i \in I}$ be a Kraus representation of Λ , i.e. a family of linear operators $K_i : \mathcal{X} \rightarrow \mathcal{Y}$ such that $\Lambda(A) = \sum_{i \in I} K_i A K_i^*$ for all $A \in \text{End}(\mathcal{X})$. The channel $\text{-}\mathcal{X}\text{-}\boxed{\Phi}\text{-}\mathcal{Y}\text{-}$ given by $\Phi(A) = \sum_{i \in I} K_i A K_i^* \otimes |i\rangle\langle i|$ is then a dilation of $\text{-}\mathcal{X}\text{-}\boxed{\Lambda}\text{-}\mathcal{Y}\text{-}$ with hidden system \mathbb{C}^I , so by dilational purity we must have $\Phi(A) = \sum_{i \in I} K_i A K_i^* \otimes \varrho$ for some state ϱ on \mathbb{C}^I . By equality of these two expressions we conclude that $K_{i_0} A K_{i_0}^* = \langle i_0 | \varrho | i_0 \rangle \sum_{i \in I} K_i A K_i^*$ for any $i_0 \in I$, and since $1 = \text{tr}(\varrho) = \sum_{i \in I} \langle i | \varrho | i \rangle$ we can pick some $i_0 \in I$ with $\langle i_0 | \varrho | i_0 \rangle > 0$, for which we consequently have

$$\Lambda(A) = \sum_{i \in I} K_i A K_i^* = \frac{1}{\langle i_0 | \varrho | i_0 \rangle} K_{i_0} A K_{i_0}^* \quad (2.20)$$

for all $A \in \text{End}(\mathcal{X})$. In other words, Λ admits a Kraus representation using the single Kraus operator $K_{i_0} / \sqrt{\langle i_0 | \varrho | i_0 \rangle}$. This operator must then be an isometry, so Λ is isometric. \square

Remark 2.2.13. (Relation of Dilational Purity to other Purity Notions.)

Whereas dilational purity is a well-defined concept in all theories that we consider, it does not generally make sense to speak of probabilistic purity as it requires a convex structure. Proposition 2.2.11 shows that in theories where the concept does make sense, it may be distinct from dilational purity: Any deterministic function whose domain has at least two elements is probabilistically pure in **CIT**, but not dilationally pure according to the proposition. (On the other hand, though we will not have occasion to be precise about this, it is easy to see that dilational purity implies probabilistic purity in theories where that concept does make sense, since to a non-trivial convex decomposition $\frac{1}{2}T_0 + \frac{1}{2}T_1$ we can associate the non-trivial dilation $\frac{1}{2}T_0 \otimes \delta_0 + \frac{1}{2}T_1 \otimes \delta_1$ which keeps as side-information a memory of which component was employed.)

The fact that these two purity notions do not coincide, also implies that dilational purity is distinct from the purity notion of Ref. [CH17], defined in terms of *weak factorisation systems*, which reduces in **CIT** to probabilistic purity.

If we restrict the condition in Definition 2.2.9 to one-sided dilations then we obtain the purity notion proposed in Ref. [Chi14b], but in general there is in fact a distinction between one- and two-sided purity.⁷ \spadesuit

2.3 Completeness and Localisability

In this section, the *completeness* and *localisability* principles are introduced, and we prove that they hold in many theories, in particular the two information theories **CIT** and **QIT**.

⁷For example, in the theory **T – REX** consisting precisely of the surjective functions between finite sets (Example 1.2.23), any one-sided dilation of a bijection is trivial, though a bijection generally has non-trivial two-sided dilations in **T – REX**.

2.3.A Complete Dilations

Dilational purity of a channel is an extreme condition, under which the dilational ordering implodes to a single level. In interesting theories, most channels will not be dilationally pure (cf. Proposition 2.2.11 and Lemma 2.2.12). They do, however, comply to another principle which is milder and still retains a remarkable simplicity:

Definition 2.3.1. (Complete Dilations.)

Let $-X\text{-}[T]\text{-}Y\text{-}$ be a channel in Θ , and let $\mathbf{D} \subseteq \text{Dil}(T)$ be a class of dilations of T . We say that a (one-sided) dilation $K \in \mathbf{D}$ is *complete for \mathbf{D}* if $K \triangleright L$ for all $L \in \mathbf{D}$. A dilation $K \in \text{Dil}(T)$ is called simply *complete* if it is complete for $\text{Dil}(T)$. ■

Definition 2.3.2. (Complete Theories.)

A theory Θ is called *complete* if every channel in Θ has a complete dilation. ■

Example 2.3.3. (Completeness and Dilational Purity.)

If $-X\text{-}[T]\text{-}Y\text{-}$ is a dilationally pure channel in Θ , then $-X\text{-}[T]\text{-}Y\text{-}$ is a complete dilation of itself; in fact, *self-completeness* is equivalent to dilational purity. ♦

Example 2.3.4. (Trivial Completeness.)

In any theory Θ , the trash $-Z\text{-}[\text{tr}]\text{-}$ has a complete dilation, namely $-Z\text{-}[\text{id}]\text{-}Z\text{-}$. ♦

Example 2.3.5. (Completeness in Thin Theories.)

The fact that a potential complete dilation is required to be one-sided means that every complete theory must have the DiVincenzo Property (Definition 2.1.17). Hence, it follows from Example 2.1.19 that the thin theory $(\mathbb{N}, \cdot, 1, \geq)$ is not complete. Some of its channels do have complete dilations, however. In fact, it is not difficult to see that the channel $-x\text{-}[\]\text{-}y\text{-}$ has a complete dilation if and only if x is divisible by y , in which case the dilation $-x\text{-}[\]\text{-}y\text{-}$ is complete. ♦

Though completeness might seem like a foolish mathematical fantasy rarely fulfilled in theories of real interest, the truth is quite the opposite. For example, we have already seen an argument to the effect that all cartesian theories are complete:

Theorem 2.3.6. (Cartesian Theories are Complete.)

Every cartesian theory Θ is complete. In fact, if $-X\text{-}[T]\text{-}Y\text{-}$ is a channel in Θ then the dilation given by copying the inputs,

$$-X\text{-}[\]\text{-}Y\text{-} \quad (2.21)$$

is complete.

Proof. It suffices to show this in the case where the interfaces X and Y are simple, but this is precisely what we did in Example 2.1.9. □

It is also the case that **CIT** and **QIT** are complete, but the most elegant way of reaching this conclusion is to start by proving completeness relative to one-sided dilations, and then lift those results in the next subsection.

To ease language, let us say that a dilation of $-X\text{-}[T]\text{-}Y\text{-}$ is *one-sided-complete* if it is complete for one-sided dilations. Let us also say that a theory Θ is *one-sided-complete* if every channel $-X\text{-}[T]\text{-}Y\text{-}$ in Θ has a one-sided-complete dilation.

Lemma 2.3.7. *(One-Sided-Completeness in CIT.)*

The theory **CIT** is one-sided-complete. Specifically, for a channel $-X-T-Y-$, the dilation given by copying both the inputs and outputs,

$$-X-\textcirclearrowleft\text{---}\boxed{T}\text{---}\textcirclearrowright-Y-, \quad (2.22)$$

is one-sided-complete.

Proof. In the case where X is trivial, this statement is exactly what was proved in Example 2.1.11: The copy-dilation $\boxed{p}\text{---}\textcirclearrowleft\text{---}\textcirclearrowright-Y-$ is complete for one-sided dilations $\boxed{p}\text{---}Y-$, by virtue of existence of conditional probabilities. A general classical channel $-X-T-Y-$ is simply a collection of states indexed by the input, so any one-sided dilation $-X-\boxed{L}\text{---}\textcirclearrowright-Y-$ can be derived from the channel (2.22), since knowing a copy of the input x essentially reduces the problem to the case of a single state (the details are left as an exercise). \square

Lemma 2.3.8. *(One-Sided-Completeness in QIT.)*

The theory **QIT** is one-sided-complete. Specifically, for a channel $-X-\Lambda-Y-$, any Stinespring dilation (i.e. any isometric dilation)

$$-X-\boxed{\Sigma}\text{---}\textcirclearrowright\mathcal{E}\text{---}\textcirclearrowright-Y- \quad (2.23)$$

of Λ is one-sided-complete.

Proof. Again, we proved this in the case where Λ is a state in Example 2.1.12: Purifications are special cases of Stinespring dilations, and by the uniqueness of purifications it was demonstrated that any purification $\boxed{\pi}\text{---}\textcirclearrowright\mathcal{E}\text{---}\textcirclearrowright-Y-$ of a state $\boxed{\varrho}\text{---}Y-$ is a complete dilation. For general channels, Stinespring's Dilation Theorem ([Sti55, NC02, Wat]), which is covered in the preliminary section of the thesis, asserts that every quantum channel has an isometric one-sided dilation and that this dilation is unique up to a channel acting on the hidden interface. By an argument analogous to that of Example 2.1.12, this implies that any one-sided dilation $-X-\boxed{\Phi}\text{---}\textcirclearrowright\mathcal{E}\text{---}\textcirclearrowright-Y-$ of $-X-\Lambda-Y-$ can be derived from a Stinespring dilation. \square

To conclude that the above dilations in **CIT** and **QIT** are not only one-sided-complete but in fact complete for all dilations, we will make use of the following observation:

Lemma 2.3.9. *(One-Sided-Completeness and DiVincenzo.)*

Let Θ be a theory. The following are equivalent:

1. Θ is complete.
2. Θ has the DiVincenzo property and Θ is one-sided-complete.

In this case, any dilation of $-X-T-Y-$ which is complete for its one-sided dilations is in fact complete for all dilations.

Proof. We have already observed that completeness implies the DiVincenzo property, so 2. clearly follows from 1. Conversely, assuming 2., if $\begin{array}{c} \boxed{K} \\ \sim \mathbb{E} \sim \\ \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \end{array}$ is complete for one-sided dilations of $\begin{array}{c} \boxed{T} \\ \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \end{array}$ then, by the DiVincenzo property, it is in fact complete for all dilations, so 1. follows, and so does the final statement in the lemma. \square

What remains in order to prove completeness is then only an argument to the effect that the DiVincenzo property holds in **CIT** and **QIT**. This will be accomplished by relating the property to the principle of *spatial localisability* which is introduced in the following subsection.

2.3.B Spatial and Temporal Localisability

Completeness is a *static* principle, in the sense that it pertains to the dilational ordering for a fixed channel T . It is natural to ask how dilations behave under the *dynamic* structure inherent in every theory: Parallel and serial composition.

Consider for example two channels $\begin{array}{c} \boxed{T_1} \\ \text{---} \mathbb{X}_1 \text{---} \boxed{T_1} \text{---} \mathbb{Y}_1 \text{---} \end{array}$ and $\begin{array}{c} \boxed{T_2} \\ \text{---} \mathbb{X}_2 \text{---} \boxed{T_2} \text{---} \mathbb{Y}_2 \text{---} \end{array}$. Clearly, for any one-sided dilations $\begin{array}{c} \boxed{L_1} \\ \sim \mathbb{E}_1 \sim \\ \text{---} \mathbb{X}_1 \text{---} \boxed{L_1} \text{---} \mathbb{Y}_1 \text{---} \end{array}$ and $\begin{array}{c} \boxed{L_2} \\ \sim \mathbb{E}_2 \sim \\ \text{---} \mathbb{X}_2 \text{---} \boxed{L_2} \text{---} \mathbb{Y}_2 \text{---} \end{array}$ of T_1 and T_2 respectively, the chan-

nel $\begin{array}{c} \boxed{L_1} \\ \sim \mathbb{E}_1 \sim \\ \text{---} \mathbb{X}_1 \text{---} \boxed{L_1} \text{---} \mathbb{Y}_1 \text{---} \\ \boxed{L_2} \\ \sim \mathbb{E}_2 \sim \\ \text{---} \mathbb{X}_2 \text{---} \boxed{L_2} \text{---} \mathbb{Y}_2 \text{---} \end{array}$ is a dilation of their parallel composition $\begin{array}{c} \boxed{T_1} \\ \text{---} \mathbb{X}_1 \text{---} \boxed{T_1} \text{---} \mathbb{Y}_1 \text{---} \\ \boxed{T_2} \\ \text{---} \mathbb{X}_2 \text{---} \boxed{T_2} \text{---} \mathbb{Y}_2 \text{---} \end{array}$. In general,

the parallel composition has other dilations than those of this form, indeed all the dilations derivable from dilations of this form. Such derivable dilations will be called *spatially localisable* since, intuitively, the side-information corresponding to them can be ‘localised’ as a combination of side-information from either T_1 or T_2 :

Definition 2.3.10. (Spatial Localisability.)

Let $\begin{array}{c} \boxed{T_1} \\ \text{---} \mathbb{X}_1 \text{---} \boxed{T_1} \text{---} \mathbb{Y}_1 \text{---} \end{array}$ and $\begin{array}{c} \boxed{T_2} \\ \text{---} \mathbb{X}_2 \text{---} \boxed{T_2} \text{---} \mathbb{Y}_2 \text{---} \end{array}$ be parallelly composable channels in Θ . We say that

a dilation of the parallel composition $\begin{array}{c} \boxed{T_1} \\ \text{---} \mathbb{X}_1 \text{---} \boxed{T_1} \text{---} \mathbb{Y}_1 \text{---} \\ \boxed{T_2} \\ \text{---} \mathbb{X}_2 \text{---} \boxed{T_2} \text{---} \mathbb{Y}_2 \text{---} \end{array}$ is *spatially localisable w.r.t. T_1 and T_2*

if it is of the form


$$\begin{array}{c} \text{---} \mathbb{X}_1 \text{---} \boxed{L_1} \text{---} \mathbb{Y}_1 \text{---} \\ \sim \mathbb{D} \sim \\ \text{---} \mathbb{X}_2 \text{---} \boxed{L_2} \text{---} \mathbb{Y}_2 \text{---} \\ \boxed{G} \\ \sim \mathbb{E} \sim \end{array} \quad (2.24)$$

for some channel G and some dilations L_1 of T_1 and L_2 of T_2 . \blacksquare

Obviously, there is a natural counterpart of localisability for serial compositions, which we name *temporal localisability*:

Definition 2.3.11. (Temporal Localisability.)

Let $\begin{array}{c} \boxed{T} \\ \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \end{array}$ and $\begin{array}{c} \boxed{S} \\ \text{---} \mathbb{Y} \text{---} \boxed{S} \text{---} \mathbb{Z} \text{---} \end{array}$ be channels in Θ . We say that a dilation of the serial composition $\begin{array}{c} \boxed{T} \\ \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \boxed{S} \text{---} \mathbb{Z} \text{---} \end{array}$ is *temporally localisable w.r.t. T and S* if it is of the form



for some channel G and some dilations L of T and M of S . ■

By a somewhat awful abuse of language, we shall moreover use the following terminology:

Definition 2.3.12. (Localisability of a Theory.)

A theory Θ is called *spatially localisable* if every one-sided dilation of a parallel composition is spatially localisable. A theory Θ is called *temporally localisable* if every one-sided dilation of a serial composition is temporally localisable. We say that Θ is *localisable* if it is both spatially and temporally localisable. ■

Remark 2.3.13. Observe that localisability is imposed on one-sided dilations only. This provides a simpler interpretation of the requirement, and also makes the conditions easier to fulfil. (It will soon be clear, however, that the conditions automatically follow for two-sided dilations if they hold for the one-sided.) ✂

It might seem a hopeless task to prove that a theory is spatially or temporally localisable. To our luck, however, we are mainly interested in theories which are one-sided-complete, and for those we have the following:

Lemma 2.3.14. (Recharacterisation of Localisability.)

Suppose that Θ is one-sided-complete. Then,

- Θ is spatially localisable if and only if for any channels $-X_1-[T_1]-Y_1-$ and $-X_2-[T_2]-Y_2-$,

there exist one-sided-complete dilations K_1 of T_1 and K_2 of T_2 , such that
is a one-sided-complete dilation of $-X_1-[T_1]-Y_1-$ and $-X_2-[T_2]-Y_2-$.
- Θ is temporally localisable if and only if for any channels $-X-[T]-Y-$ and $-Y-[S]-Z-$,

there exist one-sided-complete dilations K of T and C of S , such that
is a one-sided-complete dilation of $-X-[T]-Y-[S]-Z-$.

Proof. Obvious. □

Remark 2.3.15. It is easy to see that if Θ is spatially [temporally] localisable, then in fact the parallel [serial] composition of any complete dilations yields a complete dilation. ✂

As if tailor-made, this lemma immediately implies that our pet theories are localisable:

Theorem 2.3.16. *Every cartesian theory is localisable.*

Proof. By Theorem 2.3.6, a cartesian theory is complete, in particular one-sided-complete, so we can use Lemma 2.3.14. However, it is obvious that the parallel compositions of the dilations (2.21) obtained by copying the inputs from $T_1 : \mathbb{X}_1 \rightarrow \mathbb{Y}_1$ and $T_2 : \mathbb{X}_2 \rightarrow \mathbb{Y}_2$ yields the dilation obtained by copying the input from $T_1 \parallel T_2$, and it is also clear that the serial composition of the dilations obtained by copying from $T : \mathbb{X} \rightarrow \mathbb{Y}$ and $S : \mathbb{Y} \rightarrow \mathbb{Z}$ yields (a dilation from which we can derive) the dilation obtained by copying from $S \circ T$. The desired follows. \square

Theorem 2.3.17. *The theory CIT is localisable.*

Proof. By Lemma 2.3.7, CIT is one-sided-complete, so again Lemma 2.3.14 applies. The rest of the argument is exactly as in the previous proof, only now we have to copy the outputs as well as the inputs, cf. Eq. (2.22). \square

Theorem 2.3.18. *The theory QIT is localisable.*

Proof. By Lemma 2.3.8 QIT is one-sided complete, so we can use Lemma 2.3.14 a third time. The desired follows immediately, since the parallel and serial compositions of isometric channels are obviously isometric. \square

The status of our dilational investigations is at this point as follows: Cartesian theories are complete and localisable; CIT and QIT are one-sided-complete and localisable; thin theories need not be complete.

What we shall prove now is that temporal localisability implies the existence of reversible dilations and that spatial localisability implies the DiVincenzo property. These two consequences will be used repeatedly and they also provide us with counterexamples to localisability. Moreover, as one-sided-completeness and completeness are by Lemma 2.3.9 equivalent under the DiVincenzo property, full completeness of CIT and QIT will follow.

Proposition 2.3.19. *(Temporal Localisability implies Reversible Dilations.)*

If Θ is temporally localisable, then every channel $-\mathbb{X}-\boxed{T}-\mathbb{Y}-$ has a one-sided dilation

$$-\mathbb{X}-\boxed{R}-\mathbb{Y}- \quad \text{which is reversible.}$$

Proof. By Lemma 1.1.15,

$$-\mathbb{X}-\boxed{T}-\mathbb{Y}-\boxed{\text{tr}_{\mathbb{Y}}} = -\mathbb{X}-\boxed{\text{tr}_{\mathbb{X}}}. \quad (2.26)$$

Since $-\mathbb{X}-\boxed{\text{id}_{\mathbb{X}}}-\mathbb{X}-$ is a one-sided dilation of $\text{tr}_{\mathbb{X}}$, we must by temporal localisability find one-sided dilations R of T and M of $\text{tr}_{\mathbb{X}}$ such that

$$-\mathbb{X}-\boxed{R}-\mathbb{Y}-\boxed{M}-\mathbb{X}- \quad \text{with channel } G \quad = -\mathbb{X}-\boxed{\text{id}_{\mathbb{X}}}-\mathbb{X}- \quad (2.27)$$

for some channel G . This identity implies in particular that R is reversible. \square

Proposition 2.3.20. (*Spatial Localisability implies the DiVincenzo Property.*)

If Θ is spatially localisable, then Θ has the DiVincenzo Property. That is, every channel $-X-\boxed{T}-Y-$ which is non-signalling from X_0 to Y_0 is of the form

$$\begin{array}{c} -X \setminus X_0 - \boxed{T_1} - Y_0 - \\ | \\ -X_0 - \boxed{T_2} - Y \setminus Y_0 - \end{array} \quad (2.28)$$

for some channels T_1 and T_2 , or, equivalently, every dilation of a channel is derivable from a one-sided dilation.

Proof. Let $-X-\boxed{T}-Y-$ be non-signalling from X_0 to Y_0 . This means that there exists a channel T' such that

$$\begin{array}{c} -X \setminus X_0 - \boxed{T} - Y_0 - \\ | \\ -X_0 - \boxed{\text{tr}} - Y \setminus Y_0 - \end{array} = \begin{array}{c} -X \setminus X_0 - \boxed{T'} - Y_0 - \\ | \\ -X_0 - \boxed{\text{tr}} - Y \setminus Y_0 - \end{array} . \quad (2.29)$$

This, however, is to say that $\begin{array}{c} -X \setminus X_0 - \boxed{T} - Y_0 - \\ | \\ -X_0 - \boxed{\text{tr}} - Y \setminus Y_0 - \end{array}$ is a one-sided dilation of the parallel composition $\begin{array}{c} -X \setminus X_0 - \boxed{T'} - Y_0 - \\ | \\ -X_0 - \boxed{\text{tr}} - Y \setminus Y_0 - \end{array}$. By spatial localisability, we thus find dilations $\begin{array}{c} -X \setminus X_0 - \boxed{T_1} - Y_0 - \\ | \\ -X_0 - \boxed{M} - Y \setminus Y_0 - \end{array}$ of T' and $\begin{array}{c} -X_0 - \boxed{M} - E_2 - \\ | \\ -X_0 - \end{array}$ of tr_{X_0} , such that

$$\begin{array}{c} -X \setminus X_0 - \boxed{T} - Y_0 - \\ | \\ -X_0 - \boxed{\text{tr}} - Y \setminus Y_0 - \end{array} = \begin{array}{c} -X \setminus X_0 - \boxed{T_1} - Y_0 - \\ | \\ -X_0 - \boxed{M} - E_2 - \boxed{G} - Y \setminus Y_0 - \end{array} \quad (2.30)$$

for some channel G . Merging M and G to form T_2 we obtain Eq. (2.28) (with $H = E_1$) as desired. \square

As anticipated, we may now conclude full completeness of the information theories:

Theorem 2.3.21. (*Completeness of CIT.*)

The theory **CIT** is complete. Specifically, for a channel $-X-\boxed{T}-Y-$, the dilation given by copying both the inputs and outputs,

$$\begin{array}{c} \text{---} \\ | \\ -X - \boxed{T} - Y - \\ | \\ \text{---} \end{array}, \quad (2.31)$$

is complete.

Theorem 2.3.22. (*Completeness of QIT.*)

The theory **QIT** is complete. Specifically, for a channel $-X-\boxed{\Lambda}-Y-$, any Stinespring dilation (i.e. any isometric dilation)

$$\begin{array}{c} \text{---} \\ | \\ -X - \boxed{\Sigma} - Y - \\ | \\ \text{---} \end{array} \quad (2.32)$$

is complete.

Proof. In both cases, the proof is by one-sided-completeness (Lemma 2.3.7, Lemma 2.3.8) and spatial localisability (Theorem 2.3.17, Theorem 2.3.18), using Proposition 2.3.20 in conjunction with Lemma 2.3.9. \square

This also allows us to finally tie another loose end:

Corollary 2.3.23. (*Dilationally Pure Channels in QIT.*)
A channel in QIT is dilationally pure if and only if it is isometric.

Proof. By Lemma 2.2.12 any dilationally pure channel must be isometric, so it remains only to prove the converse. However, any isometric channel is by Theorem 2.3.22 a complete dilation of itself and thus dilationally pure. \square

Interestingly, this characterisation of the pure channels in **QIT** immediately implies a rather strong form of the No Broadcasting Theorem ([BCF+96]): Any dilation $\begin{array}{c} \sim \mathcal{E} \sim \\ \boxed{L} \\ \sim \mathcal{H} \sim \end{array}$ of an identity $\begin{array}{c} \sim \mathcal{H} \sim \\ \boxed{\text{id}} \\ \sim \mathcal{H} \sim \end{array}$ is trivial, so in particular its \mathcal{E} -marginal must trash the input \mathcal{H} .

We shall return to such considerations in more detail in Section 2.5, but for now we end this section by recalling that we have ultimately proved completeness and localisability of **CIT**, **QIT** and all cartesian theories. We have also seen that a thin theory need not be complete, and using Proposition 2.3.20 it need not be spatially localisable either, cf. Example 2.1.19 (it is also easy to find a direct example e.g. of a dilation of $\begin{array}{c} -3-\boxed{}-2- \\ -3-\boxed{}-2- \end{array}$ which cannot be derived from a parallel composition of dilations). By Proposition 2.3.19 we can similarly see that a thin theory need not be temporally localisable. The theory **T – REX** also displays lack of localisability.

2.4 Universal Dilations

Completeness and localisability give us leashes on the collections of dilations, but do not reveal what the pre-order \geq actually looks like among the one-sided dilations. This can be resolved if we slightly strengthen the notion of a complete dilation, introducing the idea of *universal dilations*.

All of the complete theories we have seen so far turn out to admit universal dilations, but this has to be proved in each case, and is actually somewhat tricky for the theory **QIT**. In fact, the result that **QIT** has universal dilations will constitute a generalisation of the injectivity of the Choi-Jamiołkowski isomorphism ([JLF13, dP67, Cho75, Jam72]) which apparently has not been observed before.

Besides providing a firm grip on the dilational order, the existence of universal dilations implies a forceful *contraction property* which will be instrumental in Chapter 4 (in the form of Lemma 4.2.5).

Definition 2.4.1. (Universal Dilations.)

Let $\begin{array}{c} -\mathbb{X}-\boxed{T}-\mathbb{Y}- \end{array}$ be a channel in a theory Θ . A one-sided dilation $\begin{array}{c} -\mathbb{X}-\boxed{U}-\mathbb{Y}- \\ \sim \mathbb{E}_0 \sim \end{array}$ is called a *universal dilation of T* if every dilation of T is of the form

$$\begin{array}{c}
\text{---}X\text{---} \\
\boxed{U} \\
\text{---}Y\text{---} \\
\text{---}E_0\text{---} \\
\boxed{G} \\
\text{---}E\text{---} \\
\text{---}D\text{---}
\end{array}
\tag{2.33}$$

for a unique channel G . ■

Definition 2.4.2. (Universal Theories.)

A theory Θ is called *universal* if every channel in Θ has a universal dilation. ■

Example 2.4.3. (Trivial Universality.)

In Example 2.3.4, we observed that $\text{---}Z\text{---}\boxed{\text{id}}\text{---}Z\text{---}$ is a complete dilation of the trash $\text{---}Z\text{---}\boxed{\text{tr}}\text{---}$ in any theory. In fact, the dilation is obviously universal. ◆

Example 2.4.4. (Universality and Dilational Purity.)

We observed in Example 2.3.3 that a dilationally pure channel $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ is *self-complete*, i.e. a complete dilation of itself. In fact, it is even *self-universal*, since by the normality requirement (Definition 1.3.7), the channel that derives a given trivial dilation is unique. ◆

We hasten to observe that cartesian theories and **CIT** are universal:

Theorem 2.4.5. *In a cartesian theory, every channel has a universal dilation.*

Proof. For any channel $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$, the complete dilation $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ is in fact universal,

since if $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ $\text{---}G\text{---}$ = $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ $\text{---}G'\text{---}$ then it follows by trashing Y that $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ $\text{---}G\text{---}$ = $\text{---}X\text{---}\boxed{\text{tr}}\text{---}$ $\text{---}G'\text{---}$ which by normality implies $G = G'$. □

Theorem 2.4.6. *In CIT, every channel has a universal dilation.*

Proof. Consider first a state $\boxed{p}\text{---}Y\text{---}$. The complete dilation $\boxed{p}\text{---}Y\text{---}$ need not be complete, since, as discussed in Example 2.1.11, the conditional distribution g_x is only uniquely determined when $p(y) > 0$. However, this also means that to obtain a universal dilation, all we need to do is cut down the hidden system to the support of p , $\text{supp}(p)$. For a channel $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ given by the distributions $(t_x)_{x \in X}$, we may similarly turn the complete dilation $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ into a universal one by cutting down the hidden system $X \times Y$ to the subset $\{(x, y) \in X \times Y \mid y \in \text{supp}(t_x)\}$. The details are left as exercise. □

Before proving that also the theory **QIT** is universal, let us prove two general results about universal dilations and universal theories.

The first result classifies all universal and complete dilations in terms of a single universal dilation:

Proposition 2.4.7. *(One determines All.)*

Let $\begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \end{array}$ be a universal dilation of $\begin{array}{c} \text{---X---} \\ \boxed{T} \\ \text{---Y---} \end{array}$.

1. For any isomorphism α , $\begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \\ \boxed{\alpha} \end{array}$ is a universal dilation of T , and every universal dilation of T is of this form for a unique isomorphism α .
2. For any reversible R , $\begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \\ \boxed{R} \end{array}$ is a complete dilation of T , and every complete dilation of T is of this form for a unique reversible R .

Proof. It is easy to see that the given channels are universal, respectively complete, for any isomorphism α , respectively reversible R . It thus suffices to prove the stated existence and uniqueness clauses. This is a standard ‘universal property in category’-argument.

The key observation is that universal dilations V have the property that if $\begin{array}{c} \text{---X---} \\ \boxed{V} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \end{array} = \begin{array}{c} \text{---X---} \\ \boxed{V} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \\ \boxed{H} \\ \sim_{\mathcal{E}_0} \end{array}$, then $H = \text{id}_{\mathcal{E}_0}$; this is simply by the uniqueness clause in the definition of universality. We can use this observation as follows:

As for item 1., let $\begin{array}{c} \text{---X---} \\ \boxed{U'} \\ \text{---Y---} \\ \sim_{\mathcal{E}'_0} \end{array}$ be any universal dilation of T . By universality of U , we must have

$$\begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \\ \boxed{G} \end{array} = \begin{array}{c} \text{---X---} \\ \boxed{U'} \\ \text{---Y---} \\ \sim_{\mathcal{E}'_0} \end{array} \quad (2.34)$$

for a unique channel G . Since however U' is also universal, we similarly find G' such that

$$\begin{array}{c} \text{---X---} \\ \boxed{U'} \\ \text{---Y---} \\ \sim_{\mathcal{E}'_0} \\ \boxed{G'} \end{array} = \begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \end{array} . \quad (2.35)$$

Together these two imply

$$\begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \\ \boxed{G} \end{array} \begin{array}{c} \text{---Y---} \\ \sim_{\mathcal{E}'_0} \\ \boxed{G'} \end{array} = \begin{array}{c} \text{---X---} \\ \boxed{U} \\ \text{---Y---} \\ \sim_{\mathcal{E}_0} \end{array} \quad \text{and} \quad \begin{array}{c} \text{---X---} \\ \boxed{U'} \\ \text{---Y---} \\ \sim_{\mathcal{E}'_0} \\ \boxed{G'} \end{array} \begin{array}{c} \text{---Y---} \\ \sim_{\mathcal{E}_0} \\ \boxed{G} \end{array} = \begin{array}{c} \text{---X---} \\ \boxed{U'} \\ \text{---Y---} \\ \sim_{\mathcal{E}'_0} \end{array} , \quad (2.36)$$

from which we by the initial observation conclude that $G' \circ G = \text{id}_{\mathcal{E}_0}$ and $G \circ G' = \text{id}_{\mathcal{E}'_0}$, i.e. the channel $G =: \alpha$ is an isomorphism with inverse G' .

Item 2. is proved analogously. \square

The second result provides a recharacterisation of the dilational ordering in the presence of universal dilations.

Let us denote by $\text{Dil}_0(T)$ the class of one-sided dilations of T with simple hidden interface (any one-sided dilation is \triangleright -equivalent to such a dilation, simply by merging systems in the environment). We have the following:

Proposition 2.4.8. *(Structure of $(\text{Dil}_0(T), \triangleright)$.)*

Suppose that $\begin{array}{c} \text{---X---} \\ \boxed{T} \\ \text{---Y---} \end{array}$ has a universal dilation for which the hidden system is \mathcal{E}_0 . Then, the pre-order $(\text{Dil}_0(T), \triangleright)$ is order-isomorphic to $(\text{Trans}_{\Theta}(\mathcal{E}_0, -), \succeq)$, where $\text{Trans}_{\Theta}(\mathcal{E}_0, -)$

denotes the class of all transformations in Θ with domain \mathcal{E}_0 , and where \succeq is the Blackwell order given by

$$-\mathcal{E}_0-\boxed{G}-\mathcal{E}- \succeq -\mathcal{E}_0-\boxed{G'}-\mathcal{E}'- \Leftrightarrow \exists -\mathcal{E}-\boxed{M}-\mathcal{E}'- : -\mathcal{E}_0-\boxed{G}-\mathcal{E}-\boxed{M}-\mathcal{E}'- = -\mathcal{E}_0-\boxed{G'}-\mathcal{E}'- . \quad (2.37)$$

Proof. If $-\mathbb{X}-\boxed{U}-\mathbb{Y}-$ is a universal dilation of T as asserted, then the map $\text{Trans}_{\Theta}(\mathcal{E}_0, -) \rightarrow \text{Dil}_0(T)$ given by $G \mapsto -\mathbb{X}-\boxed{U}-\boxed{G}-\mathbb{Y}-$ is an order-isomorphism. \square

Remark 2.4.9. (On Terminology.)

The name *Blackwell order* was pointed out to me by Tobias Fritz; I have not subsequently been able to find a reference for this, but I have no reasons not to believe him. \spadesuit

The Blackwell order is non-trivial already in simple theories like **FinSets**^{*}. The above result implies that intricacies of this order are mirrored in the dilational ordering. For example, we can now prove that even for very simple channels in **CIT** the dilational ordering has infinitely many inequivalent levels:

Example 2.4.10. (Infinite Descent in $(\text{Dil}_0(\text{id}_{\{0,1\}}), \succeq)$ in **CIT**.)

Consider the identity channel $-\{0,1\}-\boxed{\text{id}}-\{0,1\}-$ in **CIT**. It has as universal dilation the copy-channel $-\{0,1\}-\boxed{\text{id}}-\{0,1\}-$ whose hidden system is $\{0,1\}$. By Proposition 2.4.8, the pre-order of its one-sided dilations is therefore isomorphic to the Blackwell order on $\text{Trans}_{\text{CIT}}(\{0,1\}, -)$.

Consider in particular the sequence of channels $(-\{0,1\}-\boxed{G_n}-\{0,1\}-)_{n \in \mathbb{N}_0}$ given by

$$G_n(\delta_0) = \delta_0, \quad G_n(\delta_1) = \left(1 - \frac{1}{2^n}\right) \delta_0 + \frac{1}{2^n} \delta_1. \quad (2.38)$$

Clearly, $G_0 = \text{id}_{\{0,1\}}$ and $G_{n+1} = M \circ G_n$, with $M(\delta_0) = \delta_0$, $M(\delta_1) = \frac{1}{2} \delta_0 + \frac{1}{2} \delta_1$. (Plainly speaking, G_n represents n iterations of a channel which always preserves 0 but flips 1 to 0 with probability 1/2.) As such, the sequence decreases according to the Blackwell order,

$$G_0 \succeq G_1 \succeq G_2 \succeq \dots \quad (2.39)$$

However, each inequality must be strict: If for some $n \in \mathbb{N}_0$ we had $G_{n+1} \succeq G_n$, i.e. if there were a channel $N : \{0,1\} \rightarrow \{0,1\}$ with $G_n = N \circ G_{n+1}$, then by evaluating this identity in δ_0 and δ_1 we would find

$$\delta_0 = N(\delta_0), \quad \left(1 - \frac{1}{2^n}\right) \delta_0 + \frac{1}{2^n} \delta_1 = \left(1 - \frac{1}{2^{n+1}}\right) N(\delta_0) + \frac{1}{2^{n+1}} N(\delta_1), \quad (2.40)$$

which leads to $N(\delta_1) = 2\delta_1 - \delta_0$, contradicting that N is a classical channel.

(Intuitively, it is clear that no channel N could restore the damage inflicted by the erroneous channel M , and that is essentially what is asked for.) \blacklozenge

We end this section by demonstrating that **QIT** is universal:

Theorem 2.4.11. *In **QIT**, every channel has a universal dilation.*

Proof. By Theorem 2.3.22, every channel $\text{---}\mathbb{X}\text{---}\boxed{\Lambda}\text{---}\mathbb{Y}\text{---}$ in **QIT** has complete dilation given by a Stinespring dilation $\text{---}\mathbb{X}\text{---}\boxed{\Sigma}\text{---}\overset{\mathcal{E}}{\sim}\text{---}\mathbb{Y}\text{---}$. By Lemma 2.4.12 below, any minimal Stinespring dilation moreover meets the uniqueness clause in the definition of universality and is thus a universal dilation. (A *minimal* Stinespring dilation is one for which the system \mathcal{E} is as small as possible. For finite-dimensional spaces, this is simply a matter of choosing the dimension minimal, but in general it means that Σ has full rank on \mathcal{E} , in the sense that the smallest closed subspace which contains all local supports of the states $\{\Sigma(\varrho) \mid \varrho \in \text{St}(\mathcal{H}_1)\}$, where \mathcal{H}_1 is the total system corresponding to the interface \mathbb{X} , is all of \mathcal{E} ; compactly,⁸ $\bigvee_{\varrho \in \text{St}(\mathcal{H}_1)} \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\varrho)]) = \mathcal{E}$.) \square

Lemma 2.4.12. *(Generalised Choi-Jamiołkowski Isomorphism.)*

Suppose that $\text{---}\mathcal{H}_1\text{---}\boxed{\Sigma}\text{---}\overset{\mathcal{E}}{\sim}\text{---}\mathcal{K}_1\text{---}$ is an isometric channel in **QIT** (or **QIT**[∞]) with full rank on system \mathcal{E} , meaning that

$$\mathcal{E} = \bigvee_{\varrho \in \text{St}(\mathcal{H}_1)} \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\varrho)]). \quad (2.41)$$

If $\text{---}\mathcal{H}_2\text{---}\boxed{\Lambda}\text{---}\mathcal{K}_2\text{---}$ and $\text{---}\mathcal{H}_2\text{---}\boxed{\Lambda'}\text{---}\mathcal{K}_2\text{---}$ are quantum channels such that

$$\begin{array}{c} \text{---}\mathcal{H}_2\text{---} \\ \text{---}\mathcal{H}_1\text{---} \end{array} \begin{array}{c} \boxed{\Lambda} \\ \boxed{\Sigma} \end{array} \begin{array}{c} \text{---}\mathcal{K}_2\text{---} \\ \text{---}\mathcal{E}\text{---} \\ \text{---}\mathcal{K}_1\text{---} \end{array} = \begin{array}{c} \text{---}\mathcal{H}_2\text{---} \\ \text{---}\mathcal{H}_1\text{---} \end{array} \begin{array}{c} \boxed{\Lambda'} \\ \boxed{\Sigma} \end{array} \begin{array}{c} \text{---}\mathcal{K}_2\text{---} \\ \text{---}\mathcal{E}\text{---} \\ \text{---}\mathcal{K}_1\text{---} \end{array}, \quad (2.42)$$

then it holds that $\Lambda = \Lambda'$.

Remark 2.4.13. (Relation to Ordinary Choi-Jamiołkowski Isomorphism.)

The injectivity of the ordinary Choi-Jamiołkowski isomorphism ([JLF13, dP67, Cho75, Jam72]) corresponds to the special case where $\mathcal{H}_2 = \mathcal{H}_1 = \mathbb{C}$, so that Σ is simply a pure state on $\mathcal{E} \otimes \mathcal{K}_1$ with full rank on \mathcal{E} (usually taken to be maximally entangled). \spadesuit

Remark 2.4.14. (On Scope.)

Since the result holds as well in the case of infinite-dimensional Hilbert spaces, the theory **QIT**[∞] is universal by the same chain of arguments as used for **QIT**. \spadesuit

Proof. I will give two different proofs of this statement. Note that, in any case, we may assume without loss of generality that $\mathcal{H}_2 = \mathbb{C}$ (so the \mathcal{H}_2 -wire can be omitted), since if the implication holds in that case then from Eq. (2.42) we conclude that $\begin{array}{c} \boxed{\varrho}\text{---}\mathcal{H}_2\text{---} \\ \text{---}\mathcal{E}\text{---} \end{array} \boxed{\Lambda}\text{---}\mathcal{K}_2\text{---} = \begin{array}{c} \boxed{\varrho}\text{---}\mathcal{H}_2\text{---} \\ \text{---}\mathcal{E}\text{---} \end{array} \boxed{\Lambda'}\text{---}\mathcal{K}_2\text{---}$ for all states ϱ on \mathcal{H}_2 , which implies that $\Lambda = \Lambda'$ since quantum channels are determined by their action on product states. **Hence, we assume throughout that**

⁸Here, the *support* $\text{supp}(P)$ of a positive linear operator P means the range $\text{Im } P$, or, equivalently, the orthogonal complement of $\ker P$.

the domain of Λ and Λ' is the simple interface corresponding to system \mathcal{E} .

The first proof looks simple, though it actually depends on some slightly tricky arguments detailed in the footnotes (one of which is invalid in the infinite-dimensional version of the statement). This proof was found in conversation with David Pérez-García when I visited Madrid, and it works essentially by reduction to the injectivity of the usual Choi-Jamiołkowski isomorphism.

Let $\boxed{\phi}^{\mathcal{H}_1-}_{\mathcal{H}_1-}$ be a pure state with marginals of full rank. If we can argue that the pure state

$$\boxed{\tilde{\phi}}^{\mathcal{E}-}_{\mathcal{H}_1-} := \boxed{\phi}^{\mathcal{H}_1-}_{\mathcal{H}_1-} \boxed{\Sigma}^{\mathcal{E}-}_{\mathcal{K}_1-} \quad (2.43)$$

has full rank on \mathcal{E} , then the identity $\Lambda = \Lambda'$ follows from Eq. (2.42) by pre-composing with ϕ and invoking injectivity of the usual Choi-Jamiołkowski isomorphism. To show that $\tilde{\phi}$ has full rank on \mathcal{E} is to show that

$$\text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\tau)]) = \mathcal{E}, \quad (2.44)$$

where $\tau := (\text{id}_{\mathcal{H}_1} \otimes \text{tr}_{\mathcal{H}_1})(\phi)$ is the marginal of ϕ . To this end, let ϱ be any state on \mathcal{H}_1 . Since τ has full rank on \mathcal{H}_1 by assumption, there exists $p > 0$ such that $p\varrho \leq \tau$.⁹ This implies by positivity of the map $A \mapsto (\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(A)]$ that

$$p(\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\varrho)] \leq (\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\tau)], \quad (2.45)$$

and this operator inequality in turn implies¹⁰ that

$$\text{supp}(p(\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\varrho)]) \subseteq \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\tau)]). \quad (2.46)$$

Since $p \neq 0$, the containment of supports holds also when p is omitted from the left hand side, so as ϱ was arbitrary we have

$$\bigvee_{\varrho \in \text{St}(\mathcal{H}_1)} \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\varrho)]) \subseteq \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\tau)]), \quad (2.47)$$

⁹This is by the spectral theorem, using that the eigenvalues of τ are lower bounded by a strictly positive constant. It generally fails in infinite-dimensional spaces; for example, if in $\ell^2(\mathbb{N})$ the state τ is $\sum_{n=1}^{\infty} \frac{2}{3^n} |e_n\rangle\langle e_n|$ and ϱ is the pure state with vector representative $\sum_{n=1}^{\infty} \frac{1}{\sqrt{2^n}} |e_n\rangle$ (here, $(|e_n\rangle)_{n \in \mathbb{N}}$ denotes an orthonormal basis), then the operator inequality $p\varrho \leq \tau$ implies $p\langle e_m | \varrho | e_m \rangle \leq \langle e_m | \tau | e_m \rangle$, i.e. $p/2^m \leq 2/3^m$ for all m , which forces $p \leq 0$.

¹⁰To show this, we may equivalently show that the operator inequality $A \geq B \geq 0$ implies $\ker(A) \subseteq \ker(B)$. This can be proved in multiple ways, but the following elegant argument I owe to Lukas Schimmer: If $x \in \ker(A)$, then $0 = \langle x, Ax \rangle \geq \langle x, Bx \rangle \geq 0$, so $\langle x, Bx \rangle = 0$. By the spectral theorem, $B = \sum_j \lambda_j P_{\lambda_j}$, where λ_j is the j th eigenvalue of B and P_{λ_j} the orthogonal projection onto the corresponding eigenspace, and the identity $\langle x, Bx \rangle = 0$ thus becomes $\sum_j \lambda_j \|P_{\lambda_j} x\|^2 = 0$. As the eigenvalues λ_j are non-negative this implies that $P_{\lambda_j} x = 0$ for any j with $\lambda_j > 0$, but this means that $Bx = 0$, i.e. that $x \in \ker(B)$.

and the full rank assumption on Σ then yields the conclusion $\text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\tau)]) = \mathcal{E}$, as desired. This concludes the first proof.

The second proof is the one I found originally. It proceeds in three steps to reduce to the special case where $\mathcal{K}_2 = \mathcal{E}$, $\Lambda = \text{id}_{\mathcal{E}}$ and Λ' is a unitary conjugation \mathcal{E} , and then handles this special case by combining a topological argument with a consideration about the cardinality of the spectrum of a unitary operator.

Claim 1: *We may assume that Λ and Λ' are both isometric.* Indeed, suppose that Λ and Λ' are general channels satisfying (2.42). If $\hat{\Lambda}$ and $\hat{\Lambda}'$ are Stinespring dilations of Λ and Λ' respectively, then the channels

$$\begin{array}{c} \hat{\Lambda} \\ \text{---} \mathcal{K}_2 \text{---} \\ \text{---} \mathcal{K}_1 \text{---} \\ \Sigma \\ \text{---} \mathcal{H}_1 \text{---} \end{array} \quad \text{and} \quad \begin{array}{c} \hat{\Lambda}' \\ \text{---} \mathcal{K}_2 \text{---} \\ \text{---} \mathcal{K}_1 \text{---} \\ \Sigma \\ \text{---} \mathcal{H}_1 \text{---} \end{array} \quad (2.48)$$

are Stinespring dilations of the same channel, so there exist isometric channels Γ and Γ' such that

$$\begin{array}{c} \hat{\Lambda} \\ \text{---} \mathcal{K}_2 \text{---} \\ \text{---} \mathcal{K}_1 \text{---} \\ \Sigma \\ \text{---} \mathcal{H}_1 \text{---} \end{array} \begin{array}{c} \Gamma \\ \text{---} \tilde{\mathcal{F}} \text{---} \end{array} = \begin{array}{c} \hat{\Lambda}' \\ \text{---} \mathcal{K}_2 \text{---} \\ \text{---} \mathcal{K}_1 \text{---} \\ \Sigma \\ \text{---} \mathcal{H}_1 \text{---} \end{array} \begin{array}{c} \Gamma' \\ \text{---} \tilde{\mathcal{F}} \text{---} \end{array} \quad (2.49)$$

If the implication holds for isometric channels, it follows that $\begin{array}{c} \hat{\Lambda} \\ \text{---} \mathcal{K}_2 \text{---} \\ \text{---} \mathcal{E} \text{---} \end{array} \begin{array}{c} \Gamma \\ \text{---} \tilde{\mathcal{F}} \text{---} \end{array} = \begin{array}{c} \hat{\Lambda}' \\ \text{---} \mathcal{K}_2 \text{---} \\ \text{---} \mathcal{E} \text{---} \end{array} \begin{array}{c} \Gamma' \\ \text{---} \tilde{\mathcal{F}} \text{---} \end{array}$, and trashing $\tilde{\mathcal{F}}$ we then conclude $\Lambda = \Lambda'$.

Claim 2: *We may assume that $\mathcal{K}_2 = \mathcal{E} \otimes \mathcal{R}$ for some system \mathcal{R} , and that Λ and Λ' are both isometric, with Λ of the form $\begin{array}{c} \sigma \\ \text{---} \mathcal{R} \text{---} \\ \text{---} \mathcal{E} \text{---} \end{array}$ for some pure state σ .* For this, first observe that we may by the previous claim assume both channels to be isometric. Clearly, we may also assume that $\mathcal{K}_2 = \mathcal{E} \otimes \mathcal{R}$ for some system \mathcal{R} , by possibly isometrically embedding the isometries into a larger space. Finally, any isometric conjugation from \mathcal{E} to $\mathcal{E} \otimes \mathcal{R}$ is of the form $\begin{array}{c} \sigma \\ \text{---} \mathcal{R} \text{---} \\ \text{---} \mathcal{E} \text{---} \end{array} \begin{array}{c} \mathcal{U} \\ \text{---} \mathcal{R} \text{---} \\ \text{---} \mathcal{E} \text{---} \end{array}$ for some pure state σ on \mathcal{R} and some unitary conjugation \mathcal{U} , so the desired follows by moving \mathcal{U} to the other side of the identity, invoking injectivity in the special case, and moving \mathcal{U} back.

Claim 3: *We may assume that $\mathcal{K}_2 = \mathcal{E}$, that $\Lambda = \text{id}_{\mathcal{E}}$, and that Λ' is a unitary conjugation on \mathcal{E} .* Indeed, assume that the implication holds in this case. Let Λ and Λ' be of the form from the case in the previous claim – i.e. $\mathcal{K}_2 = \mathcal{E} \otimes \mathcal{R}$ for some system \mathcal{R} , Λ isometric of the form $\begin{array}{c} \sigma \\ \text{---} \mathcal{R} \text{---} \\ \text{---} \mathcal{E} \text{---} \end{array}$ for some pure state σ , and Λ' arbitrary isometric from \mathcal{E} to $\mathcal{E} \otimes \mathcal{R}$ – and suppose that Λ and Λ' satisfy the identity

$$\begin{array}{c} \sigma \\ \text{---} \mathcal{R} \text{---} \\ \text{---} \mathcal{E} \text{---} \\ \Sigma \\ \text{---} \mathcal{K}_1 \text{---} \end{array} = \begin{array}{c} \Lambda' \\ \text{---} \mathcal{R} \text{---} \\ \text{---} \mathcal{E} \text{---} \\ \Sigma \\ \text{---} \mathcal{K}_1 \text{---} \end{array} ; \quad (2.50)$$

we must show that $\Lambda = \Lambda'$. By tracing out system \mathcal{K}_1 in (2.50), we get

$$\begin{array}{c} \sigma \\ \hline \Sigma \\ \hline \mathcal{K}_1 \text{---tr} \end{array} \begin{array}{c} \mathcal{R} \\ \hline \mathcal{E} \end{array} = \begin{array}{c} \Lambda' \\ \hline \Sigma \\ \hline \mathcal{K}_1 \text{---tr} \end{array} \begin{array}{c} \mathcal{R} \\ \hline \mathcal{E} \end{array}, \quad (2.51)$$

and by inserting various states $\varrho \in \text{St}(\mathcal{H}_1)$, this implies that

$$\begin{array}{c} \sigma \\ \hline \mu \\ \hline \mathcal{E} \end{array} \begin{array}{c} \mathcal{R} \\ \hline \mathcal{E} \end{array} = \begin{array}{c} \Lambda' \\ \hline \mu \\ \hline \mathcal{E} \end{array} \begin{array}{c} \mathcal{R} \\ \hline \mathcal{E} \end{array} \quad (2.52)$$

for all states μ in the set $\mathcal{S}_0 := \{(\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})[\Sigma(\varrho)] \mid \varrho \in \text{St}(\mathcal{H}_1)\} \subseteq \text{St}(\mathcal{E})$. The full rank assumption on Σ does not guarantee that \mathcal{S}_0 contains every state on \mathcal{E} , but it does ensure that all of \mathcal{E} is covered by the supports, in the sense that

$$\bigvee_{\mu \in \mathcal{S}_0} \text{supp}(\mu) = \mathcal{E}. \quad (2.53)$$

Now, letting $|\sigma\rangle \in \mathcal{R}$ be a vector representative of the pure state σ , and letting $W : \mathcal{E} \rightarrow \mathcal{E} \otimes \mathcal{R}$ be an isometry representing Λ' (i.e. $\Lambda'(A) = WAW^*$), Eq. (2.52) implies that

$$\text{supp}(\mu) \otimes \text{span}\{|\sigma\rangle\} = W(\text{supp}(\mu)) \quad (2.54)$$

for all $\mu \in \mathcal{S}_0$. Eq. (2.53) then yields

$$\mathcal{E} \otimes \text{span}\{|\sigma\rangle\} = \bigvee_{\mu \in \mathcal{S}_0} W(\text{supp}(\mu)) = W\left(\bigvee_{\mu \in \mathcal{S}_0} \text{supp}(\mu)\right) = \text{Im}(W), \quad (2.55)$$

using linearity and continuity of W for the middle equality. This identity, however, implies that $W = U \otimes |\sigma\rangle$ for some unitary operator $U : \mathcal{E} \rightarrow \mathcal{E}$, and thus $\Lambda' = \mathcal{U} \otimes \sigma$ for some unitary conjugation \mathcal{U} on \mathcal{E} . By the assumption that injectivity holds in the special unitary case, we must have $\mathcal{U} = \text{id}_{\mathcal{E}}$, and this implies that $\Lambda = \Lambda'$ as desired.

Claim 4: *The implication of the lemma holds in the case described in the previous claim. The assumption is that \mathcal{U} is a unitary conjugation on \mathcal{E} such that*

$$\begin{array}{c} \Sigma \\ \hline \mathcal{K}_1 \end{array} \begin{array}{c} \mathcal{E} \\ \hline \mathcal{E} \end{array} = \begin{array}{c} \Sigma \\ \hline \mathcal{K}_1 \end{array} \begin{array}{c} \mathcal{E} \text{---}\mathcal{U}\text{---}\mathcal{E} \\ \hline \mathcal{E} \end{array}, \quad (2.56)$$

and the objective is to show that $\mathcal{U} = \text{id}_{\mathcal{E}}$.

Let $U : \mathcal{E} \rightarrow \mathcal{E}$ be a unitary operator which represents \mathcal{U} (i.e. $\mathcal{U}(A) = UAU^*$). We must show that $U = \zeta_0 \mathbb{1}_{\mathcal{E}}$, where $\zeta_0 \in \mathbb{C}$ has modulus 1 (here, $\mathbb{1}_{\mathcal{E}}$ denotes the identity operator on \mathcal{E}). From (2.56) it follows that

$$\begin{array}{c} \Sigma(\psi) \\ \hline \mathcal{K}_1 \end{array} \begin{array}{c} \mathcal{E} \\ \hline \mathcal{E} \end{array} = \begin{array}{c} \Sigma(\psi) \\ \hline \mathcal{K}_1 \end{array} \begin{array}{c} \mathcal{E} \text{---}\mathcal{U}\text{---}\mathcal{E} \\ \hline \mathcal{E} \end{array} \quad (2.57)$$

for any pure state $\psi \in \text{St}(\mathcal{H}_1)$. This identity in turn lends itself to a use of the ordinary Choi-Jamiołkowski isomorphism, or at least a variation of it: If $|\phi\rangle \in \mathcal{E} \otimes \mathcal{K}_1$ is a vector representative of the pure state $\phi := \Sigma(\psi)$, and if $|\phi\rangle = \sum_{j=1}^r \sqrt{p(j)} |\phi^{\mathcal{E}}(j)\rangle \otimes |\phi^{\mathcal{K}_1}(j)\rangle$ is a Schmidt decomposition with $p(j) > 0$ for all $j = 1, \dots, r$, then Eq. (2.57), which asserts the equality of two pure state, reads in terms of vector representatives

$$\sum_{j=1}^r \sqrt{p(j)} [U |\phi^{\mathcal{E}}(j)\rangle] \otimes |\phi^{\mathcal{K}_1}(j)\rangle = \zeta(\psi) \sum_{j=1}^r \sqrt{p(j)} |\phi^{\mathcal{E}}(j)\rangle \otimes |\phi^{\mathcal{K}_1}(j)\rangle \quad (2.58)$$

for some (unique) phase $\zeta(\psi) \in \mathbb{C}$ with $|\zeta(\psi)| = 1$. As $(|\phi^{\mathcal{K}_1}(j)\rangle)_{j=1, \dots, r}$ is an orthonormal system, this implies that $U |\phi^{\mathcal{E}}(j)\rangle = \zeta(\psi) |\phi^{\mathcal{E}}(j)\rangle$ for all $j = 1, \dots, r$, or, equivalently, that

$$U |\phi\rangle = \zeta(\psi) |\phi\rangle \quad (2.59)$$

for all vectors $|\phi\rangle$ in the subspace

$$\text{span}\{|\phi^{\mathcal{E}}(j)\rangle \mid j = 1, \dots, r\} = \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})(\phi)) = \text{supp}((\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})(\Sigma(\psi))). \quad (2.60)$$

Now, if $\phi = \Sigma(\psi)$ had full rank on \mathcal{E} , i.e. if this span were all of \mathcal{E} (for some pure state ψ on \mathcal{H}_1), then we would be done, since (2.59) would then assert that U acts as a multiple of the identity operator globally. In general, however, the full-rank assumption on Σ merely implies that \mathcal{E} can be ‘patched up’ from subspaces on which U acts as (possibly different) multiples of the identity operator. We are saved by properties of a topological nature:

Since U acts as in (2.59) on the non-zero subspace (2.60), the function $\zeta(\cdot)$ admits the explicit expression

$$\zeta(\psi) = \text{tr}(U(\text{id}_{\mathcal{E}} \otimes \text{tr}_{\mathcal{K}_1})(\Sigma(\psi))) \quad (2.61)$$

and for each ψ , $\zeta(\psi)$ is an eigenvalue of U (whose eigenspace contains the subspace (2.60)). Now, the function ζ defined by Eq. (2.61) is clearly continuous w.r.t. the topology induced by the trace-distance and the ordinary topology on \mathbb{C} . Also, the set of pure states on \mathcal{H}_1 is path-connected with the topology induced by the trace-distance (meaning that for any pure states ψ_0 and ψ_1 on \mathcal{H}_1 , we can find a continuous map $\gamma : [0, 1] \rightarrow \text{St}(\mathcal{H}_1)$ whose range contains only pure states, with $\gamma(0) = \psi_0$ and $\gamma(1) = \psi_1$). Therefore, the image of this set under $\zeta(\cdot)$ must be a non-empty path-connected subset of \mathbb{C} . However, every non-empty path-connected subset of \mathbb{C} is either uncountable or consists of a single point, and since $\text{Im} \zeta$ is a set of eigenvalues of U , it cannot be uncountable if \mathcal{E} is only finite-dimensional, or even separable. Consequently, $\text{Im} \zeta$ contains only a single point, i.e. $\zeta(\cdot)$ is constant. With ζ_0 denoting this constant value, we conclude that $U = \zeta_0 \mathbf{1}_{\mathcal{E}}$ globally, as desired. \square

2.5 Purification

The results of the previous sections imply that **CIT**, **QIT** and all cartesian theories are localisable and universal. In particular, all of the dilational properties we have considered

have been shared by the two information theories **CIT** and **QIT** alike. One dilational property, however, will set them apart, and that property is the topic of this section.

Picturesquely speaking, complete dilations are *omniscient*: If K is a complete dilation of T , then K knows everything there is to know about T . Both **CIT** and **QIT** admit complete dilations.

In **QIT**, however, there exist complete dilations of T which, in addition to knowing everything about T , also know everything there is to know *about themselves*. Indeed, we have already seen that isometric channels in **QIT** are dilationally pure (Corollary 2.3.23), that is, *self-complete*, and every quantum channel admits an isometric dilation, namely its Stinespring dilation. This phenomenon – the fact that a quantum channel can be *purified* – abruptly terminates the process of forming dilations. *Exhaustive knowledge is possible*.

In contrast, dilational purity – or, self-completeness – in **CIT** is a property reserved for the dull; any channel which is not a pure state has non-trivial dilations (Proposition 2.2.11). As such, there is no ceiling to the formation of dilations: A complete dilation K_0 of a channel T (given by copying all inputs and outputs) is not a complete dilation of itself; of course, K_0 too has a complete dilation, K_1 , but K_1 is not a complete dilation of itself either – and so it continues, ad infinitum, with complete dilations K_2 of K_1 , K_3 of K_2 , and so forth. *There is always more to know*.

The goal of this section is to demonstrate that this distinction between **CIT** and **QIT** is not a randomly chosen one, but rather one that can be seen as responsible for many features which distinguish the two theories. As such, the significance of the isometric channels in the so-called ‘Church of the larger Hilbert space’¹¹ is not that they are reversible, or per se that they are isometric – the significance is that *in the church, (dilational) purity is obtained*. Jokingly, one might say that **QIT** is like Christianity whereas **CIT** is like Buddhism.

‘Purification’ as an axiom has been considered in the literature before, with the profound conclusion that the property more or less characterises quantum theory uniquely ([CDP10, CDP11]). However, there are differences in the statement of that principle, and the categorical framework used here is technically simpler and independent of probabilistic structure. Moreover, the main results of this chapter do not have counterparts in Refs. [CDP10, CDP11], and they are derived from fewer principles, none of which are not about dilations. (Further details are given in Remark 2.5.4.)

Definition 2.5.1. (Purifiable Theories.)

A theory Θ is called *purifiable* if every channel in Θ has a dilationally pure dilation. ■

Example 2.5.2. (Purification in the Information Theories.)

The theory **QIT** is purifiable, since Stinespring dilations are dilationally pure. The theory **CIT** is not purifiable, since the only channels which are dilationally pure are probabilistically pure states. ♦

Example 2.5.3. (Purifiable Thin Theories.)

A thin theory described by the monoid $(M, \star, 1, \succeq)$ is purifiable if \star satisfies the cancellation law $z \star u \succeq z \star v \Rightarrow u \succeq v$. Indeed, for any channel $-x-\square-y-$ the dilation $\begin{array}{c} \sim y \sim \\ -x-\square-y- \\ \sim x \sim \end{array}$ is dilationally pure by Example 2.2.10 if \star satisfies the cancellation law. Thus, for example, the thin theory $(\mathbb{N}, \cdot, 1, \geq)$ is purifiable. ♦

¹¹A phrase coined by John A. Smolin ([Chu]) about the possibility – by virtue of Stinespring’s dilation theorem – of always regarding a quantum channel as the marginal of an isometric channel into a larger space.

Remark 2.5.4. (Relation to the Purification Postulate of Refs. [CDP10, CDP11].)

The main result of Ref. [CDP11] is that, within a large ground class of theories (defined by a list of ‘standing assumptions’), five reasonable axioms determine a subclass of reasonable information theories, and an additional ‘Purification Postulate’ (first introduced in Ref. [CDP10]) uniquely identifies the theory **QIT**. In our language, the Purification Postulate of Refs. [CDP10, CDP11] is the requirement that every state has a dilationally pure one-sided dilation, and that any two such dilations with the same hidden interface are related by an isomorphism on the hidden interface (so-called ‘uniqueness’ of purifications). This requirement is clearly from the same womb as that of Definition 2.5.1, but there are important differences.

The Purification Postulate as formulated in Refs. [CDP10, CDP11] is about *probabilistic* purity and as such does not a priori pertain to the general theories considered here; however, it can be reformulated equivalently in terms of dilational purity, so we may ignore that difference.

More importantly, the Purification Postulate concerns only states, and though this seems to make it more general, the truth is actually the opposite:

Firstly, according to Thm. 15 of Ref. [CDP10], the Purification Postulate implies (within their large ground class of theories) that any channel $\text{---}\mathbb{X}\text{---}\boxed{T}\text{---}\mathbb{Y}\text{---}$ has a dilation of the form

$\text{---}\mathbb{X}\text{---}\boxed{\phi}\text{---}\boxed{\alpha}\text{---}\mathbb{Y}\text{---}$, where ϕ is a pure state and α an isomorphism. (This can be thought of as a generalisation of Stinespring’s theorem, with α a unitary conjugation). By virtue of their ‘dynamically faithful states’ (these are like the full rank states in the Choi-Jamiołkowski isomorphism) one can easily show that the channel $\text{---}\mathbb{X}\text{---}\boxed{\phi}\text{---}\boxed{\alpha}\text{---}\mathbb{Y}\text{---}$ is dilationally pure, by

converting the argument to an argument about purity of states. Thus, within their ground class of theories, the Purification Postulate of Refs. [CDP10, CDP11] is a **stronger requirement** than purifiability in the sense of Definition 2.5.1. (I do not know whether they are actually equivalent within this ground class; this depends on whether their uniqueness clause follows from purifiability in the sense of Definition 2.5.1.)

But secondly, not only is Definition 2.5.1 less restrictive than the Purification Postulate, its effective scope is also larger: Indeed, the ground class of theories in Refs. [CDP10, CDP11] assumes among other things that transformations are determined by their action on states (excluding thin theories and some cartesian theories), and that theories are non-deterministic (excluding cartesian theories), cf. Def. 2 in [CDP10]. These standing assumptions are used repeatedly when deriving quantum-like consequences of the Purification Postulate, and this has the subtle side-effect that some theories not complying to the standing assumptions satisfy the Purification Postulate without being anything like quantum theory; for example, every cartesian theory, e.g. **Sets***, satisfies the Purification Postulate, for the simple reason that its states are already pure. As such, purifiability in the sense of Definition 2.5.1 seems more robust in that it does not accidentally include such theories.

As a final remark it should be noted that the uniqueness clause in the Purification Postulate can quite easily be used to show that the dilations $\text{---}\mathbb{X}\text{---}\boxed{\phi}\text{---}\boxed{\alpha}\text{---}\mathbb{Y}\text{---}$ are complete (in the same way we demonstrated completeness of Stinespring dilations). Hence, the Purification Postulate ultimately implies both purifiability and completeness, whereas the framework of this chapter separates the two.

✦

In the remainder of the section, we state and prove three general aspects of theories which are purifiable and comply to various subsets of our previous dilational principles. All

of these aspects are traditionally considered ‘quantum’, but the main point here is that they can be seen rather as consequences of a few simple and abstract principles about dilations.

The first result (Proposition 2.5.5) has to do with isomorphisms in a purifiable theory, and we prove in particular a rather simple recharacterisation of universal purifiable theories (Proposition 2.5.10).

Secondly, we prove a structure theorem for reversible channels in a universal, temporally localisable and purifiable theory (Theorem 2.5.11). In the case of **QIT**, it reproduces the not entirely trivial result that a quantum channel is reversible precisely if it is the tensoring with a (possibly mixed) ancillary state, followed by an isometric conjugation.

Finally, we discuss how purifiability entails (in conjunction with the other dilational principles) the notion of *complementarity* between channels, generalising (by Theorem 2.5.17) the concept of complementary quantum channels ([DS05]). In particular, we recover an abstract version of the complementarity between reversible and ‘completely forgetful’ channels (Theorem 2.5.18). This complementarity will be generalised in Chapter 3 to an approximate setting.

2.5.A Isomorphisms and Purifiability

We begin by proving that in a purifiable theory, isomorphisms not only have dilationally pure dilations, but must in fact already themselves be dilationally pure:

Proposition 2.5.5. (*Isomorphisms and Purifiability.*)

If Θ is a purifiable theory, every isomorphism in Θ is dilationally pure.

Remark 2.5.6. (Complete Characterisation of Purifiable Thin Theories.)

Example 2.5.3 shows that for a thin theory $(M, \star, 1, \succeq)$, the cancellation law $x \star y \succeq x \star z \Rightarrow y \succeq z$ implies purifiability. Proposition 2.5.5 shows the opposite implication, since to assert that the identity $x \succeq x$ is dilationally pure is precisely to assert that cancellation law. \spadesuit

Proof. Let us start by showing that any identity $\text{id}_{\mathcal{X}}$ is dilationally pure. By assumption, $-\mathcal{X}-\boxed{\text{id}}-\mathcal{X}-$ has a dilationally pure dilation, say $\begin{array}{c} \sim \mathbb{D}_0 \sim \\ \sim \mathcal{X} - \boxed{P} - \mathbb{E}_0 \sim \\ \sim \mathcal{X} - \end{array}$. Let $\begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} - \boxed{L} - \mathbb{E} \sim \\ \sim \mathcal{X} - \end{array}$ be any dilation of $-\mathcal{X}-\boxed{\text{id}}-\mathcal{X}-$; we show that it is trivial.

The channel

$$\begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} - \boxed{L} - \mathbb{E} \sim \\ \sim \mathcal{X} - \boxed{P} - \mathbb{E}_0 \sim \\ \sim \mathbb{D}_0 \sim \end{array} \quad (2.62)$$

is a dilation of P with hidden interfaces \mathbb{D} and \mathbb{E} (trashing \mathbb{E} yields $P \boxplus \text{tr}_{\mathbb{D}}$, since L dilates $\text{id}_{\mathcal{X}}$). However, as P is dilationally pure, this dilation must be trivial, i.e. of the form

$$\begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} - \boxed{M} - \mathbb{E} \sim \\ \sim \mathcal{X} - \boxed{P} - \mathbb{E}_0 \sim \\ \sim \mathbb{D}_0 \sim \end{array} \quad (2.63)$$

for some channel M . Equating (2.62) with (2.63) and trashing \mathbb{E}_0 , we derive the identity $\begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} \sim \boxed{L} \sim \mathcal{X} \sim \\ \sim \mathbb{E} \sim \end{array} = \begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} \sim \boxed{\text{id}_{\mathcal{X}}} \sim \mathcal{X} \sim \\ \sim \mathbb{E} \sim \end{array}$, since P dilates $\text{id}_{\mathcal{X}}$ and the theory is normal (Definition 1.3.7). This demonstrates the desired.

Now, if $\sim \mathcal{X} \sim \boxed{\alpha} \sim \mathcal{Y} \sim$ is an arbitrary isomorphism then

$$\sim \mathcal{X} \sim \boxed{\alpha} \sim \mathcal{Y} \sim \boxed{\alpha^{-1}} \sim \mathcal{X} \sim = \sim \mathcal{X} \sim \boxed{\text{id}_{\mathcal{X}}} \sim \mathcal{X} \sim , \quad (2.64)$$

so for any dilation L of α , dilational purity of $\text{id}_{\mathcal{X}}$ entails that

$$\begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} \sim \boxed{L} \sim \mathcal{Y} \sim \boxed{\alpha^{-1}} \sim \mathcal{X} \sim \\ \sim \mathbb{E} \sim \end{array} = \begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} \sim \boxed{\text{id}_{\mathcal{X}}} \sim \mathcal{X} \sim \\ \sim \mathbb{E} \sim \end{array} \quad (2.65)$$

for some M . Composing with α on both sides gives $\begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} \sim \boxed{L} \sim \mathcal{Y} \sim \\ \sim \mathbb{E} \sim \end{array} = \begin{array}{c} \sim \mathbb{D} \sim \\ \sim \mathcal{X} \sim \boxed{\alpha} \sim \mathcal{Y} \sim \\ \sim \mathbb{E} \sim \end{array}$, and consequently α is dilationally pure. \square

We immediately recover as a special case the *No Broadcasting Theorem* ([BCF⁺96]):

Corollary 2.5.7. (*No Broadcasting.*)

If Θ is a purifiable theory, and if $\sim \mathcal{X} \sim \boxed{T} \sim \mathcal{X} \sim$ is a channel for which one marginal is $\text{id}_{\mathcal{X}}$, then the other marginal must be of the form $\sim \mathcal{X} \sim \boxed{\text{tr}} \boxed{s} \sim \mathcal{X} \sim$ for some state s on \mathcal{X} .

Remark 2.5.8. In a purifiable thin theory $(M, \star, 1, \succeq)$ (i.e. one for which \star satisfies the cancellation law), the ‘No Broadcasting’ theorem simply says that if $x \succeq x \star x$, then $1 \succeq x$. \star

We also recover the following non-trivial classification of isomorphisms in **QIT**:

Corollary 2.5.9. (*Isomorphisms in QIT.*)

A transformation in **QIT** is an isomorphism if and only if it is a unitary conjugation.

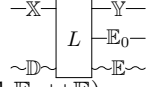
Proof. It is clear that all unitary conjugations are isomorphisms. Conversely, since **QIT** is purifiable, any isomorphism in **QIT** is purifiable by Proposition 2.5.5, hence an isometric conjugation by Corollary 2.3.23. By a dimensional argument, an isometric conjugation has an inverse only if it is unitary. \square

Finally, we can augment the result to yield a surprising recharacterisation of purifiability for universal theories:

Proposition 2.5.10. (*Recharacterisation of Purifiable Theories.*)

A universal theory Θ is purifiable if and only if all identities $\text{id}_{\mathcal{X}}$ are dilationally pure. Moreover, in this case every universal dilation of any channel is dilationally pure.

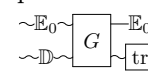
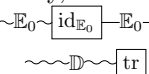
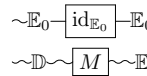
Proof. The ‘only if’-direction follows from Proposition 2.5.5. As for the ‘if’-direction, assume that all identities in Θ are dilationally pure. Let $\sim \mathcal{X} \sim \boxed{T} \sim \mathcal{Y} \sim$ be a channel and let $\sim \mathcal{X} \sim \boxed{U} \sim \mathcal{Y} \sim$ be a universal dilation of T . We show that U is dilationally pure. To this end,

let  be any dilation of U . Since L is a dilation of \underline{T} (with hidden interfaces \mathbb{D} and $\mathbb{E}_0 \cup \mathbb{E}$), we must have

$$\begin{array}{c} \text{---}X\text{---} \\ | \\ \boxed{L} \\ | \\ \text{---}Y\text{---} \\ | \\ \text{---}E_0\text{---} \\ | \\ \text{---}E\text{---} \\ | \\ \text{---}D\text{---} \end{array} = \begin{array}{c} \text{---}X\text{---} \\ | \\ \boxed{U} \\ | \\ \text{---}E_0\text{---} \\ | \\ \boxed{G} \\ | \\ \text{---}E\text{---} \\ | \\ \text{---}D\text{---} \end{array} \quad (2.66)$$

for some channel G . As L dilates U , however,

$$\begin{array}{c} \text{---}X\text{---} \\ | \\ \boxed{U} \\ | \\ \text{---}E_0\text{---} \\ | \\ \boxed{G} \\ | \\ \text{---}E\text{---} \\ | \\ \text{---}D\text{---} \end{array} \text{---}Y\text{---} = \begin{array}{c} \text{---}X\text{---} \\ | \\ \boxed{L} \\ | \\ \text{---}E_0\text{---} \\ | \\ \text{---}E\text{---} \\ | \\ \text{---}D\text{---} \end{array} \text{---}Y\text{---} = \begin{array}{c} \text{---}X\text{---} \\ | \\ \boxed{U} \\ | \\ \text{---}E_0\text{---} \\ | \\ \boxed{\text{id}_{E_0}} \\ | \\ \text{---}E_0\text{---} \\ | \\ \text{---}D\text{---} \end{array} \text{---}Y\text{---} \text{---} \text{tr} \quad (2.67)$$

and comparing the first and last expression of this identity, the uniqueness clause in the definition of universality implies that  = , i.e. G is a dilation of id_{E_0} . But by assumption id_{E_0} is dilationally pure, so must G must be  for some M , and inserting this into Eq. (2.66) we see that L is a trivial dilation of U . Hence, U is dilationally pure as asserted. \square

We will often use this result to ensure in a purifiable theory the existence of one-sided pure dilations (as every universal dilation is one-sided).

2.5.B Reversibles and Purifiability

By Proposition 2.5.5, every isomorphism in a purifiable theory is dilationally pure. The converse is not necessarily the case – for example, a non-unitary isometric channel in **QIT** is not an isomorphism. It is, however, reversible, and as we shall now see, a few dilational axioms will guarantee in general that dilationally pure channels are reversible. In fact, we can use the dilationally pure channels in a purifiable theory to precisely understand the class of reversible channels:

Theorem 2.5.11. (*Structure of Reversible Channels.*)

Let Θ be a universal, temporally localisable and purifiable theory. Then, a channel $\text{---}X\text{---}\boxed{T}\text{---}Y\text{---}$ in Θ is reversible if and only if it is of the form

$$\begin{array}{c} \text{---}X\text{---} \\ | \\ \boxed{P} \\ | \\ \text{---}Y\text{---} \\ | \\ \boxed{r} \\ | \\ \text{---}Z\text{---} \end{array} \quad (2.68)$$

for some state r and some dilationally pure P .

Proof. Temporal localisability is only used to the effect that every channel has a reversible one-sided dilation (Proposition 2.3.19).

If P is dilationally pure, then any reversible dilation is trivial, so P must already be reversible itself. Hence, it is clear that channels of the form (2.68) are reversible. The converse implication is a more intricate gymnastic exercise:

Assume that $-\mathcal{X}-\boxed{T}-\mathcal{Y}-$ is reversible, and let $-\mathcal{Y}-\boxed{T^-}-\mathcal{X}-$ be a channel with

$$-\mathcal{X}-\boxed{T}-\mathcal{Y}-\boxed{T^-}-\mathcal{X}- = -\mathcal{X}-\boxed{\text{id}}-\mathcal{X}- \quad . \quad (2.69)$$

Let $-\mathcal{X}-\boxed{\check{T}}-\mathcal{Y}-$ be a universal dilation of T ; by Proposition 2.5.10 it is dilationally pure. Let moreover $-\mathcal{Y}-\boxed{R_{T^-}}-\mathcal{X}-$ be a reversible dilation of T^- . The channel

$$-\mathcal{X}-\boxed{\check{T}}-\mathcal{Y}-\boxed{R_{T^-}}-\mathcal{X}- \quad (2.70)$$

is a dilation of $-\mathcal{X}-\boxed{\text{id}}-\mathcal{X}-$, and since Θ is purifiable, $\text{id}_{\mathcal{X}}$ is dilationally pure by Proposition 2.5.5, so

$$-\mathcal{X}-\boxed{\check{T}}-\mathcal{Y}-\boxed{R_{T^-}}-\mathcal{X}- = -\mathcal{X}-\boxed{\text{id}}-\mathcal{X}- \quad (2.71)$$

for some state s . Now, pick a left-inverse \tilde{P} of the reversible R_{T^-} , and observe that by composing both channels with \tilde{P} we obtain the identity

$$-\mathcal{X}-\boxed{\check{T}}-\mathcal{Y}- = -\mathcal{X}-\boxed{s}-\mathcal{Y}-\boxed{\tilde{P}}-\mathcal{Y}- \quad . \quad (2.72)$$

If we knew that \tilde{P} were dilationally pure, we could simply trash the system \mathbb{E} and conclude the form (2.68), but this conclusion is not within reach. The trick is to express \boxed{s} as $\boxed{u}-\boxed{G}-\mathbb{E}$, where u is a universal dilation of the marginal $\boxed{s}-\boxed{\text{tr}}$, thus obtaining

$$-\mathcal{X}-\boxed{\check{T}}-\mathcal{Y}- = -\mathcal{X}-\boxed{u}-\mathcal{Y}-\boxed{P}-\mathcal{Y}- \quad , \quad (2.73)$$

with P the composition of \tilde{P} with G . We can now argue that \underline{P} is dilationally pure: Since P has a complete (namely universal) one-sided dilation, it suffices to argue that any one-sided dilation of P is trivial. However, any one-sided dilation $-\mathcal{X}-\boxed{L}-\mathcal{Y}-$ of P (with hidden interface \mathbb{G}) gives rise, by virtue of Eq. (2.73), to a dilation of \check{T} , and since the latter is dilationally pure, we must have

$$\begin{array}{c}
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---} \\
\boxed{L} \\
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---} \\
\boxed{P} \\
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---} \\
\boxed{t} \\
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---}
\end{array}
=
\begin{array}{c}
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---} \\
\boxed{P} \\
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---} \\
\boxed{t} \\
\text{---}\mathcal{X}\text{---} \\
\text{---}\mathcal{Y}\text{---} \\
\text{---}\mathbb{E}\text{---} \\
\text{---}\mathbb{G}\text{---}
\end{array}
\quad (2.74)$$

for some state t . By the universality property of u , however, this implies the identity $\text{---}\mathcal{X}\text{---}\boxed{L}\text{---}\mathcal{Y}\text{---} = \text{---}\mathcal{X}\text{---}\boxed{P}\text{---}\mathcal{Y}\text{---}$, in other words, P is dilationally pure, as desired. Trashing the system \mathbb{E} in Eq. (2.73) finally yields the conclusion of the theorem. \square

The following follows immediately:

Corollary 2.5.12. (*Reversibles in QIT.*)

A quantum channel $\text{---}\mathcal{X}\text{---}\boxed{\Lambda}\text{---}\mathcal{Y}\text{---}$ is reversible if and only if it is of the form $\text{---}\mathcal{X}\text{---}\boxed{\Sigma}\text{---}\mathcal{Y}\text{---}$ for some state ρ and some isometric channel Σ .

Proof. Immediate from Theorem 2.5.11, since by Proposition 2.5.5 the dilationally pure channels in QIT are precisely the isometric channels. \square

The proof of Theorem 2.5.11 required surprising assumptions additional to purifiability of the theory, namely the existence of reversible dilations and of universal dilations. The following is left unanswered:

Open Problem 2.5.13. *Can the assumptions of Theorem 2.5.11 be weakened without compromising the conclusion?*

2.5.C Complementarity

If a channel knows everything about itself, does it necessarily know everything about any channel that it dilates? Moreover formally, if $\text{---}\mathcal{X}\text{---}\boxed{T}\text{---}\mathcal{Y}\text{---}$ is a channel and $\text{---}\mathcal{X}\text{---}\boxed{P}\text{---}\mathcal{Y}\text{---}$ is dilationally pure, must then P be a complete dilation of T ?

This is not obvious. If true, though, it has a remarkable consequence for purifiable theories: In QIT, Stinespring dilations creates a *complementarity* between certain pairs of quantum channels ([DS05]), and this is ultimately because it is a complete dilation of any channel that it dilates. Thus, an affirmative answer to the above question will foster a similar notion of complementarity in general. This will specialise in particular to a complementarity between reversible channels and *completely forgetful* channels (that is, channels of the form $\text{---}\mathcal{X}\text{---}\boxed{\text{tr}}\boxed{s}\text{---}\mathcal{Y}\text{---}$ for some state s) as known from the theory QIT, and this immediately implies a rather long list of impossibility (‘no go’) theorems.

Let us start by answering the introductory question. It is a special case (corresponding to $K = L$) of the following more general question: If L is a (one-sided) dilation of T and K is a complete dilation of L , is then K a complete dilation of T ?

This is indeed true under suitable conditions:

Lemma 2.5.14. *(Completeness is Hereditary.)*

Suppose that Θ is complete and temporally localisable. Let $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \mathbb{E} \\ \sim \end{array}$ be a dilation of

$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}$. If $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{K} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \mathbb{G} \\ \sim \\ \mathbb{E} \\ \sim \end{array}$ is a complete dilation of L , then K is also a complete dilation of T .

Remark 2.5.15. The converse to this statement is true as well, but trivial: If K is a complete dilation of T , then it is a complete dilation of L simply for the reason that the dilations of L form a sub-class of the dilations of T . What makes the above statement interesting is that, in general, T has dilations which are not dilations of L . \spadesuit

Proof. By assumption, $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}$ has some complete dilation, say $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \hat{\mathbb{E}} \\ \sim \end{array}$, and it is enough to show that $K \triangleright \hat{T}$. Since \hat{T} is a complete dilation of T , there exists a channel G such that

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \hat{\mathbb{E}} \\ \sim \\ \mathbb{G} \\ \sim \\ \mathbb{E} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \mathbb{E} \\ \sim \end{array} . \quad (2.75)$$

Now, by temporal localisability G has a reversible dilation, say R , and since $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \hat{\mathbb{E}} \\ \sim \\ \mathbb{R} \\ \sim \\ \mathbb{Z} \end{array}$ is a dilation of L , we find a channel F such that

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \hat{\mathbb{E}} \\ \sim \\ \mathbb{R} \\ \sim \\ \mathbb{Z} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{K} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \mathbb{E} \\ \sim \\ \mathbb{G} \\ \sim \\ \mathbb{F} \\ \sim \\ \mathbb{Z} \end{array} , \quad (2.76)$$

by completeness of K . But now we can apply a left-inverse to R on both sides, and the desired relation $K \triangleright \hat{T}$ is evident. \square

Proposition 2.5.16. *(Pure Dilations are Complete.)*

Suppose that Θ is complete and temporally localisable. If $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{P} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \mathbb{E} \\ \sim \end{array}$ is a pure dilation of $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}$, then P is a complete dilation of T . In particular, if two channels $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}$ and $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{E}\text{---} \end{array}$ have a common pure dilation, then this dilation is a complete dilation of them both.

Proof. Take $K = L = P$ in Lemma 2.5.14. \square

The point of Proposition 2.5.16 is that if we now define two channels $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}$ and $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{E}\text{---} \end{array}$ to be *complementary* if they admit a common pure dilation $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{P} \\ \text{---}\mathbb{Y}\text{---} \end{array} \begin{array}{c} \sim \\ \mathbb{E} \\ \sim \end{array}$, then complementarity is ‘well-formed’ in a sense which is not obvious from the definition itself:

Theorem 2.5.17. *(Complementarity.)*

If $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}} \\ \text{---}\mathbb{E}\text{---} \end{array}$ and $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\hat{T}'} \\ \text{---}\mathbb{E}\text{---} \end{array}$ are both complementary to $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}$, then they are equivalent in the Blackwell order, i.e. there exist channels G and G' such that

$$-\mathbb{X}-\boxed{\tilde{T}'}\sim\mathbb{E}'\sim = -\mathbb{X}-\boxed{\tilde{T}}\sim\mathbb{E}\sim\boxed{G}\sim\mathbb{E}'\sim \quad \text{and} \quad -\mathbb{X}-\boxed{\tilde{T}}\sim\mathbb{E}\sim = -\mathbb{X}-\boxed{\tilde{T}'}\sim\mathbb{E}'\sim\boxed{G'}\sim\mathbb{E}\sim . \quad (2.77)$$

Moreover, complementarity is a duality, in the sense that if a channel T' is complementary to a channel complementary to T , then T' is equivalent to T in the Blackwell order.

Proof. The first statement follows from Proposition 2.5.16; the two pure dilations giving rise to \tilde{T} and \tilde{T}' are both complete, hence equivalent in the dilational order, which precisely implies equivalence of \tilde{T} and \tilde{T}' in the Blackwell order. The second statement is a consequence of the first, since complementarity is clearly a symmetric relation. \square

The point of Theorem 2.5.17 is that there is a strain, or balance, between channels $-\mathbb{X}-\boxed{T}-\mathbb{Y}-$ and their complementary channels $-\mathbb{X}-\boxed{\tilde{T}}\sim\mathbb{E}\sim$.

In the case where the theory is furthermore universal, this strain implies a generalisation of the information-disturbance duality.

Let us call a channel *completely forgetful* if it is of the form $-\mathbb{X}-\boxed{\text{tr}}\boxed{s}-\mathbb{Y}-$ for some state s . It is clear that the property of being completely forgetful is invariant under equivalence in the Blackwell order, and so is the property of being reversible. We have the following:

Theorem 2.5.18. (*Duality between Reversible and Completely Forgetful Channels.*)

Let Θ be a universal, localisable and purifiable theory. Then, a channel $-\mathbb{X}-\boxed{T}-\mathbb{Y}-$ is reversible if and only if the complementary channels $-\mathbb{X}-\boxed{\tilde{T}}\sim\mathbb{E}\sim$ are completely forgetful, and vice versa.

Proof. One might prove this by restarting the argument from the proof of Theorem 2.5.11, and we shall do so when proving its approximate generalisation (Theorem 3.4.8) in Chapter 3; however, for the sake of variation we instead use here an argument based on the statement of Theorem 2.5.11 itself.

If the channel $-\mathbb{X}-\boxed{T}-\mathbb{Y}-$ is reversible, then by Theorem 2.5.11 it is of the form $\frac{\text{---}\mathbb{X}\text{---}}{\boxed{r}\text{---}\mathbb{Z}\text{---}}\boxed{P}\text{---}\mathbb{Y}\text{---}$ for some pure channel P and some state r . Choosing a pure dilation v of r , the channel

$$\frac{\text{---}\mathbb{X}\text{---}}{\boxed{v}\text{---}\mathbb{E}\sim}\boxed{P}\text{---}\mathbb{Y}\text{---} \quad (2.78)$$

is then by localisability a pure dilation of T . Evidently, the corresponding complementary channel is then completely forgetful.

Conversely, if $-\mathbb{X}-\boxed{T}-\mathbb{Y}-$ is completely forgetful, i.e. of the form $-\mathbb{X}-\boxed{\text{tr}}\boxed{s}-\mathbb{Y}-$, then, letting v be a pure dilation of s , T has a pure dilation of the form

$$\frac{\text{---}\mathbb{X}\text{---}}{\boxed{v}\text{---}\mathbb{E}\sim}\boxed{\text{id}}\text{---}\mathbb{X}\text{---} , \quad (2.79)$$

for which the corresponding complementary channel is obviously reversible. \square

Let me conclude this section by observing that the complementarity result framed by Theorem 2.5.18 implies not only the No Broadcasting Theorem (which we have already seen based on weaker assumptions), and hence the No Cloning Theorem ([WZ82]) and the No Deletion Theorem ([PB00]), but also the more recent No Hiding ([BP07]) and No Masking ([MPSS18]) Theorems, which assert precisely the non-existence of pure channels both of whose marginals are completely forgetful.

2.6 Summary and Outlook

In this chapter, we have seen a general theory of dilations. At the heart of this theory is the *dilational ordering* (Definition 2.2.1), which in particular facilitates the notions of *complete dilations* (Definition 2.3.2) and of *spatial* and *temporal localisability* (Definition 2.3.10 and Definition 2.3.11).

These three *dilational principles* have information-theoretic interpretations, and they formally have useful consequences (e.g. they imply the DiVincenzo Property and the existence of reversible dilations). They seem to hold in all theories that are remotely physical, e.g. in the information theories **CIT** and **QIT** (as well as **QIT**[∞]) and in cartesian theories, but they are not derivable from the axioms which define a general theory, since they are violated for example in the thin theory $(\mathbb{N}, \cdot, 1, \geq)$.

A strengthening of completeness in the guise of *universality* (Definition 2.4.1) holds in all of the above complete theories, and universal dilations in principle allows us to fully characterise the dilational ordering among one-sided dilations (Proposition 2.4.8).

The theory **QIT** distinguishes itself from the theory **CIT** by being *purifiable* (Definition 2.5.1), a notion which can be seen as generalising the ‘Purification Postulate’ of Refs. [CDP10, CDP11]. The purifiability principle implies, especially in conjunction with the other dilational principles, significant ‘quantum’ features of a theory, most prominently the notion of complementarity and duality between reversible and completely forgetful channels (Theorem 2.5.18).

One conclusion of this chapter rises above all, namely: *It is possible to prove significant consequences from principles entirely about the structure of dilations.*

An obvious question for future work is whether even more properties can be captured by principles about dilations. Another question is whether there are naturally occurring examples of ‘information theories’ in which some of the dilational principles presented here fail.

Thirdly, it is reasonable to investigate whether additional impossibility theorems known from quantum information theory can be derived using the principle of purifiability. In particular, some such impossibility results regard the non-existence of certain *protocols*, for example the *No Bit Commitment Theorem* ([May97, LC97]) and the *insecurity of one-sided* ([Lo97]) and *two-sided* ([BCS12]) *two-party computation*. Proof sketches lead me to believe that such results should be possible, but a proper coverage of impossibility theorems for protocols requires introducing an array of further concepts, and they must apparently be treated with the care of a soldier traversing a minefield ([DKSW07]).

Finally, for the results of Section 2.5 to be of interest beyond the virtue of simplicity, the following question is relevant: What are the theories besides **QIT** and **QIT**[∞] which are localisable, universal and purifiable?

Chapter 3

Metric Theories

§1. Introduction and Outline – Comparison to Existing Literature.

The framework of theories as presented in Chapter 1 and Chapter 2 is purely algebraic. All of its notions, primary and derived, are defined entirely in terms of the serial and parallel compositions in a theory Θ . In this chapter, I present some initial thoughts on how to create a *metric* version of this framework. That idea is not simply one of mathematical curiosity, but is physically natural too – indeed, our two most important examples of theories, **CIT** and **QIT**, are both endowed with natural metrics, namely the *variational* (or *statistical*) *distance* (Example 3.2.7) and the *diamond-distance* (Example 3.2.8), respectively.

Whereas metric aspects of general theories have been considered before (e.g. in Ref. [CDP10]), the defined metrics seem to be always grounded in the concept of *optimal distinguishing probability*. As we have made no assumptions about probabilistic structure in the definition of a theory, an appropriate definition of metric theory also should not rely on such structure. Rather, we abstract from such ‘distinguishing metrics’ just two properties which can be phrased in terms of the serial and parallel composition in the theory. The main mantra of the ideas presented in this chapter, however, is the insistence that the metric structure comply to the architecture of dilations. This not only allows us to smoothly transfer results of Chapter 2 to an approximate setting, but also contextualises sporadic definitions and observations in the literature, particularly of Refs. [TCR10, Tom12] and Refs. [KSW08a, KSW08b].

Topological Structure. For the sake of completeness, we start in Section 3.1 by defining what one would call a *topological theory*. It is more or less obvious how to do this: Simply posit that for any $\mathcal{X}, \mathcal{Y} \in \text{Sys}_\Theta$, the set¹ of transformations from \mathcal{X} to \mathcal{Y} , $\text{Trans}_\Theta(\mathcal{X}, \mathcal{Y})$, is equipped with a topology $\mathcal{T}_{\mathcal{X}, \mathcal{Y}}$, and that the various topologies cohere such that the serial and parallel compositions are continuous maps. For example, when $\Theta = \mathbf{QIT}$ there is a natural topology on the set of quantum channels $\text{Trans}_{\mathbf{QIT}}(\mathcal{H}, \mathcal{K})$, since it is a convex subset of a finite-dimensional vector space. In general, note that a topological structure on a theory entails in particular a topology on the set of states $\text{St}(\mathcal{X}) = \text{Trans}_\Theta(\mathbf{1}, \mathcal{X})$ for any system \mathcal{X} .

For reasons to come, it is actually more appropriate to define topologies and metrics on the sets of channels, $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$, between interfaces \mathbb{X} and \mathbb{Y} , and this is what we shall

¹We assume throughout this chapter that the category Θ is *locally small*, i.e. that for any $\mathcal{X}, \mathcal{Y} \in \text{Sys}_\Theta$ the class of transformations $\text{Trans}_\Theta(\mathcal{X}, \mathcal{Y})$ is not a proper class, but merely a set. I know of no physically relevant example where this is not the case.

do (Definition 3.1.1); in principle, one could therefore have different topologies or metrics on $\text{Chan}_{\Theta}(\mathbb{X}, \mathbb{Y})$ and $\text{Chan}_{\Theta}(\mathbb{X}', \mathbb{Y}')$, even if the total systems corresponding to \mathbb{X} and \mathbb{X}' , respectively \mathbb{Y} and \mathbb{Y}' , are the same,² but in the examples given in this chapter this will never be the case.

Metric Structure and Compositional Compatibility. In defining a *metric theory*, we must clearly endow the sets of channels $\text{Chan}_{\Theta}(\mathbb{X}, \mathbb{Y})$ with metrics, $d_{\mathbb{X}, \mathbb{Y}}$, rather than topologies, $\mathcal{T}_{\mathbb{X}, \mathbb{Y}}$. This is the topic of Section 3.2. However, whereas it is obvious in the case of topologies that the correct coherence condition across various channel-topologies is to require continuity of serial and parallel composition, the ‘correct’ coherence conditions for metrics are not handed unambiguously to us – for example, it is a priori unclear whether to impose on a metric d in the case of serial composition that $d(S \circ T, \tilde{S} \circ \tilde{T}) \leq d(\tilde{S}, S) + d(\tilde{T}, T)$ or that $d(S \circ T, \tilde{S} \circ \tilde{T}) \leq \max\{d(\tilde{S}, S), d(\tilde{T}, T)\}$, or something else.

The above notion of a topological theory can be seen as a special case of categorical *enrichment* ([ML13], [Enr]), where one basically replaces the sets of morphisms by objects from a category (e.g. topological spaces of morphisms). One might therefore suspect that the ‘mathematically correct’ notion of a metric theory would similarly arise by enriching over a suitable category of metric spaces. F. W. Lawvere has argued that the most ‘natural’ category of metric spaces is that of *Lawvere metric spaces* and *short maps* between them ([Law73]); following this idea leads to the requirement $d(S \circ T, \tilde{S} \circ \tilde{T}) \leq \max\{d(\tilde{S}, S), d(\tilde{T}, T)\}$. It turns out, however, that this cannot be ‘physically correct’, since it implies for example that $d(T^n, \tilde{T}^n) \leq d(T, \tilde{T})$ for all $n \in \mathbb{N}$, for any channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{X}$ (here, T^n is the n -fold iterated serial composition), and it is easy to give examples in **CIT** with the total variation distance (or in **QIT** with the diamond-distance) for which $d(T, \tilde{T}) < 1$ whereas $d(T^n, \tilde{T}^n) \rightarrow 1$ for $n \rightarrow \infty$.³

Rather, to state the correct conditions, we take inspiration from the metrics defined by optimal distinguishing probability, which suggest the requirement $d(S \circ T, \tilde{S} \circ \tilde{T}) \leq d(\tilde{S}, S) + d(\tilde{T}, T)$ (and a similar condition for parallel composition). Metrics on Θ which adhere to the two coherence conditions for serial and parallel composition will be called *monotone* (Definition 3.2.2). Both the variational metric d_1 on **CIT** and the diamond metric d_{\diamond} on **QIT** are examples of monotone metrics. (On the other hand, the variational (trace) distance d_1 on **QIT** is not monotone, since, as is well-known, it fails to be invariant under parallel composition.) There are also more mathematical examples, however, which could not have been cast in terms of distinguishing probability, such as metrics deriving from operator norms in the cartesian theory $\mathbf{Vect}_{\mathbb{R}}$ or $\mathbf{Vect}_{\mathbb{C}}$ (Example 3.2.5).

Dilationality. Monotonicity of a metric ensures compatibility with the serial and parallel composition in the theory, but it generally fails to guarantee an affirmative answer to the following problem:⁴ Suppose that $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ are channels in a theory Θ and that d is a monotone metric on Θ ; if $d(T, \tilde{T})$ is small, we tend to think of T as being ‘close to’ \tilde{T} . Does it follow that dilations of T are close to dilations of \tilde{T} ? More precisely, given a dilation L of T , can we find a dilation \tilde{L} of \tilde{T} such that $d(L, \tilde{L}) \leq d(T, \tilde{T})$? (The inequality ‘ \geq ’ will be automatic from monotonicity.)

In Section 3.3, we will name this property *dilationality* (Definition 3.3.1), and, importantly, it does constitute an additional requirement: Whereas monotone metrics in cartesian

²And this might indeed be relevant if one were to generalise the theory to the causal setting of Chapter 4.

³For instance, take $T : \{0, 1\} \rightarrow \{0, 1\}$ to be the identity and $\tilde{T} : \{0, 1\} \rightarrow \{0, 1\}$ given by $\tilde{T}(\delta_0) = \delta_0$, $\tilde{T}(\delta_1) = \frac{1}{2}\delta_0 + \frac{1}{2}\delta_1$.

⁴The fact that we want to talk about dilations is why we want our metrics to be defined on sets of channels rather than sets of transformations.

theories and in **CIT** are always dilational cf. Example 3.3.7, ultimately due to the particular nature of complete dilations in these theories, it is known that e.g. the diamond-distance d_\diamond on **QIT** is not dilational. Thus arises the question of how we might construct dilational, monotone metrics.

The childishly naive guess of defining $\hat{d}(T, \tilde{T}) = \sup_L \inf_{\tilde{L}} d(L, \tilde{L})$ with \tilde{L} ranging over dilations of \tilde{T} and L ranging over (compatible, i.e. with same hidden interfaces) dilations of T does not necessarily produce a dilational metric \hat{d} . Indeed, \hat{d} may not only fail to be symmetric (this issue could be ignored), but, more importantly, may even fail to be dilational; sure, it was defined so as to mimic dilationality of d , but the new standard that is must be held to – namely dilationality of \hat{d} – is distinct. To better come to terms with this issue, we derive a result (Proposition 3.3.6) to the effect that dilationality of a (symmetric) metric d can in a sufficiently nice theory be reformulated as the requirement that $d(T, \tilde{T}) = \inf_{K, \tilde{K}} d(K, \tilde{K})$, where K, \tilde{K} range over complete dilations of T, \tilde{T} . This property will be termed the *Generalised Uhlmann Property*, as it imitates an infimum over Stinespring dilations, but rather using complete dilations which are more general.

Proposition 3.3.6 could seem to suggest that a definition like $\hat{d}(T, \tilde{T}) = \inf_{K, \tilde{K}} d(K, \tilde{K})$, which is now certainly symmetric, could yield a dilational metric \hat{d} , but, once again, the new standard that it must meet is rather $\hat{d}(T, \tilde{T}) = \inf_{K, \tilde{K}} \hat{d}(K, \tilde{K})$; it appears as if the potential to always dilate further is dragged along wherever we go and hindering dilationality. What we need is to *short circuit* this phenomenon, and we have already seen in Chapter 2 what such a short circuit looks like:

Purified Distances. Indeed, the trick is to use pure dilations in lifting a metric d to a dilational metric \check{d} , and this is possible in a purifiable theory, such as **QIT**, as we will see in Section 3.4. If we do this, then the self-completeness of pure dilations along with the hereditary nature of completeness (as mixed in Proposition 2.5.16) form the key to achieving dilationality. In particular, we obtain in the case of the diamond metric d_\diamond on **QIT** a metric $P_\diamond := \check{d}_\diamond$ which I will call the *purified diamond-distance* and which generalises the purified distance for states as introduced in Refs. [TCR10, Tom12].

In Section 3.4 we also prove (Theorem 3.4.8) that a dilational metric in a purifiable theory implies a tight approximate converse to Theorem 2.5.18, which can be seen as an abstract statement of the *information-disturbance trade-off* of Ref. [KSW08b]. This conceptually has nothing to do with purified distances, but seems to fit in well at that point when we have established the existence of dilational metrics in purifiable theories.

The Curious Case of QIT. In Section 3.5, we investigate the newly proposed diamond-distance P_\diamond more closely.

We start in Section 3.5.A by translating to P_\diamond the continuity result from Refs. [KSW08a, KSW08b] which is there stated in terms of the so-called *Bures distance* β (as inspired by Ref. [Bur69]), whose definition is based on a similar idea as P_\diamond . Specifically, we infer that the purified diamond-distance satisfies $d_\diamond \leq P_\diamond \leq \sqrt{2} d_\diamond$, an inequality which was known in the case of states ([TCR10, Tom12]) but which in the general case relies on the non-trivial result from Refs. [KSW08a, KSW08b] that $d_\diamond \leq \beta \leq \sqrt{2} d_\diamond$.

So, what is the difference between P_\diamond and β ? The two quantities share their most important features (cf. Remark 3.5.1), but the main argument in favour of P_\diamond is that its definition is in terms of the diamond-distance (which is operational as optimal distinguishing probability) and thus quite theory-independent, whereas β makes explicit and non-operational reference to the operator formalism of quantum information theory. In Section 3.5.B we tackle the question of determining a more quantitative relation between the two. We start by deriving

two formulas which relate P_\diamond and β to two fidelity-like quantities (Proposition 3.5.4). Using a minimax-theorem, we can then link those two fidelity-like quantities (Lemma 3.5.7), and from this ultimately derive a link between P_\diamond and β (Theorem 3.5.8). The quantitative difference between them is small for small values (in fact, $P_\diamond = \beta - \beta^3/8 + O(\beta^5)$ in the limit $\beta \rightarrow 0$), but for large values they can differ by a factor as large as $\sqrt{2}$.

§2. Contributions.

The original contributions of this chapter are the following:

- Identifying a minimal framework for the discussion of topological (Definition 3.1.1) and metric (Definition 3.2.2) aspects of general theories in the sense of Definition 1.1.6.
- Articulating the property of *dilationality* of a metric, and demonstrating its equivalence to a ‘*Generalised Uhlmann property*’ (Proposition 3.3.6).
- Presenting a recipe for the construction in purifiable theories of dilational metrics called *purified* (Theorem 3.4.3), and observing their useful behaviour with regards to complementarity (Proposition 3.4.7).
- Proving an approximate counterpart (Theorem 3.4.8) to the duality between reversible and completely forgetful channels, which can be seen as an abstract statement of the *information-disturbance trade-off* ([KSW08b]).
- Defining in the specific theory **QIT** the *purified diamond-distance* P_\diamond , which directly generalises purified distance between states of Refs. [TCR10, Tom12], and relates to the *Bures distance* of Refs. [KSW08a, KSW08b, Bur69] but distinguishes itself by a more stable definition.

3.1 Topological Structure

Let us begin by considering what would be a sensible notion of a *topological theory*. Recall that a topology on a set S is a system of subsets of S which is closed under arbitrary unions and finite intersections, and which contains the sets \emptyset and S .

Let $\text{Int}(\Theta)$ denote the class of interfaces in Θ and let $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$ denote the set⁵ of channels in Θ from the interface \mathbb{X} to the interface \mathbb{Y} .

Definition 3.1.1. (Topological Theory.)

A *topological structure* on a theory Θ is a collection $\mathcal{T} = (\mathcal{T}_{\mathbb{X}, \mathbb{Y}})_{\mathbb{X}, \mathbb{Y} \in \text{Int}(\Theta)}$, where, for each pair of interfaces \mathbb{X}, \mathbb{Y} in Θ , $\mathcal{T}_{\mathbb{X}, \mathbb{Y}}$ is a topology on the set $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$, such that the following compatibility conditions hold:

1. For all interfaces $\mathbb{X}, \mathbb{Y}, \mathbb{Z}$ in Θ , the serial composition

$$(S, T) \mapsto S \circ T \tag{3.1}$$

is continuous from $\text{Chan}_\Theta(\mathbb{Y}, \mathbb{Z}) \times \text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$ to $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Z})$ when the former is equipped with the product topology of $\mathcal{T}_{\mathbb{Y}, \mathbb{Z}}$ and $\mathcal{T}_{\mathbb{X}, \mathbb{Y}}$, and the latter with the topology $\mathcal{T}_{\mathbb{X}, \mathbb{Z}}$.

2. For all parallelly composable interfaces $\mathbb{X}_1, \mathbb{X}_2$ and $\mathbb{Y}_1, \mathbb{Y}_2$ in Θ , the parallel composition

$$(T_1, T_2) \mapsto T_1 \parallel T_2 \tag{3.2}$$

is continuous from $\text{Chan}_\Theta(\mathbb{X}_1, \mathbb{Y}_1) \times \text{Chan}_\Theta(\mathbb{X}_2, \mathbb{Y}_2)$ to $\text{Chan}_\Theta(\mathbb{X}_1 \cup \mathbb{X}_2, \mathbb{Y}_1 \cup \mathbb{Y}_2)$ when the former is equipped with the product topology of $\mathcal{T}_{\mathbb{X}_1, \mathbb{Y}_1}$ and $\mathcal{T}_{\mathbb{X}_2, \mathbb{Y}_2}$, and the latter with the topology $\mathcal{T}_{\mathbb{X}_1 \cup \mathbb{X}_2, \mathbb{Y}_1 \cup \mathbb{Y}_2}$.

A *topological theory* (Θ, \mathcal{T}) is a theory Θ together with a topological structure \mathcal{T} .⁶ ■

Plainly speaking, what we have done is to replace in the definition of a theory the sets of channels with topological spaces of channels, and to require continuity of the two modes of composition in the theory.

As mentioned in the introduction, the topologies $\mathcal{T}_{\mathbb{X}, \mathbb{Y}}$ will in all our cases really be topologies $\mathcal{T}_{\mathcal{X}, \mathcal{Y}}$ which depend only on the total systems \mathcal{X} and \mathcal{Y} associated to the interfaces \mathbb{X} and \mathbb{Y} .

Example 3.1.2. (QIT as a Topological Theory.)

The theory **QIT** has an obvious topological structure: $\text{Chan}_{\text{QIT}}(\mathbb{X}, \mathbb{Y}) \cong \text{Trans}_{\text{QIT}}(\mathcal{X}, \mathcal{Y})$ is a set of linear maps $\Lambda : \text{End}(\mathcal{X}) \rightarrow \text{End}(\mathcal{Y})$ between finite-dimensional vector spaces, and as such is naturally topologised. The continuity of the two modes of composition is clear. ◆

Example 3.1.3. (CIT as a Topological Theory.)

The theory **CIT** also has a natural topological structure; it can be most succinctly described as the subspace topology that arises from regarding **CIT** as a sub-theory of **QIT**. This topology can also be given a more concrete description in terms of the natural topology on the sets of probability distributions. ◆

⁵Recall the assumption that Θ is locally small.

⁶Strictly speaking, if we work in a foundational framework where proper classes are predicates with existence only in the metalanguage, then there is formally no such thing as the ‘pair’ (Θ, \mathcal{T}) . We ignore this point.

Example 3.1.4. (Discrete Topological Theory.)

Any theory Θ can be endowed with the *discrete topological structure*, for which the topology $\mathcal{T}_{\mathbb{X}, \mathbb{Y}}$ on $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$ is simply the discrete topology, i.e. the topology consisting of all subsets of $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$. In this case, the two continuity requirements are trivially satisfied. \blacklozenge

Example 3.1.5. (Thin Topological Theories.)

Let Θ be a thin theory. For any interfaces \mathbb{X}, \mathbb{Y} in Θ , the set $\text{Chan}_\Theta(\mathbb{X}, \mathbb{Y})$ is either empty or contains a single element, in which case it can be topologised in only one way, namely with the discrete topology. Thus, topological structure is utterly uninteresting for thin theories. \blacklozenge

Example 3.1.6. (Topologised Topology.)

If we consider a sub-theory of the cartesian theory \mathbf{Top}^* whose systems are sufficiently nice topological spaces (e.g. Hausdorff and locally compact), and if we equip the sets of continuous maps from X to Y with the *compact-open topology*, then both modes of composition are continuous and so we have a topological theory. (A sub-base for this topology is given by the collection of sets $\{f : X \rightarrow Y \mid f(K) \subseteq U\}$ with $K \subseteq X$ compact and $U \subseteq Y$ open.) \blacklozenge

As a warm-up to metric theories, it is worth observing that, by properties of continuous maps, in a topological theory the map $T \mapsto S_0 \circ T$ is continuous for any fixed channel S_0 , as is the map $S \mapsto S \circ T_0$ for any fixed T_0 . It follows from this that also the swapping map $T_1 \parallel T_2 \mapsto T_2 \parallel T_1$ is continuous, since $T_2 \parallel T_1 = \sigma_{\mathbb{Y}_1, \mathbb{Y}_2} \circ (T_1 \parallel T_2) \circ \sigma_{\mathbb{X}_2, \mathbb{X}_1}$ cf. Definition 1.1.2. In short, all sensible operations involving the algebraic operations in Θ are continuous.

Let us to also note that, under condition 1., condition 2. in Definition 3.1.1 can be weakened to continuity, for all fixed \mathbb{Z}_2 , of the map $T_1 \mapsto T_1 \parallel \text{id}_{\mathbb{Z}_2}$. Indeed, from this follows continuity of $T_2 \mapsto \text{id}_{\mathbb{Z}_1} \parallel T_2$ like above, and general parallel composition can then be realised by means of a serial composition, as $T_1 \parallel T_2 = (T_1 \parallel \text{id}_{\mathbb{Y}_2}) \circ (\text{id}_{\mathbb{X}_1} \parallel T_2)$.

3.2 Metric Structure

Though it is possible and potentially beneficial to work with very general kinds of metrics (e.g. those in the sense of Lawvere, which need neither be symmetric nor non-degenerate), we shall content ourselves with the usual definition of a metric:

Definition 3.2.1. (Metric.)

Let S be a set. A map $d : S \times S \rightarrow [0, \infty]$ is called a *metric on S* if it satisfies the following four requirements:

- (Nullity.) For all $x \in S$, $d(x, x) = 0$
- (Non-Degeneracy.) For all $x, y \in S$, $d(x, y) = d(y, x) = 0$ implies $x = y$
- (Symmetry.) For all $x, y \in S$, $d(x, y) = d(y, x)$.
- (Triangle Inequality.) For all $x, y, z \in S$, $d(x, y) \leq d(x, z) + d(z, y)$.

(By symmetry, non-degeneracy can be weakened to $d(x, y) = 0 \Rightarrow x = y$.) \blacksquare

Most of our metrics will moreover be *normalised*, meaning that $\sup_{x, y \in S} d(x, y) = 1$ (for example, this is the case for the metrics derived from optimal distinguishing probability), but we need not impose this.

Definition 3.2.2. (Monotone Metrics on a Theory.)

A *metric* on a theory Θ is a collection $d = (d_{\mathbb{X}, \mathbb{Y}})_{\mathbb{X}, \mathbb{Y} \in \text{Int}(\Theta)}$, where, for each pair of interfaces \mathbb{X}, \mathbb{Y} in Θ , $d_{\mathbb{X}, \mathbb{Y}}$ is a metric on the set $\text{Chan}_{\Theta}(\mathbb{X}, \mathbb{Y})$. A metric is called

1. *serially monotone*, if

$$d_{\mathbb{X}, \mathbb{Z}}(S \circ T, \tilde{S} \circ \tilde{T}) \leq d_{\mathbb{X}, \mathbb{Y}}(T, \tilde{T}) + d_{\mathbb{Y}, \mathbb{Z}}(S, \tilde{S}) \quad (3.3)$$

for all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ and $S, \tilde{S} : \mathbb{Y} \rightarrow \mathbb{Z}$;

2. *parallelly invariant*, if

$$d_{\mathbb{X} \cup \mathbb{Z}, \mathbb{Y} \cup \mathbb{W}}(T \parallel T_0, \tilde{T} \parallel T_0) = d_{\mathbb{X}, \mathbb{Y}}(T, \tilde{T}) \quad (3.4)$$

for all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ and $T_0 : \mathbb{Z} \rightarrow \mathbb{W}$.

A metric is called *monotone* if it is both serially monotone and parallelly invariant.⁷

A *metric theory* (Θ, d) is a theory Θ together with a monotone metric d . ■

Remark 3.2.3. (Notational Convention I.)

From now on, we always omit the subscripts on d indicating the interfaces, and write e.g. $d(T, \tilde{T})$ rather than $d_{\mathbb{X}, \mathbb{Y}}(T, \tilde{T})$ for channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$. ✚

Remark 3.2.4. (Notational Convention II.)

We will in this chapter move quite liberally between the algebraic notation in terms of operations ‘ \circ ’ and ‘ \parallel ’ and the pictorial notation, depending on the situation at hand. For reference, note that the conditions of serial monotonicity and parallel invariance are pictorially stated as

$$d(\begin{array}{c} \boxed{-\mathbb{X}-\boxed{T}-\mathbb{Y}-\boxed{S}-\mathbb{Z}-} \\ \boxed{-\mathbb{X}-\boxed{\tilde{T}}-\mathbb{Y}-\boxed{\tilde{S}}-\mathbb{Z}-} \end{array} , \begin{array}{c} \boxed{-\mathbb{X}-\boxed{T}-\mathbb{Y}-} \\ \boxed{-\mathbb{X}-\boxed{\tilde{T}}-\mathbb{Y}-} \end{array}) \leq d(\begin{array}{c} \boxed{-\mathbb{X}-\boxed{T}-\mathbb{Y}-} \\ \boxed{-\mathbb{Y}-\boxed{S}-\mathbb{Z}-} \end{array} , \begin{array}{c} \boxed{-\mathbb{X}-\boxed{\tilde{T}}-\mathbb{Y}-} \\ \boxed{-\mathbb{Y}-\boxed{\tilde{S}}-\mathbb{Z}-} \end{array}) \quad (3.5)$$

respectively

$$d(\begin{array}{c} \boxed{-\mathbb{X}-\boxed{T}-\mathbb{Y}-} \\ \boxed{-\mathbb{Z}-\boxed{T_0}-\mathbb{W}-} \end{array} , \begin{array}{c} \boxed{-\mathbb{X}-\boxed{\tilde{T}}-\mathbb{Y}-} \\ \boxed{-\mathbb{Z}-\boxed{T_0}-\mathbb{W}-} \end{array}) = d(\begin{array}{c} \boxed{-\mathbb{X}-\boxed{T}-\mathbb{Y}-} \\ \boxed{-\mathbb{X}-\boxed{\tilde{T}}-\mathbb{Y}-} \end{array}). \quad (3.6)$$

✚

Like for topological theories, it will in our examples always be the case that the metric is really defined on the set of transformations, rather than channels.

As such, we will see shortly that the *variational distance* and the *diamond-distance* are monotone metrics on the theories **CIT** and **QIT**, respectively; this can be realised already now by the acquainted reader, since both of these metrics are defined using optimal distinguishing probability. However, the conditions of serial monotonicity and parallel invariance abstract away just the most essential features of such metrics, and they consequently accommodate examples which have nothing to do with distinguishing probability:

⁷The word ‘monotone’ seems to not do justice to the requirement of parallel invariance, but that requirement is in fact equivalent (in the presence of serial monotonicity) to the ‘monotonicity’ requirement that results from replacing in Eq. (3.4) the equality with ‘ \leq ’, provided the theory has states on every system. We shall use the term generally even so.

Example 3.2.5. (Operator Norms.)

Recall from Example 1.2.8 the cartesian theory whose systems are vector spaces over k and whose transformations are k -linear operators, composing serially by functional composition and parallelly by the direct sum, \oplus . Suppose that $k = \mathbb{R}$ or $k = \mathbb{C}$, and consider instead the cartesian theory $\mathbf{Vect}_k^{\|\cdot\|}$ whose systems are normed vector spaces and whose transformations $A : (V, \|\cdot\|_V) \rightarrow (W, \|\cdot\|_W)$ are linear operators with operator norm $\|A\|_\infty \leq 1$. The serial and parallel compositions are as in \mathbf{Vect}_k , with the additional detail that the norm on $V_1 \oplus V_2$ is given by $\max\{\|\cdot\|_{V_1}, \|\cdot\|_{V_2}\}$.

Now, consider the metric d given on the set of transformations from $(V, \|\cdot\|_V)$ to $(W, \|\cdot\|_W)$ by $d(A, \tilde{A}) = \|A - \tilde{A}\|_\infty$. It is easy to verify that d is serially monotone (using the fact that all allowed transformations have operator norm at most 1) and that d is parallelly invariant (since $A \oplus A_0 - \tilde{A} \oplus A_0 = (A - \tilde{A}) \oplus 0$), so d defines a monotone metric on the theory $\mathbf{Vect}_k^{\|\cdot\|}$. \blacklozenge

The following observation is often helpful in demonstrating that a given metric is serially monotone and parallelly invariant:

Lemma 3.2.6. (*Recharacterisation of Monotone Metrics.*)

A metric d on Θ is monotone (i.e. serially monotone and parallelly invariant) if and only if the following three conditions hold:

- d is **monotone under serial pre-composition**, meaning that

$$d(S_0 \circ T, S_0 \circ \tilde{T}) \leq d(T, \tilde{T}) \quad (3.7)$$

for all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ and $S_0 : \mathbb{Y} \rightarrow \mathbb{Z}$;

- d is **monotone under serial post-composition**, meaning that

$$d(S \circ T_0, \tilde{S} \circ T_0) \leq d(S, \tilde{S}) \quad (3.8)$$

for all channels $S, \tilde{S} : \mathbb{Y} \rightarrow \mathbb{Z}$ and $T_0 : \mathbb{X} \rightarrow \mathbb{Y}$;

- d is **parallelly invariant under identities**, meaning that

$$d(T \parallel \text{id}_{\mathbb{Z}}, \tilde{T} \parallel \text{id}_{\mathbb{Z}}) = d(T, \tilde{T}) \quad (3.9)$$

for all interfaces \mathbb{Z} and all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$.

Proof. To see that all three conditions must hold if d is a monotone metric, simply observe that they are special instances of the general conditions (e.g. letting $S = \tilde{S} = S_0$ and using nullity of the metric).

Conversely, if d is monotone under serial pre- and post-composition then

$$\begin{aligned} d(S \circ T, \tilde{S} \circ \tilde{T}) &\leq d(S \circ T, \tilde{S} \circ T) + d(\tilde{S} \circ T, \tilde{S} \circ \tilde{T}) \\ &\leq d(S, \tilde{S}) + d(T, \tilde{T}) \end{aligned} \quad (3.10)$$

by the triangle inequality, so d is monotone under serial composition. If in addition d is invariant under parallel identities, then

$$\begin{aligned} d(T_1 \parallel T_2, \tilde{T}_1 \parallel \tilde{T}_2) &= d([T_1 \parallel \text{id}_{\mathbb{Y}_2}] \circ [\text{id}_{\mathbb{X}_1} \parallel T_2], [\tilde{T}_1 \parallel \text{id}_{\mathbb{Y}_2}] \circ [\text{id}_{\mathbb{X}_1} \parallel \tilde{T}_2]) \\ &\leq d(T_1 \parallel \text{id}_{\mathbb{Y}_2}, \tilde{T}_1 \parallel \text{id}_{\mathbb{Y}_2}) + d(\text{id}_{\mathbb{X}_1} \parallel T_2, \text{id}_{\mathbb{X}_1} \parallel \tilde{T}_2) \\ &= d(T_1, \tilde{T}_1) + d(T_2, \tilde{T}_2). \end{aligned} \quad (3.11)$$

□

Whereas the conditions of Lemma 3.2.6 are often easier to verify than those of Definition 3.2.2 (the reader may revisit Example 3.2.5 in this light), a further simplification in the form of Lemma 3.2.9 more clearly yields the argument that the usual metrics in **CIT** and **QIT** are monotone. Before observing this result, however, let us recall the definition of these metrics (for more details, the reader may consult Refs. [NC02, Wat]):

Example 3.2.7. (The Variational Distance d_1 on **CIT**.)

Consider the theory **CIT**. On $\text{St}(Y)$, the set of probability distributions on Y , let d_1 be the metric given by

$$d_1(p, q) = \frac{1}{2} \sum_{y \in Y} |p(y) - q(y)|. \quad (3.12)$$

The metric d_1 is called the *total variation distance*, or simply the *variational distance*. It can be shown that $d_1(p, q) = \max_{A \subseteq Y} |p(A) - q(A)|$, where we use the notation $r(A)$ as shorthand for $\sum_{y \in A} r(y)$, and as such $d_1(p, q)$ quantifies the largest possible difference in probability that p and q assign to any subset of Y . It can also be shown that $\frac{1+d_1(p, q)}{2}$ is the optimal probability of distinguishing between the probability distributions p and q based on an outcome drawn according to either p or q , with equal probability.

We can easily extend d_1 to a metric on general classical channels, by defining, for $T = (t_x)_{x \in X} : X \rightarrow Y$ and $\tilde{T} = (\tilde{t}_x)_{x \in X} : X \rightarrow Y$,

$$d_1(T, \tilde{T}) = \max_{x \in X} d_1(t_x, \tilde{t}_x). \quad (3.13)$$

As such, d_1 provides a metric on the theory **CIT**. It is not eye-catching that d_1 is serially monotone and parallelly invariant (or, equivalently, satisfies the conditions of Lemma 3.2.6). A convexity argument shows however that $d_1(T, \tilde{T}) = \sup_{p \in \text{St}(X)} d_1(T \circ p, \tilde{T} \circ p)$, and more generally that $d_1(T, \tilde{T}) = \sup_{p \in \text{St}(X \times R)} d_1([T \otimes \text{id}_R] \circ p, [\tilde{T} \otimes \text{id}_R] \circ p)$ for any system R , which implies one of the conditions stated in our next result, Lemma 3.2.9; as the other condition can also quite easily be demonstrated, that result will imply the desired. Thus, (**CIT**, d_1) is a metric theory. ♦

Example 3.2.8. (The Diamond-Distance d_\diamond on **QIT**.)

Consider the theory **QIT**. On $\text{St}(\mathcal{K}) = \mathcal{D}(\mathcal{K})$, the set of density matrices on \mathcal{K} , let d_1 be the metric given by

$$d_1(\varrho, \sigma) = \frac{1}{2} \|\varrho - \sigma\|_1, \quad (3.14)$$

where $\|\cdot\|_1$ denotes the trace-norm on $\text{End}(\mathcal{K})$. This metric, called *trace distance*, is denoted by the same symbol as the variational distance in Example 3.2.7, since under the natural embedding of **CIT** in **QIT** it is easily seen to reproduce d_1 as defined in **CIT**. It can be shown that also in **QIT** the quantity $\frac{1+d_1(\varrho,\sigma)}{2}$ is the optimal probability of distinguishing between the states ϱ and σ if one of these is given at random with equal probability.

If we were to define the trace-distance between general quantum channels $\Lambda, \tilde{\Lambda} : \mathcal{H} \rightarrow \mathcal{K}$ as $d_1(\Lambda, \tilde{\Lambda}) = \sup_{\varrho \in \text{St}(\mathcal{H})} d_1(\Lambda_1(\varrho), \tilde{\Lambda}_1(\varrho))$, then d_1 would be a metric on **QIT** which is serially monotone. As is well-known, however, this metric is not parallelly invariant.⁸ The fix is to define for channels instead the *diamond-distance* d_\diamond given by

$$d_\diamond(\Lambda, \tilde{\Lambda}) = \sup_{\mathcal{R} \in \text{Sys}_{\mathbf{QIT}}} \sup_{\varrho \in \text{St}(\mathcal{H} \otimes \mathcal{R})} d_1([\Lambda \otimes \text{id}_{\mathcal{R}}](\varrho), [\tilde{\Lambda} \otimes \text{id}_{\mathcal{R}}](\varrho)) \quad (3.15)$$

for $\Lambda, \tilde{\Lambda} : \mathcal{H} \rightarrow \mathcal{K}$. Again, by Lemma 3.2.9, d_\diamond will be serially monotone and parallelly invariant, so we have a metric theory (**QIT**, d_\diamond). ◆

The above two examples illustrate that naturally occurring metrics in a theory are often constructed on the basis of metrics on the sets of states. The following result will be applicable in such circumstances:

Lemma 3.2.9. (*Sufficient Conditions for Monotonicity.*)

Suppose that d is a metric on Θ which is **state-determined**, meaning that for all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$,

$$d(T, \tilde{T}) = \sup_{\mathcal{R}, s} d \left(\begin{array}{c} \boxed{s} \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \\ \text{---} \mathcal{R} \text{---} \end{array}, \begin{array}{c} \boxed{s} \text{---} \mathbb{X} \text{---} \boxed{\tilde{T}} \text{---} \mathbb{Y} \text{---} \\ \text{---} \mathcal{R} \text{---} \end{array} \right). \quad (3.16)$$

Then, d is monotone (i.e. serially monotone and parallelly invariant) if and only if d is **state-monotone**, meaning that

$$d \left(\begin{array}{c} \boxed{r} \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \\ \text{---} \mathcal{R} \text{---} \end{array}, \begin{array}{c} \boxed{\tilde{r}} \text{---} \mathbb{X} \text{---} \boxed{T} \text{---} \mathbb{Y} \text{---} \\ \text{---} \mathcal{R} \text{---} \end{array} \right) \leq d \left(\begin{array}{c} \boxed{r} \text{---} \mathbb{X} \text{---} \\ \text{---} \mathcal{R} \text{---} \end{array}, \begin{array}{c} \boxed{\tilde{r}} \text{---} \mathbb{X} \text{---} \\ \text{---} \mathcal{R} \text{---} \end{array} \right) \quad (3.17)$$

for all channels $T : \mathbb{X} \rightarrow \mathbb{Y}$ and all states r, \tilde{r} on \mathbb{X} .

Remark 3.2.10. (Metrics form Distinguishing Probability.)

Any metric which like d_1 and d_\diamond quantifies optimal distinguishing probability will be (quite evidently) state-determined and state-monotone. Therefore, any such metric is a monotone metric in the sense of Definition 3.2.2. ✂

Proof. It is clear that state-monotonicity follows from general serial monotonicity. Assume conversely that d is state-determined and state-monotone. We need to verify the conditions in Lemma 3.2.6, but it suffices to prove parallel monotonicity rather than parallel invariance, since there are states on every system in Θ .

The monotonicity under parallel composition with identities is built into the condition of being state-determined. So is monotonicity under serial post-composition, since the supremum is after composition with T_0 confined to a smaller set. Finally, monotonicity under serial pre-composition follows from d being both state-determined and state-monotone. □

⁸See e.g. Ref. [Wat], where the diamond norm is referred to as the *completely bounded* norm.

The fact that d_1 on **CIT** and d_\diamond on **QIT** are state-monotone are proved in most introductory texts, see e.g. Ref. [NC02]. (Alternatively, the reader might find the demonstration to be a nice exercise.) Thus, we confirm that indeed d_1 is a monotone metric on **CIT**, and d_\diamond a monotone metric on **QIT**.

3.3 Dilationality and the Generalised Uhlmann Property

Thus far, we have stated a general definition of metric theories and provided two important examples, (**CIT**, d_1) and (**QIT**, d_\diamond). The usefulness of the monotonicity properties is likely well-known to anyone who has experience with these two examples, but one additional property of metrics is desirable, and that property has do with dilations:

Suppose that $\text{-}\mathbb{X}\text{-}\boxed{T}\text{-}\mathbb{Y}\text{-}$ and $\text{-}\mathbb{X}\text{-}\boxed{\tilde{T}}\text{-}\mathbb{Y}\text{-}$ are channels in Θ , and that d is a monotone metric on Θ . If $\begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{L}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array}$ and $\begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{\tilde{L}}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array}$ are *compatible* (meaning with the same hidden interfaces) dilations of T and \tilde{T} respectively, then

$$\begin{aligned} d(\text{-}\mathbb{X}\text{-}\boxed{T}\text{-}\mathbb{Y}\text{-}, \text{-}\mathbb{X}\text{-}\boxed{\tilde{T}}\text{-}\mathbb{Y}\text{-}) &= d\left(\begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{T}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array}, \begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{\tilde{T}}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array}\right) \\ &= d\left(\begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{L}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array} \begin{array}{c} \boxed{\text{tr}} \\ \sim\mathbb{E}\sim \end{array}, \begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{\tilde{L}}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array} \begin{array}{c} \boxed{\text{tr}} \\ \sim\mathbb{E}\sim \end{array}\right) \\ &\leq d\left(\begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{L}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array}, \begin{array}{c} \sim\mathbb{D}\sim \\ \sim\mathbb{X}\text{-}\boxed{\tilde{L}}\text{-}\mathbb{Y}\text{-} \\ \sim\mathbb{E}\sim \end{array}\right) \end{aligned} \quad (3.18)$$

by parallel invariance and serial monotonicity of d . In other words, *distance never decreases under dilations*. Of course, even when T and \tilde{T} are identical, the dilations L and \tilde{L} might be far from each other; what is desirable, however, is that for any dilation L of T we can find some compatible dilation \tilde{L} of \tilde{T} such that $d(T, \tilde{T}) = d(L, \tilde{L})$.

Somewhat more forgivingly, we may ask that merely $d(T, \tilde{T}) = \inf_{\tilde{L}} d(L, \tilde{L})$, where the infimum is over dilations of \tilde{T} compatible with L . The universal quantification over L can be then replaced by a supremum for a closed statement:

Definition 3.3.1. (Dilationality.)

A monotone metric d on Θ is called *dilational* if for all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$, it holds that

$$d(T, \tilde{T}) = \sup_{L \in \text{Dil}(T)} \inf_{\tilde{L} \in \text{Dil}(\tilde{T})} d(L, \tilde{L}), \quad (3.19)$$

where the inner infimum is over dilations \tilde{L} *compatible with* L , i.e. with the same hidden interfaces as L . ■

Example 3.3.2. (Dilationality of the Discrete Metric.)

Consider on any theory Θ the discrete metric d_0 , for which $d_0(T, \tilde{T}) = 1$ if $T \neq \tilde{T}$. The metric d_0 is trivially dilational. ◆

Before presenting more interesting examples of dilational metrics, we demonstrate a number of properties surrounding dilationality. (Though I encourage the reader already at this point to ponder why a monotone metric in a cartesian theory, or in **CIT**, is necessarily dilational.)

First, let us observe the following useful fact (which has nothing to do with dilationality):

Lemma 3.3.3. (*Monotonicity of Derivability.*)

Let d be a monotone metric on Θ . Suppose that $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ are channels, and let $L \in \text{Dil}(T)$ and $\tilde{L} \in \text{Dil}(\tilde{T})$ be compatible one-sided dilations. For any dilation $L' \trianglelefteq L$ of T , there exists a compatible dilation $\tilde{L}' \trianglelefteq \tilde{L}$ such that $d(L', \tilde{L}') \leq d(L, \tilde{L})$.

Proof. Since $L' \trianglelefteq L$, we find a channel G such that $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}'} \\ \boxed{L'} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}'} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}'} \\ \boxed{L} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}'} \\ \boxed{G} \\ \sim_{\mathbb{E}'} \end{array} .$

But then defining \tilde{L}' to be $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}'} \\ \boxed{\tilde{L}} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}'} \\ \boxed{G} \\ \sim_{\mathbb{E}'} \end{array}$, we have $d(L', \tilde{L}') \leq d(L, \tilde{L})$ by serial and parallel monotonicity of d . \square

Among other things, Lemma 3.3.3 has as consequence that the supremum in the condition (3.19) can be restricted to one-sided dilations if every dilation is derivable from a one-sided dilation:

Proposition 3.3.4. (*Restricting to One-Sided Dilations.*)

Suppose that Θ has the DiVincenzo Property (e.g. by being spatially localisable). Then, a monotone metric d on Θ is dilational precisely if for all channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$, and for all one-sided dilations L of T , it holds that

$$d(T, \tilde{T}) = \inf_{\tilde{L}} d(L, \tilde{L}), \quad (3.20)$$

where the infimum is over all one-sided dilations \tilde{L} of \tilde{T} which are compatible with L .

Proof. Obvious from Lemma 3.3.3. \square

We will often use this result more or less implicitly.

Remark 3.3.5. (*On the Importance of Dilationality for the Establishment of Dilations.*)

One of the useful consequences of dilationality has to do with the very establishment of the concept of a dilation.

In the operational narrative, we have pretended that by interaction with a pair of open interfaces \mathbb{X} and \mathbb{Y} we can certify exactly the occurrence of a given channel $\text{---}\mathbb{X}\text{---}\boxed{T_0}\text{---}\mathbb{Y}\text{---}$. In reality, we can of course, even under an i.i.d. assumption, only verify that we interact with a channel T close to T_0 . Now, the agent (or, ‘Nature’) who implements T is really employing some channel $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}} \\ \boxed{M} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}} \end{array}$ between the interfaces accessible to us and additional interfaces in the environment which are inaccessible to us. Not only have we pretended that $T = T_0$, we have also pretended that M must in fact be a dilation of T_0 . In reality, however, the channel we ‘see’ when M is employed is the channel $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}} \\ \boxed{M} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}} \text{---}\boxed{\text{tr}} \end{array}$, and we control only the input to the interface \mathbb{X} . By doing this, we may (under an i.i.d. assumption) certify that

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}} \\ \boxed{M} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}} \text{---}\boxed{\text{tr}} \end{array} \approx_{\varepsilon}^d \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim_{\mathbb{D}} \\ \boxed{T_0} \\ \text{---}\mathbb{Y}\text{---} \\ \sim_{\mathbb{E}} \text{---}\boxed{\text{tr}} \end{array} , \quad (3.21)$$

where \approx_ε^d signifies that the two channels are at most ε apart w.r.t. d . Hence, what we actually need for all of our mathematical pretendings to be operationally justifiable is that T_0 have a two-sided dilation which is at most ε apart from M w.r.t. d .

Dilationality hands us precisely this feature: If d is dilational, then according to Eq. (3.21), some one-sided dilation of $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T_0} \\ \text{---}\mathbb{Y}\text{---} \end{array}$ will be ε -close to M , and this dilation will be a dilation of T_0 when including the input interface \mathbb{D} .

(It is worth observing that another way of phrasing this discussion is by saying that dilationality guarantees that a channel which is almost non-signalling is close to a channel which is perfectly non-signalling.)

✦

Now, it is known that the dilational property fails for the diamond-distance d_\diamond in **QIT** already in the case of distances between states. This is one of the reasons for the introduction of the ‘purified distance’ ([TCR10, Tom12]), and will ultimately also be the reason for us to introduce a generalisation of this distance to arbitrary channels in **QIT**, the *purified diamond-distance* (Definition 3.4.2).

In contrast, the dilational property is satisfied for the variational d_1 on **CIT**, and will in fact be satisfied for any monotone metric on **CIT**. To better understand the mechanism behind this, it is helpful to observe the following recharacterisation of dilationality, which is also of independent interest:

Proposition 3.3.6. (*Dilationality means the Generalised Uhlmann Property.*)

Suppose that Θ is complete and localisable, and suppose that d is a monotone metric on Θ . Then, the following are equivalent:

1. The metric d is dilational, that is, for any channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$, and any one-sided dilation L of T , it holds that

$$d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\tilde{T}} \\ \text{---}\mathbb{Y}\text{---} \end{array}\right) = \inf_{\tilde{L}} d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{L} \\ \text{---}\mathbb{Y}\text{---} \end{array} \overset{\sim}{\mathbb{E}} \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\tilde{L}} \\ \text{---}\mathbb{Y}\text{---} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\tilde{L}} \\ \text{---}\mathbb{Y}\text{---} \end{array}\right), \quad (3.22)$$

where the infimum is over all dilations \tilde{L} of \tilde{T} , compatible with L .

2. The metric d has the Generalised Uhlmann Property, that is, for any channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$, it holds that

$$d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\tilde{T}} \\ \text{---}\mathbb{Y}\text{---} \end{array}\right) = \inf_{K, \tilde{K}} d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{K} \\ \text{---}\mathbb{Y}\text{---} \end{array} \overset{\sim}{\mathbb{E}_0} \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\tilde{K}} \\ \text{---}\mathbb{Y}\text{---} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\tilde{K}} \\ \text{---}\mathbb{Y}\text{---} \end{array}\right), \quad (3.23)$$

where the infimum is over all compatible pairs of complete dilations K of T and \tilde{K} of \tilde{T} .

Proof. By the monotonicity properties, the inequalities ‘ \leq ’ are clear in both (3.22) and (3.23). Hence, for each implication, it suffices to show the inequality ‘ \geq ’.

First assume that d has the Generalised Uhlmann Property. Let $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ be channels and let L be a dilation of T . Given $\varepsilon > 0$, pick, according to the Generalised Uhlmann Property, compatible complete dilations K, \tilde{K} such that

$$d(K, \tilde{K}) < d(T, \tilde{T}) + \varepsilon. \quad (3.24)$$

Now, since by completeness $L \sqsubseteq K$, it follows from Lemma 3.3.3 that

$$\inf_{\tilde{L}} d(L, \tilde{L}) \leq d(K, \tilde{K}) < d(T, \tilde{T}) + \varepsilon, \quad (3.25)$$

where the infimum is over dilations \tilde{L} of \tilde{T} , compatible with L . Since $\varepsilon > 0$ was arbitrary the desired inequality follows.

The converse implication is more interesting. Suppose that d is dilational, and let $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ be channels. Now, choose some complete dilation K_0 of T . Given $\varepsilon > 0$, pick, according to dilationality, a dilation \tilde{L} of \tilde{T} such that

$$d(K_0, \tilde{L}) < d(T, \tilde{T}) + \varepsilon. \quad (3.26)$$

Now, by symmetry of d we may switch the order of arguments in the condition of dilationality. Therefore, if we pick a complete dilation \tilde{K} of \tilde{L} , we find by dilationality some dilation K of K_0 , compatible with \tilde{K} , such that

$$d(K, \tilde{K}) < d(K_0, \tilde{L}) + \varepsilon. \quad (3.27)$$

Combining Eqs. (3.26) and (3.27) yields

$$d(K, \tilde{K}) < d(T, \tilde{T}) + 2\varepsilon, \quad (3.28)$$

and so we are done if we can argue that K and \tilde{K} are complete dilations of T and \tilde{T} , respectively.

However, K dilates K_0 which was chosen complete for T , and therefore K is trivially a complete dilation of T . More interestingly, \tilde{K} is a complete dilation of \tilde{L} , which dilates \tilde{T} , and is therefore by the fact that completeness is hereditary (Lemma 2.5.14) a complete dilation of \tilde{T} . The desired follows again since $\varepsilon > 0$ was arbitrary. \square

Example 3.3.7. (Dilationality in **CIT**.)

In **CIT**, a complete dilation K_T of T can be canonically obtained by copying inputs and outputs (Theorem 2.3.21). Since K_T is given from T by tensoring with identities and serially composing with surrounding channels, we must have $d(T, \tilde{T}) = d(K_T, K_{\tilde{T}})$ for any monotone metric d . Consequently, d trivially has the Generalised Uhlmann Property, and d is therefore necessarily dilational. In particular, the variational distance d_1 on **CIT** is dilational. \blacklozenge

By a completely similar argument, any monotone metric on a cartesian theory is dilational.

In general, however, complete dilations are not computable simply by pre- and post-composing cleverly with channels – for example, this is not the case in **QIT** where completeness goes through Stinespring dilations. A natural question is thus the following: Given a monotone metric d on a theory Θ , can we construct from d a monotone metric \hat{d} which is also dilational?

Proposition 3.3.6 suggests that the quantity

$$\hat{d}(T, \tilde{T}) = \inf_{K, \tilde{K}} d(T, \tilde{T}) \quad (3.29)$$

could be a sensible candidate (where the infimum ranges over compatible pairs of complete dilations). Though \hat{d} can be proven to be a metric under circumstances which ensure the triangle inequality, and proven to be also monotone, it seems to me that (as mentioned in the introduction) dilationality could well fail for \hat{d} , although I know of no specific counterexample.

Open Problem 3.3.8. (*Construction of Dilational Metrics.*)

Suppose that Θ is complete and localisable and that d is a monotone metric on Θ . Is the quantity \hat{d} given by Eq. (3.29) a monotone dilational metric on Θ ? If not, are there tractable additional conditions which ensure that it is?

3.4 Purified Dilationality – the Uhlmann Property

In this section, we shall see that purifiability allows us to squeeze out a dilational, monotone metric starting from one which is just monotone. This is expressed by Theorem 3.4.3. We will also prove an approximate version of the duality of Theorem 2.5.18 which applies to any dilational monotone metric in a purifiable theory.

More precisely, let us assume from now on that the theory Θ is localisable, universal and purifiable. In particular, every channel T in Θ has a one-sided pure dilation Σ by Proposition 2.5.10. (We will in the remainder of the chapter use the notation ‘ Σ ’ rather than ‘ P ’ for pure dilations, since P will be soon used for something else; in **QIT**, the pure dilations Σ are Stinespring dilations.)

Let us also make the technical assumption that if $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$ are channels, then there exist pure one-sided dilations $\Sigma, \tilde{\Sigma} : \mathbb{X} \rightarrow \mathbb{Y} \cup \mathbb{E}$ which are compatible, i.e. have the same hidden interface. This is true, for example, in **QIT**, and more generally whenever the theory Θ has a pure state on every system.⁹

We make the following definition:

Definition 3.4.1. (Purified Distance.)

Let d be a monotone metric on Θ . We define the *purified d -distance* between channels $-_{\mathbb{X}}\boxed{T}\text{-}_{\mathbb{Y}-}$ and $-_{\mathbb{X}}\boxed{\tilde{T}}\text{-}_{\mathbb{Y}-}$ by

$$\check{d}(-_{\mathbb{X}}\boxed{T}\text{-}_{\mathbb{Y}-}, -_{\mathbb{X}}\boxed{\tilde{T}}\text{-}_{\mathbb{Y}-}) = \inf_{\Sigma, \tilde{\Sigma}} d\left(-_{\mathbb{X}}\boxed{\Sigma}\text{-}_{\mathbb{Y}-}^{\sim\mathbb{E}\sim}, -_{\mathbb{X}}\boxed{\tilde{\Sigma}}\text{-}_{\mathbb{Y}-}^{\sim\mathbb{E}\sim}\right), \quad (3.30)$$

where the infimum is over all compatible pure one-sided dilations Σ of T and $\tilde{\Sigma}$ of \tilde{T} . ■

⁹If $\Sigma_0 : \mathbb{X} \rightarrow \mathbb{Y} \cup \mathbb{E}_0$ is a pure dilation of T and $\tilde{\Sigma}_0 : \mathbb{X} \rightarrow \mathbb{Y} \cup \tilde{\mathbb{E}}_0$ a pure dilation of \tilde{T} , then we obtain compatible pure dilations $\Sigma, \tilde{\Sigma} : \mathbb{X} \rightarrow \mathbb{Y} \cup \mathbb{E}_0 \cup \tilde{\mathbb{E}}_0$ by parallelly composing with any pure states on the appropriate systems.

Observe that in the special case where $\Theta = \mathbf{QIT}$, the infimum is over compatible Stinespring dilations of the channels. Taking in this case $d = d_\diamond$ (the diamond-distance) is our most important example, and we will give it a special name:

Definition 3.4.2. (Purified Diamond-Distance.)

The *purified diamond-distance in QIT* is the quantity $P_\diamond := \check{d}_\diamond$. ■

In Section 3.5 we shall discuss how to calculate $P_\diamond(T, \tilde{T})$ and relate it to the so-called *Bures distance* of Refs. [KSW08a, Bur69]. For the moment, let us simply observe that in the case where the channels are states, the quantity $P_\diamond(\varrho, \tilde{\varrho})$ is already known in the literature as the *purified distance* ([TCR10, Tom12]).

Right now, we are confronted with the much more pressing problem of demonstrating that in general the purified distance \check{d} defines a monotone metric, so that indeed its name is sensible. Moreover – and this was the reason for its introduction – it is dilational:

Theorem 3.4.3. (*Properties of the Purified Distance.*)

Assume that Θ is localisable, universal and purifiable.¹⁰ For any monotone metric d on Θ , the purified distance \check{d} defines a monotone, dilational metric on Θ . We moreover have the inequality $d(T, \tilde{T}) \leq \check{d}(T, \tilde{T})$ with equality when T and \tilde{T} are dilationally pure.

Proof. The inequality $d(T, \tilde{T}) \leq \check{d}(T, \tilde{T})$ is clear from the definition of \check{d} by monotonicity of d . It is moreover obvious that we have the reverse inequality (and hence equality) when the channels T and \tilde{T} are dilationally pure.

As for the claim that \check{d} is a metric, symmetry follows from symmetry of d , and non-degeneracy from non-degeneracy of d along with the inequality $d \leq \check{d}$. It thus only remains to show the triangle inequality, which is also the hardest part. To this end, let $T_1, T_2, T_3 : \mathbb{X} \rightarrow \mathbb{Y}$ be channels, and assume to arrive at a contradiction that $\check{d}(T_1, T_2) + \check{d}(T_2, T_3) < \check{d}(T_1, T_3)$. Then, we can pick by definition of \check{d} pure dilations Σ_1 and Σ_2 of T_1 and T_2 , respectively, and pure dilations Σ'_2 and Σ_3 of T_2 and T_3 , respectively, such that

$$d(\Sigma_1, \Sigma_2) + d(\Sigma'_2, \Sigma_3) < \check{d}(T_1, T_3). \quad (3.31)$$

Now, if it were the case that $\Sigma_2 = \Sigma'_2$ then the triangle inequality for d would imply that $d(\Sigma_1, \Sigma_3)$ is a lower bound to the left hand side, in conflict with the definition of $\check{d}(T_1, T_3)$. The challenge is, however, that this need not be the case. In **QIT**, we can apply isometries to ‘align’ Σ_2 and Σ'_2 , and the general remedy is indeed a generalisation of this trick.

First, pick a universal dilation U_2 of T_2 . Then there exist channels V and V' such that Σ_2 equals $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{U_2} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{V} \end{array}$ and Σ'_2 equals $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{U_2} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{V'} \end{array}$. By universality, V and V' must moreover be dilationally pure since Σ_2 and Σ'_2 are, so by Lemma 3.4.4 below we find pure channels \tilde{V} and \tilde{V}' such that $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{V} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{V'} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}$, whence we actually have $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_2} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma'_2} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}$. In other words, though Σ_2 and Σ'_2 are not necessarily identical, we can apply pure channels in the environment and thereby align them. This now implies by the triangle inequality for d that

¹⁰And satisfies the technical assumption about existence of compatible pure dilations.

$$\begin{aligned}
d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_1} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_3} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}\right) &\leq d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_1} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_2} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array}\right) \\
&+ d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_2'} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_3} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}\right), \tag{3.32}
\end{aligned}$$

and because the channels \tilde{V} and \tilde{V}' are reversible, monotonicity of d implies that the two distances on the right hand side coincide with $d(\Sigma_1, \Sigma_3)$ and $d(\Sigma_2', \Sigma_3)$, respectively, so by Eq. (3.31) we actually have

$$d\left(\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_1} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array}, \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_3} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}\right) < \check{d}(T_1, T_3). \tag{3.33}$$

Finally, however, as \tilde{V} and \tilde{V}' are pure, the two dilations $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_1} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}} \end{array}$ and $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{\Sigma_3} \\ \text{---}\mathbb{Y}\text{---} \\ \boxed{\tilde{V}'} \end{array}$ are (by localisability) pure dilations of T_1 and T_3 , respectively, so this contradicts the definition of $\check{d}(T_1, T_3)$. This altogether demonstrates the triangle inequality. (Incidentally, the specialisation of this proof to the case of states yields a different proof for the triangle inequality of the purified distance than the one given in Ref. [Tom12].)

To show that \check{d} is serially monotone, it suffices to observe that by localisability the serial composition of pure dilations is a pure dilation and then use serial monotonicity of d . Parallel invariance of \check{d} is proved by using parallel invariance of d and the fact that identities are dilationally pure (Proposition 2.5.5).

Lastly, to see the \check{d} is dilational, we may by Proposition 3.3.6 equivalently demonstrate that \check{d} has the Generalised Uhlmann Property. The inequality $\check{d}(T, \tilde{T}) \leq \inf_{K, \tilde{K}} \check{d}(K, \tilde{K})$ holds simply because d is monotone; the non-trivial inequality is the converse, $\check{d}(T, \tilde{T}) \geq \inf_{K, \tilde{K}} \check{d}(K, \tilde{K})$, and it owes its validity to the hereditary property of completeness in the guise of Proposition 2.5.16. Indeed, we have

$$\check{d}(T, \tilde{T}) = \inf_{\Sigma, \tilde{\Sigma}} d(\Sigma, \tilde{\Sigma}) = \inf_{\Sigma, \tilde{\Sigma}} \check{d}(\Sigma, \tilde{\Sigma}) \tag{3.34}$$

by definition of \check{d} and the introductory observation, and because every pure dilation of a channel is complete (Proposition 2.5.16), the above quantity is trivially an upper bound to $\inf_{K, \tilde{K}} \check{d}(K, \tilde{K})$. □

Lemma 3.4.4. *(Pure Channels can be Purely Aligned.)*

Suppose that Θ is localisable, universal and purifiable. Then for any pure channels $\begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{V} \\ \text{---}\mathcal{Y}\text{---} \end{array}$ and $\begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{V'} \\ \text{---}\mathcal{Y}'\text{---} \end{array}$, there exist pure channels $\begin{array}{c} \text{---}\mathcal{Y}\text{---} \\ \boxed{\tilde{V}} \\ \text{---}\mathcal{Z}\text{---} \end{array}$ and $\begin{array}{c} \text{---}\mathcal{Y}'\text{---} \\ \boxed{\tilde{V}'} \\ \text{---}\mathcal{Z}\text{---} \end{array}$ such that

$$\begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{V} \\ \text{---}\mathcal{Y}\text{---} \\ \boxed{\tilde{V}} \\ \text{---}\mathcal{Z}\text{---} \end{array} = \begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{V'} \\ \text{---}\mathcal{Y}'\text{---} \\ \boxed{\tilde{V}'} \\ \text{---}\mathcal{Z}\text{---} \end{array} . \tag{3.35}$$

Proof. By Theorem 2.5.11, V and V' are reversible, hence we can find S and S' with

$$\begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{V} \\ \text{---}\mathcal{Y}\text{---} \\ \boxed{S} \\ \text{---}\mathcal{X}\text{---} \end{array} = \begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{\text{id}} \\ \text{---}\mathcal{X}\text{---} \end{array} = \begin{array}{c} \text{---}\mathcal{X}\text{---} \\ \boxed{V'} \\ \text{---}\mathcal{Y}'\text{---} \\ \boxed{S'} \\ \text{---}\mathcal{X}\text{---} \end{array} . \tag{3.36}$$

Now, S and S' have universal dilations, say U and U' , which by Proposition 2.5.10 are dilationally pure, so the channels

$$-\mathcal{X}-\boxed{V}-\boxed{U}-\mathcal{X}-, \quad -\mathcal{X}-\boxed{V'}-\boxed{U'}-\mathcal{X}- \quad (3.37)$$

are pure dilations of $-\mathcal{X}-\boxed{\text{id}}-\mathcal{X}-$. Since identities are dilationally pure by Proposition 2.5.5, these dilations must consequently take the form

$$-\mathcal{X}-\boxed{\text{id}}-\mathcal{X}-, \quad -\mathcal{X}-\boxed{\text{id}}-\mathcal{X}-, \quad (3.38)$$

$$\boxed{t}, \quad \boxed{t'}$$

respectively, for pure states t and t' . Now, those two channels can clearly be purely aligned, simply by tensoring with t' and t , respectively. Pre-composing with U and U' we altogether obtain pure channels \tilde{V} and \tilde{V}' which align V and V' , as desired. \square

Corollary 3.4.5. *The purified diamond-distance P_\diamond is a monotone, dilational metric on the theory QIT.*

It is quite easy to identify the characteristic property of purified distances:¹¹

Proposition 3.4.6. *(Purified means Uhlmann.)*

Suppose that Θ is a localisable, universal and purifiable theory, and suppose that D is a monotone metric on Θ . Then the following are equivalent:

1. $D = \check{d}$ for some monotone metric d .
2. D has the Uhlmann Property, that is, for any channels $T, \tilde{T} : \mathbb{X} \rightarrow \mathbb{Y}$, it holds that

$$D(-\mathbb{X}-\boxed{T}-\mathbb{Y}-, -\mathbb{X}-\boxed{\tilde{T}}-\mathbb{Y}-) = \inf_{\Sigma, \tilde{\Sigma}} D\left(-\mathbb{X}-\boxed{\Sigma}-\mathbb{Y}-, -\mathbb{X}-\boxed{\tilde{\Sigma}}-\mathbb{Y}-\right), \quad (3.39)$$

where the infimum is over all compatible pairs of pure one-sided dilations Σ of T and $\tilde{\Sigma}$ of \tilde{T} .

Proof. If $D = \check{d}$, then $D(T, \tilde{T}) = \inf_{\Sigma, \tilde{\Sigma}} d(\Sigma, \tilde{\Sigma}) = \inf_{\Sigma, \tilde{\Sigma}} \check{d}(\Sigma, \tilde{\Sigma}) = \inf_{\Sigma, \tilde{\Sigma}} D(\Sigma, \tilde{\Sigma})$ since $\check{d}(\Sigma, \tilde{\Sigma}) = d(\Sigma, \tilde{\Sigma})$ by the statement in Theorem 3.4.3. If conversely D has the Uhlmann property, then obviously $D = \check{D}$. \square

Now that we have established the purified distance \check{d} as a monotone, dilational metric, and understood its characteristic trait, let us observe what could be an important quality of it, namely the fact that it tightly captures the approximate theory of complementary channels.

Recall from Section 2.5 that the channels $-\mathbb{X}-\boxed{T}-\mathbb{Y}-$ and $-\mathbb{X}-\boxed{T^c}-\mathbb{E}-$ are said to be *complementary* if they admit a common pure dilation $-\mathbb{X}-\boxed{\Sigma}-\mathbb{Y}-$. We have the following:

¹¹We will not need this result, I simply produce it for the reader to gain further intuition.

Proposition 3.4.7. (*Approximate Complementarity.*)

Suppose that $D = \check{d}$ is a purified distance, or, equivalently, has the Uhlmann property. For any channels $-X-\boxed{T}-Y-$ and $-X-\boxed{\tilde{T}}-Y-$, it holds that

$$\inf_{T^c, \tilde{T}^c} \check{d}(-X-\boxed{T^c}-\mathbb{E}\sim, -X-\boxed{\tilde{T}^c}-\mathbb{E}\sim) \leq \check{d}(-X-\boxed{T}-Y-, -X-\boxed{\tilde{T}}-Y-), \quad (3.40)$$

where the infimum is over all pairs of channels T^c complementary to T and \tilde{T}^c complementary to \tilde{T} , with the same codomain.

Proof. Simply observe that by monotonicity of \check{d} under marginalisation we have $\check{d}(T^c, \tilde{T}^c) \leq \check{d}(\Sigma, \tilde{\Sigma})$ for pure dilations Σ and $\tilde{\Sigma}$ which witness the complementarity. \square

This result immediately implies an approximate version of Theorem 2.5.18, in the sense that a channel T is close to a reversible channel if and only if its complementary channel T^c is close to being completely forgetful.

We can, however, improve on this observation by discarding Proposition 3.4.7 and using instead mere dilationality, and with this I will conclude the section. The improvement will reside in the fact that there is a natural notion of a channel T being ‘approximately reversible’, and this notion is weaker than T being approximately equal to a channel which is genuinely reversible.¹²

The improved result can be seen as an abstract statement of the ‘information-disturbance trade off’ of Ref. [KSW08b]:

Theorem 3.4.8. (*Approximate Duality between Reversible and Completely Forgetful Channels.*)

Let Θ be a universal, localisable and purifiable theory, and let D be a monotone and dilational metric on Θ . Then, for any channel $-X-\boxed{T}-Y-$ and any complementary channel $-X-\boxed{T^c}-\mathcal{E}\sim$, we have

$$\inf_{T^-} D(-X-\boxed{T}-Y-\boxed{T^-}-X-, -X-\boxed{\text{id}}-X-) = \inf_s D(-X-\boxed{T^c}-\mathcal{E}\sim, -X-\boxed{\text{tr}} \boxed{s}-\mathcal{E}\sim), \quad (3.41)$$

where the first infimum is over all channels $-Y-\boxed{T^-}-X-$ and the second infimum over all states $\boxed{s}-\mathcal{E}\sim$.

Proof. Let $-Y-\boxed{T^-}-X-$ be a channel and define

$$\varepsilon := D(-X-\boxed{T}-Y-\boxed{T^-}-X-, -X-\boxed{\text{id}}-X-). \quad (3.42)$$

If as in the proof of Theorem 2.5.11 we let $-X-\boxed{\check{T}}-\mathbb{E}\sim$ be a universal (and hence by Proposition 2.5.10 pure) dilation of T , then by dilationality of D we find a state s' such that

$$-X-\boxed{\check{T}}-Y-\boxed{T^-}-X- \underset{\varepsilon}{\approx} -X-\boxed{\text{id}}-X- \underset{\varepsilon}{\approx} -X-\boxed{s'}-\mathcal{E}\sim \quad (3.43)$$

¹²To be fair, it constitutes an improvement only in one direction of the bi-implication.

(since $\text{id}_{\mathcal{X}}$ is dilationally pure). Here, ' \approx_{ε}^D ' signifies that the distance between the two channels as measured by D is at most ε . Now, trashing \mathcal{X} yields on the left hand side some channel $- \mathcal{X} - \boxed{T'} - \check{\varepsilon} -$ complementary to T , and on the right hand side the completely forgetful channel $- \mathcal{X} - \boxed{\text{tr}} \boxed{s'} - \check{\varepsilon} -$ which by monotonicity is ε -close to $- \mathcal{X} - \boxed{T'} - \check{\varepsilon} -$ w.r.t. D . By Theorem 2.5.17, the specific complementary channel $- \mathcal{X} - \boxed{T^c} - \check{\varepsilon} -$ must equal $- \mathcal{X} - \boxed{T'} - \check{\varepsilon} - \boxed{G} - \check{\varepsilon} -$ for some channel G , and by monotonicity it is then ε -close w.r.t. D to the completely forgetful channel $- \mathcal{X} - \boxed{\text{tr}} \boxed{s} - \check{\varepsilon} -$, with $\boxed{s} - \check{\varepsilon} - = \boxed{s'} - \check{\varepsilon} - \boxed{G} - \check{\varepsilon} -$. Consequently,

$$D(- \mathcal{X} - \boxed{T^c} - \check{\varepsilon} - , - \mathcal{X} - \boxed{\text{tr}} \boxed{s} - \check{\varepsilon} -) \leq \varepsilon = D(- \mathcal{X} - \boxed{T} - \mathcal{Y} - \boxed{T^-} - \mathcal{X} - , - \mathcal{X} - \boxed{\text{id}} - \mathcal{X} -), \quad (3.44)$$

and therefore the inequality ' \geq ' holds in Eq. (3.41).

For the converse inequality, let $\boxed{s} - \check{\varepsilon} -$ be a state and define

$$\varepsilon := D(- \mathcal{X} - \boxed{T^c} - \check{\varepsilon} - , - \mathcal{X} - \boxed{\text{tr}} \boxed{s} - \check{\varepsilon} -). \quad (3.45)$$

By arguments similar to the above (or better, by Proposition 3.4.7 and Theorem 2.5.18), we find a channel $(T^c)^c$ complementary to T^c and a reversible channel R such that $D((T^c)^c, R) \leq \varepsilon$. As such, there exists a channel R^- with $R^- \circ R = \text{id}_{\mathcal{X}}$, so by monotonicity of D ,

$$R^- \circ (T^c)^c \approx_{\varepsilon}^D R^- \circ R = \text{id}_{\mathcal{X}}. \quad (3.46)$$

By Theorem 2.5.17 again, $(T^c)^c$ must equal $G \circ T$ for some channel G , so the above implies $(R^- \circ G) \circ T \approx_{\varepsilon}^D \text{id}_{\mathcal{X}}$, so with $T^- := R^- \circ G$ we finally have

$$D(- \mathcal{X} - \boxed{T} - \mathcal{Y} - \boxed{T^-} - \mathcal{X} - , - \mathcal{X} - \boxed{\text{id}} - \mathcal{X} -) \leq \varepsilon = D(- \mathcal{X} - \boxed{T^c} - \check{\varepsilon} - , - \mathcal{X} - \boxed{\text{tr}} \boxed{s} - \check{\varepsilon} -), \quad (3.47)$$

implying the inequality ' \leq ' in Eq. (3.41). □

The significance of Theorem 3.4.8 is that, when measured using an appropriate notion of distance – namely a monotone, dilational metric – the duality of Theorem 2.5.18 tightens under approximations. By an interpretation similar to the one usually employed in the case of **QIT**, we might say that the degree to which information is preserved by a channel T (i.e. the degree to which it can be reversed) equals the degree to which it leaks no information to the environment.

3.5 The Purified Diamond-Distance P_{\diamond}

In this section, we cast our attention on the *purified diamond-distance* which we defined above as the metric $P_{\diamond} := \check{d}_{\diamond}$ in **QIT** (Definition 3.4.2).

First, we will observe (owing fully to the result of Refs. [KSW08a, KSW08b]) that P_{\diamond} can be bounded non-trivially in terms of d_{\diamond} (Theorem 3.5.2). Then, we will compare more systematically the metric P_{\diamond} to the *Bures distance* β used in Refs. [KSW08a, KSW08b], in particular establishing a quantitative relationship between the two (Theorem 3.5.8).

Recall that the purified diamond-distance between quantum channels Λ_1, Λ_2 was defined as

$$P_\diamond(\Lambda_1, \Lambda_2) = \inf_{\Sigma_1, \Sigma_2} d_\diamond(\Sigma_1, \Sigma_2), \quad (3.48)$$

where the infimum is over all compatible Stinespring dilations Σ_1 of Λ_1 and Σ_2 of Λ_2 .¹³

On the other hand, the so-called *Bures-distance* ([KSW08a, KSW08b], generalising [Bur69]) between Λ_1 and Λ_2 , is given by $\beta(\Lambda_1, \Lambda_2) = \inf_{S_1, S_2} \|S_1 - S_2\|_\infty$, where the infimum is again over compatible Stinespring dilations of Λ_1, Λ_2 , but this time in terms of isometries S_j that represent the isometric dilations Σ_j (i.e. $\Sigma_j(A) = S_j A S_j^*$). The isometry representing a given isometric channel is unique up to a phase, and therefore we can equivalently define β as

$$\beta(\Lambda_1, \Lambda_2) = \inf_{\Sigma_1, \Sigma_2} d_\infty(\Sigma_1, \Sigma_2), \quad (3.49)$$

with d_∞ given by

$$d_\infty(\Sigma_1, \Sigma_2) := \inf_{\lambda \in \mathbb{T}} \|S_1 - \lambda S_2\|_\infty \quad (3.50)$$

for isometric channels $\Sigma_j(\cdot) = S_j(\cdot)S_j^*$, with \mathbb{T} denoting the unit circle in \mathbb{C} . The quantity $d_\infty(\Sigma_1, \Sigma_2)$ is independent of the choice of representatives S_1 and S_2 , and the definition (3.49) is more easily compared to (3.48).

Remark 3.5.1. (Qualitative Comparison of P_\diamond and β .)

Already before we quantitatively compare P_\diamond and β , a more qualitative comparison is possible and appropriate.

It can be checked (by more or less the same arguments as used in the proof of Theorem 3.4.3, but specialised to isometric channels in **QIT**), that β defines a monotone and dilational metric on **QIT**, even one which has the Uhlmann Property. This set of properties (which P_\diamond also has) therefore does not give a useful evaluation of the two against each other. For this reason, it also probably does not matter much in **QIT** which metric is used. However, the purified diamond-distance P_\diamond is defined in terms of dilations and the operational metric δ_\diamond , whereas the Bures distance β lacks a similar operational definition. Indeed, its definition is derived from operator algebra, and therefore contingent on a particular formalism of quantum information theory. It should also be noted that P_\diamond restricts in the case of states to the well-known purified distance ([TCR10, Tom12]), whereas the Bures distance does not.

It is finally worth remarking that, somewhat curiously, while P_\diamond always agrees with d_\diamond on isometric channels, β and d_∞ need not agree on isometric channels (cf. Remark 3.5.10). \blacktimes

¹³I will use a notation with number indices (Σ_1, Σ_2) rather than with tildes ($\Sigma, \tilde{\Sigma}$) in this section for better legibility.

3.5.A P_\diamond versus d_\diamond

The weightiest result about the purified distance P_\diamond is a rip-off from Refs. [KSW08a, KSW08b], whose formulation is in terms of the Bures distance β .

Whereas dilationality can act as a supply of ‘magic’ in manipulations (as exemplified by Theorem 3.4.8), that result will allow us to connect the magic of P_\diamond to the more mundane world of d_\diamond :

Theorem 3.5.2. (*Equivalence of d_\diamond and P_\diamond .*)

For any quantum channels Λ_1, Λ_2 , we have the inequalities

$$d_\diamond(\Lambda_1, \Lambda_2) \leq P_\diamond(\Lambda_1, \Lambda_2) \leq \sqrt{2 d_\diamond(\Lambda_1, \Lambda_2)}. \quad (3.51)$$

Remark 3.5.3. When Λ_1 and Λ_2 are states, so that P_\diamond is the purified distances for states, this result specialises to Lem. 6 in Ref. [TCR10] (or Prop. 3.3 in Ref. [Tom12]).

✂

Proof. The inequality $d_\diamond \leq P_\diamond$ follows from Theorem 3.4.3 (and its proof was easy). The hard part is the inequality $P_\diamond \leq \sqrt{2 d_\diamond}$.

The main result (Thm. 1) of Ref. [KSW08a] states in our language that

$$\frac{\|\Lambda_1 - \Lambda_2\|_\diamond}{\sqrt{\|\Lambda_1\|_\diamond} + \sqrt{\|\Lambda_2\|_\diamond}} \leq \beta(\Lambda_1, \Lambda_2) \leq \sqrt{\|\Lambda_1 - \Lambda_2\|_\diamond}, \quad (3.52)$$

where $\|\cdot\|_\diamond$ is the diamond norm and β the Bures distance. Now, any quantum channel has diamond-norm 1, and the diamond-norm is related to our diamond-distance by a factor of 2, so the above is the statement that

$$d_\diamond(\Lambda_1, \Lambda_2) \leq \beta(\Lambda_1, \Lambda_2) \leq \sqrt{2 d_\diamond(\Lambda_1, \Lambda_2)}. \quad (3.53)$$

To finish, it therefore suffices to prove $P_\diamond(\Lambda_1, \Lambda_2) \leq \beta(\Lambda_1, \Lambda_2)$, and by Eqs. (3.48) and (3.49) this will follow if $d_\diamond(\Sigma_1, \Sigma_2) \leq d_\infty(\Sigma_1, \Sigma_2)$ for any isometric channels Σ_1 and Σ_2 . In particular, it suffices to show that for any isometries $S_1, S_2 : \mathcal{H} \rightarrow \mathcal{K}$, we have $d_\diamond(\Sigma_1, \Sigma_2) \leq \|S_1 - S_2\|_\infty$, where Σ_j denotes the channel $A \mapsto S_j A S_j^*$.

But this follows from elementary calculations: For any state $\varrho \in \mathcal{D}(\mathcal{H} \otimes \mathcal{H})$, and any linear operators $A, B : \mathcal{H} \otimes \mathcal{H} \rightarrow \mathcal{L}$, we have

$$\begin{aligned} \|A\varrho A^* - B\varrho B^*\|_1 &\leq \|(A - B)\varrho A^*\|_1 + \|B\varrho(A^* - B^*)\|_1 \\ &\leq \|A - B\|_\infty \|\varrho\|_1 \|A^*\|_\infty + \|B\|_\infty \|\varrho\|_1 \|A^* - B^*\|_\infty \\ &= \|A - B\|_\infty (\|A\|_\infty + \|B\|_\infty), \end{aligned} \quad (3.54)$$

so in particular, for isometries S_1, S_2 ,

$$\begin{aligned} 2 d_1((\Sigma_1 \otimes \text{id}_\mathcal{H})(\varrho), (\Sigma_2 \otimes \text{id}_\mathcal{H})(\varrho)) &= \|(S_1 \otimes \mathbf{1}_\mathcal{H})\varrho(S_1^* \otimes \mathbf{1}_\mathcal{H}) - (S_2 \otimes \mathbf{1}_\mathcal{H})\varrho(S_2^* \otimes \mathbf{1}_\mathcal{H})\|_1 \\ &\leq \|S_1 \otimes \mathbf{1}_\mathcal{H} - S_2 \otimes \mathbf{1}_\mathcal{H}\|_\infty (\|S_1 \otimes \mathbf{1}_\mathcal{H}\|_\infty + \|S_2 \otimes \mathbf{1}_\mathcal{H}\|_\infty) \\ &= \|S_1 - S_2\|_\infty (1 + 1), \end{aligned} \quad (3.55)$$

from which it follows that

$$d_\diamond(\Sigma_1, \Sigma_2) = \sup_{\varrho \in \mathcal{S}(\mathcal{H} \otimes \mathcal{H})} d_1((\Sigma_1 \otimes \text{id}_{\mathcal{H}})(\varrho), (\Sigma_2 \otimes \text{id}_{\mathcal{H}})(\varrho)) \leq \|S_1 - S_2\|_\infty, \quad (3.56)$$

as desired. \square

Theorem 3.5.2 implies that we may (if we so desire) reformulate the result of Theorem 3.4.8, which applies to the dilational metric $D = P_\diamond$, as a result about diamond-distances, though in doing so we of course lose a square root (the resulting bound is one of the main results of Ref. [KSW08b]).

Theorem 3.5.2 also implies more abstract features, such as the non-trivial fact that the topologies induced by d_\diamond and P_\diamond coincide. (In particular, since sets of channels are compact w.r.t. this topology, the infimum in the dilationality condition for P_\diamond is always attained.)

In translating the proof of Refs. [KSW08a, KSW08b] from β to P_\diamond , we demonstrated the inequality $P_\diamond \leq \beta$. In the next and final subsection, we will examine the relationship between these two metrics more carefully.

3.5.B P_\diamond versus β

The complications of calculating P_\diamond and β is two-fold; both of determining an infimum over Stinespring dilations, and in turn of determining the d_\diamond -distance, respectively d_∞ -distance, itself between Stinespring dilations. For the latter distances, we have the following formulas:

Proposition 3.5.4. (*Calculating d_\diamond - and d_∞ -distances between Isometric Channels.*)

Let $\Sigma_1, \Sigma_2 : \mathcal{H} \rightarrow \mathcal{K}$ be isometric quantum channels. Define the fidelity, respectively fake fidelity, between Σ_1 and Σ_2 as

$$F(\Sigma_1, \Sigma_2) = \inf_{\varrho \in \mathcal{D}(\mathcal{H})} |\text{tr}(S_1^* S_2 \varrho)|, \text{ respectively} \quad (3.57)$$

$$FF(\Sigma_1, \Sigma_2) = \sup_{\lambda \in \mathbb{T}} \inf_{\varrho \in \mathcal{D}(\mathcal{H})} \text{Re}[\lambda \text{tr}(S_1^* S_2 \varrho)], \quad (3.58)$$

where S_1, S_2 represent Σ_1, Σ_2 in the sense that $\Sigma_j = S_j(\cdot)S_j^*$. It then holds that

$$d_\diamond(\Sigma_1, \Sigma_2) = \sqrt{1 - F(\Sigma_1, \Sigma_2)^2}, \text{ and} \quad (3.59)$$

$$d_\infty(\Sigma_1, \Sigma_2) = \sqrt{2 - 2FF(\Sigma_1, \Sigma_2)}. \quad (3.60)$$

Remark 3.5.5. (Relation to Fidelity between States.)

When the domain \mathcal{H} is trivial, Σ_1 and Σ_2 are pure states on \mathcal{K} , say ψ_1 and ψ_2 , so the infima in Eqs. (3.57) and (3.58) are over a one-element set, and both fidelities reduce to the ordinary fidelity for states, $|\langle \psi_1 | \psi_2 \rangle|$. This means in particular that we have an equational relationship for states ϱ_1, ϱ_2 between β and P_\diamond , given by

$$\sqrt{1 - P_\diamond(\varrho_1, \varrho_2)^2} = 1 - \frac{\beta(\varrho_1, \varrho_2)^2}{2}, \quad \text{or} \quad P_\diamond(\varrho_1, \varrho_2) = \beta(\varrho_1, \varrho_2) \sqrt{1 - \frac{\beta(\varrho_1, \varrho_2)^2}{4}}. \quad (3.61)$$

It is not clear that this relationship should hold for general channels, since the two fidelities might differ, but we shall demonstrate that this is in fact the case (Theorem 3.5.8).

\blacklozenge

Proof. As for d_\diamond , we have

$$d_\diamond(\Sigma_1, \Sigma_2) = \sup_{\psi \in \mathcal{D}(\mathcal{H} \otimes \mathcal{H}) \text{ pure}} d_1((\Sigma_1 \otimes \text{id}_{\mathcal{H}})(\psi), (\Sigma_2 \otimes \text{id}_{\mathcal{H}})(\psi)), \quad (3.62)$$

since by convexity the supremum over all states is attained already on the set of pure states. Using the well-known formula (see e.g. p. 415 in [NC02]) $d_1(\phi_1, \phi_2) = \sqrt{1 - |\langle \phi_1 | \phi_2 \rangle|^2}$ for pure states ϕ_1 and ϕ_2 , we find

$$\begin{aligned} d_1((\Sigma_1 \otimes \text{id}_{\mathcal{H}})(\psi), (\Sigma_2 \otimes \text{id}_{\mathcal{H}})(\psi)) &= \sqrt{1 - |\langle \psi | (S_1^* S_2 \otimes \mathbf{1}_{\mathcal{H}}) | \psi \rangle|^2} \\ &= \sqrt{1 - |\text{tr}[(S_1^* S_2 \otimes \mathbf{1}_{\mathcal{H}})\psi]|^2} \\ &= \sqrt{1 - |\text{tr}[S_1^* S_2 \pi_1(\psi)]|^2}, \end{aligned} \quad (3.63)$$

where $\pi_1(\psi) := [\text{id}_{\mathcal{H}} \otimes \text{tr}](\psi)$ signifies the first marginal of ψ . As ψ ranges over all pure states on $\mathcal{H} \otimes \mathcal{H}$, the quantity $\pi_1(\psi)$ ranges precisely over all states ϱ on \mathcal{H} . Therefore,

$$d_\diamond(\Sigma_1, \Sigma_2) = \sup_{\varrho \in \mathcal{D}(\mathcal{H})} \sqrt{1 - |\text{tr}(S_1^* S_2 \varrho)|^2}, \quad (3.64)$$

from which the formula (3.59) follows.

As for d_∞ , we find (using the formula $\|\phi_1 - \phi_2\|^2 = \|\phi_1\|^2 + \|\phi_2\|^2 - 2 \text{Re}(\langle \phi_1 | \phi_2 \rangle)$), that

$$\begin{aligned} \|S_1 - \lambda S_2\|_\infty^2 &= \sup_{|\psi\rangle \in \mathcal{H}, \|\psi\|=1} \|S_1 |\psi\rangle - \lambda S_2 |\psi\rangle\|^2 = \sup_{|\psi\rangle \in \mathcal{H}, \|\psi\|=1} (1 + 1 - 2 \text{Re}(\lambda \langle \psi | S_1^* S_2 | \psi \rangle)) \\ &= \sup_{\psi \in \mathcal{D}(\mathcal{H}) \text{ pure}} (2 - 2 \text{Re}(\lambda \text{tr}(S_1^* S_2 \psi))) = \sup_{\varrho \in \mathcal{D}(\mathcal{H})} (2 - 2 \text{Re}[\lambda \text{tr}(S_1^* S_2 \varrho)]), \end{aligned} \quad (3.65)$$

where the last equality is due to convex-linearity of the map $\varrho \mapsto \text{Re}[\lambda \text{tr}(S_1^* S_2 \varrho)]$. Consequently,

$$d_\infty(\Sigma_1, \Sigma_2) = \inf_{\lambda \in \mathbb{T}} \|S_1 - \lambda S_2\|_\infty = \inf_{\lambda \in \mathbb{T}} \sup_{\varrho \in \mathcal{D}(\mathcal{H})} \sqrt{2 - 2 \text{Re}[\lambda \text{tr}(S_1^* S_2 \varrho)]}, \quad (3.66)$$

from which the formula (3.60) follows. This finishes the proof. \square

From Proposition 3.5.4 follow the formulas

$$P_\diamond(\Lambda_1, \Lambda_2) = \sqrt{1 - \sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2)^2}, \quad (3.67)$$

$$\beta(\Lambda_1, \Lambda_2) = \sqrt{2 - 2 \sup_{\Sigma_1, \Sigma_2} FF(\Sigma_1, \Sigma_2)}, \quad (3.68)$$

where the suprema range over pairs of compatible Stinespring dilations Σ_1 of Λ_1 and Σ_2 of Λ_2 . It is not evident from those formulas how P_\diamond and β might be related, since we lack a relation between the fidelity and the fake fidelity, except for the obvious inequality $FF(\Sigma_1, \Sigma_2) \leq F(\Sigma_1, \Sigma_2)$ (which, incidentally, yields a different argument for the inequality $d_\diamond(\Sigma_1, \Sigma_2) \leq d_\infty(\Sigma_1, \Sigma_2)$ used in the proof of Theorem 3.5.2).

We will establish in a moment that $\sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2) = \sup_{\Sigma_1, \Sigma_2} FF(\Sigma_1, \Sigma_2)$, but it is worth observing that the identity $F(\Sigma_1, \Sigma_2) = FF(\Sigma_1, \Sigma_2)$ does not necessarily hold; in fact, the fake fidelity $FF(\Sigma_1, \Sigma_2)$ can even be negative (hence its name). To gain a feeling for this, and for the two fidelities in general, we will consider below an example in which Σ_1 and Σ_2 are unitary conjugations.

A key observation in analysing the fidelities is that for any linear operator $A : \mathcal{H} \rightarrow \mathcal{H}$, the set

$$C(A) := \{\operatorname{tr}(A\rho) \mid \rho \in \mathcal{D}(\mathcal{H})\}, \quad (3.69)$$

known as the *numerical range of A* , is a convex compact subset of \mathbb{C} (contained in the disc of radius $\|A\|_\infty$), since it is a continuous linear image of the convex compact set $\mathcal{D}(\mathcal{H})$. In particular, this holds for the linear operator $A = S_1^* S_2$.

The following (rather long) example exhibits a calculation of the fidelity and fake fidelity in the case where Σ_1 and Σ_2 are unitary. It can be skipped without losing coherence.

Example 3.5.6. (Fidelity and Fake Fidelity between Unitary Conjugations.)

Suppose that Σ_1 and Σ_2 are unitary conjugations. Let S_1 and S_2 be some choice of unitaries representing Σ_1 and Σ_2 , respectively. Then, the operator $S_1^* S_2$ is unitary, and thus admits a basis of eigenvectors. This means that for any unit vector $|\psi\rangle \in \mathcal{H}$, the quantity $\operatorname{tr}(S_1^* S_2 |\psi\rangle\langle\psi|) = \langle\psi| S_1^* S_2 |\psi\rangle$ is a convex combination of eigenvalues of $S_1^* S_2$, and any particular eigenvalue is obtained by a suitable choice of $|\psi\rangle$. It follows that $C(S_1^* S_2)$ as given by Eq. (3.69) coincides with $\operatorname{Conv}(\operatorname{spec}(S_1^* S_2))$, the convex hull of the spectrum of $S_1^* S_2$.

Calculating the Fidelity. Now, the fidelity $F(\Sigma_1, \Sigma_2)$ is geometrically the distance from the origin $0 \in \mathbb{C}$ to the set $C(S_1^* S_2)$. We can determine this distance by simple considerations. All of the eigenvalues of $S_1^* S_2$ will lie on the unit circle in \mathbb{C} ; let $\gamma(\Sigma_1, \Sigma_2) \in [0, 2\pi)$ denote the length of the smallest closed arc containing all of the eigenvalues. (Equivalently, $\gamma(\Sigma_1, \Sigma_2) = 2\pi - \bar{\gamma}(\Sigma_1, \Sigma_2)$, where $\bar{\gamma}(\Sigma_1, \Sigma_2)$ is the length of the largest of the finitely many open arcs into which the unit circle is divided by the eigenvalues.) Observe that this quantity indeed depends only on the channels Σ_1 and Σ_2 , and not on the chosen representatives S_1 and S_2 , since another choice will merely have the effect of rotating the set of eigenvalues around the origin. It turns out that $F(\Sigma_1, \Sigma_2)$ is a function of $\gamma(\Sigma_1, \Sigma_2)$ alone:

If $\gamma(\Sigma_1, \Sigma_2) \geq \pi$, no straight line can strictly separate $\operatorname{spec}(S_1^* S_2)$ from the point 0, and by a standard result in convex analysis, we must therefore have $0 \in C(S_1^* S_2)$; consequently, $F(\Sigma_1, \Sigma_2) = 0$.

If $\gamma(\Sigma_1, \Sigma_2) < \pi$, all of the eigenvalues lie on one side of some straight line through 0. Now, if $C(S_1^* S_2)$ consists of a single point (i.e. if $\gamma(\Sigma_1, \Sigma_2) = 0$), we clearly have $F(\Sigma_1, \Sigma_2) = 1$. If $C(S_1^* S_2)$ does not consist of a single point, the infimal distance from 0 to a point in $C(S_1^* S_2)$ must be attained on one of those faces of the polytope $C(S_1^* S_2)$ which is a straight line segment between two distinct eigenvalues. In fact, by simple trigonometry, it

must be attained precisely in the midpoint of such a line segment. It is intuitively obvious that the two eigenvalues, for which this midpoint is closest to 0, are the ones which have the entire arc of length $\gamma(\Sigma_1, \Sigma_2)$ between them, say λ_a and λ_b . We can also see this formally: By definition of $\gamma(\Sigma_1, \Sigma_2)$, any two distinct eigenvalues are separated by an arc of length $\alpha \leq \gamma(\Sigma_1, \Sigma_2) (< \pi)$. For a pair of points on the circle separated by an arc of length $\alpha \in [0, \pi]$, the distance from 0 to the midpoint of the line segment between them is given by $\cos(\alpha/2)$. Since this expression decreases as α increases, it follows that indeed the desired infimal distance is $\cos(\gamma(\Sigma_1, \Sigma_2)/2)$.

Summarising these findings, the fidelity is given by

$$F(\Sigma_1, \Sigma_2) = \begin{cases} \cos\left(\frac{\gamma(\Sigma_1, \Sigma_2)}{2}\right) & \text{if } \gamma(\Sigma_1, \Sigma_2) \in [0, \pi) \\ 0 & \text{if } \gamma(\Sigma_1, \Sigma_2) \in [\pi, 2\pi) \end{cases} = \cos(\gamma(\Sigma_1, \Sigma_2)/2)^+, \quad (3.70)$$

where, as usual, $u^+ := \max\{0, u\}$ for $u \in \mathbb{R}$.

Calculating the Fake Fidelity. As for the fake fidelity, the analysis is also rather easily carried out using elementary arguments. $FF(\Sigma_1, \Sigma_2)$ is by definition the minimal real part of a point in the rotated set $\lambda C(S_1^* S_2)$, for a choice of λ which gives the largest possible such minimal real part. Again, this quantity turns out to depend only on $\gamma(\Sigma_1, \Sigma_2)$:

If $\gamma(\Sigma_1, \Sigma_2) < \pi$, then, as observed before, all of $\text{spec}(S_1^* S_2)$ will lie on one side of some straight line through 0. It is intuitively clear, that the largest possible smallest real part of a rotation of $C(S_1^* S_2)$ is in this case coincident with the distance from $C(S_1^* S_2)$ to 0, i.e. equals $F(\Sigma_1, \Sigma_2) (= \cos(\gamma(\Sigma_1, \Sigma_2)/2))$, but a formal argument is also rather simple: $FF(\Sigma_1, \Sigma_2) \leq F(\Sigma_1, \Sigma_2)$ always, and $FF(\Sigma_1, \Sigma_2) \geq F(\Sigma_1, \Sigma_2)$ follows by choosing $\lambda \in \mathbb{T}$ such that the straight line through 0 and some 0-nearest point $z_0 \in C(S_1^* S_2)$ becomes horizontal; z_0 will then also be a point with minimal real part, since if $z' \in C(S_1^* S_2)$ had strictly smaller real part, then some point on the line segment between z_0 and z' (which by convexity belongs to $C(S_1^* S_2)$) would be strictly closer to 0 than z_0 is, contradicting its choice.

If, on the other hand, $\gamma(\Sigma_1, \Sigma_2) \geq \pi$, a new situation emerges. In this case, the rotation with the largest possible minimal real part is the one which positions the arc of length $\gamma(\Sigma_1, \Sigma_2)$ such that its mid-point is the point $1 \in \mathbb{C}$, or, equivalently, positions the eigenvalue-free arc of length $\bar{\gamma}(\Sigma_1, \Sigma_2)$ such that its mid-point is $-1 \in \mathbb{C}$ (thus with its endpoints having identical real parts). To see this, simply note that if some choice of λ yielded a strictly larger minimal real part, say $r \in (-1, 0]$, then the rotated set of eigenvalues, $\lambda \text{spec}(S_1^* S_2)$, would be contained in the half-plane to the right of the straight line $\text{Re}(z) = r$, and this would imply that the eigenvalue-free arc to the left of this line had length strictly greater than $\bar{\gamma}(\Sigma_1, \Sigma_2)$, contradicting its definition. Having established this, we see by easy trigonometric considerations that the relevant minimal real part is $\cos\left(\frac{\gamma(\Sigma_1, \Sigma_2)}{2}\right)$, so that altogether we conclude that, globally for $\gamma(\Sigma_1, \Sigma_2) \in [0, 2\pi)$, we have

$$FF(\Sigma_1, \Sigma_2) = \cos(\gamma(\Sigma_1, \Sigma_2)/2). \quad (3.71)$$

Comparison. We observe the relationship $F(\Sigma_1, \Sigma_2) = \max\{0, FF(\Sigma_1, \Sigma_2)\}$, which implies that for $\gamma(\Sigma_1, \Sigma_2) \in [0, \pi)$ the two fidelities coincide. However, by choosing Σ_1 and Σ_2 appropriately, we can design the value of $\gamma(\Sigma_1, \Sigma_2)$ at will – e.g. letting $S_1 = \mathbb{1}_{\mathcal{H}}$ and letting S_2 have eigenvalues which are spread out suitably along the unit circle. (Note, though, that since $S_1^* S_2$ has at most $\dim \mathcal{H}$ distinct eigenvalues, the arc length $\gamma(\Sigma_1, \Sigma_2)$ is at most $2\pi - \frac{2\pi}{\dim \mathcal{H}}$). In particular, by choosing $\gamma(\Sigma_1, \Sigma_2) > \pi$ the fidelity $F(\Sigma_1, \Sigma_2)$ will be

constantly 0, whereas the fake fidelity $FF(\Sigma_1, \Sigma_2)$ approaches -1 as $\gamma(\Sigma_1, \Sigma_2)$ approaches 2π . \blacklozenge

The relationship $F(\Sigma_1, \Sigma_2) = \max\{0, FF(\Sigma_1, \Sigma_2)\}$ observed in Example 3.5.6 for unitary channels turns out to be no coincidence. To see this, we use a minimax-theorem:

Lemma 3.5.7. (*Fake it Till You Make it.*)
For any isometric quantum channels $\Sigma_1, \Sigma_2 : \mathcal{H} \rightarrow \mathcal{K}$, it holds that

$$F(\Sigma_1, \Sigma_2) = \max\{0, FF(\Sigma_1, \Sigma_2)\}. \quad (3.72)$$

In particular, $FF(\Sigma_1, \Sigma_2) = F(\Sigma_1, \Sigma_2)$ as soon as $FF(\Sigma_1, \Sigma_2) \geq 0$.

Proof. Let S_1 and S_2 be isometries representing Σ_1 and Σ_2 , respectively. As already observed, the set

$$C(S_1^* S_2) := \{\text{tr}(S_1^* S_2 \varrho) \mid \varrho \in \mathcal{D}(\mathcal{H})\} \quad (3.73)$$

is a convex compact subset of \mathbb{C} . Now, we can write the two fidelities as

$$F(\Sigma_1, \Sigma_2) = \inf_{z \in C(S_1^* S_2)} |z|, \quad (3.74)$$

$$FF(\Sigma_1, \Sigma_2) = \sup_{\lambda \in \mathbb{T}} \inf_{z \in C(S_1^* S_2)} \text{Re}(\lambda z). \quad (3.75)$$

Clearly, $|z| = \sup_{\alpha \in \mathbb{D}} \text{Re}(\alpha z)$ for any $z \in \mathbb{C}$, where $\mathbb{D} := \{\alpha \in \mathbb{C} \mid |\alpha| \leq 1\}$ denotes the closed unit disc in \mathbb{C} , so we really have

$$F(\Sigma_1, \Sigma_2) = \inf_{z \in C(S_1^* S_2)} \sup_{\alpha \in \mathbb{D}} \text{Re}(\alpha z). \quad (3.76)$$

Now, the map $(z, \alpha) \mapsto \text{Re}(\alpha z)$ is convex-linear in both of its arguments, and both of the sets $C(S_1^* S_2)$ and \mathbb{D} are convex and compact. Therefore, the order of optimisation can be interchanged by (e.g.) von Neumann's minimax-theorem, so in fact

$$F(\Sigma_1, \Sigma_2) = \sup_{\alpha \in \mathbb{D}} \inf_{z \in C(S_1^* S_2)} \text{Re}(\alpha z) = \sup_{r \in [0,1]} \sup_{\lambda \in \mathbb{T}} \inf_{z \in C(S_1^* S_2)} \text{Re}(r \lambda z) = \sup_{r \in [0,1]} r \cdot FF(\Sigma_1, \Sigma_2). \quad (3.77)$$

The formula (3.72) now follows, since if $FF(\Sigma_1, \Sigma_2) \leq 0$ the last supremum above equals 0, and if $FF(\Sigma_1, \Sigma_2) > 0$ it equals $FF(\Sigma_1, \Sigma_2)$. \square

As demonstrated by Example 3.5.6, the condition $FF(\Sigma_1, \Sigma_2) \geq 0$ is not a void one. It turns out, however, that we nevertheless obtain a functional relationship between P_\diamond and β , and with this result we conclude the section:

Theorem 3.5.8. (P_\diamond versus β .)

For any quantum channels Λ_1, Λ_2 , it holds that

$$\sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2) = \sup_{\Sigma_1, \Sigma_2} FF(\Sigma_1, \Sigma_2), \quad (3.78)$$

where the suprema range over pairs of compatible Stinespring dilations Σ_1 of Λ_1 and Σ_2 of Λ_2 . In particular, we have the relationship

$$\sqrt{1 - P_\diamond(\Lambda_1, \Lambda_2)^2} = 1 - \frac{\beta(\Lambda_1, \Lambda_2)^2}{2}, \quad (3.79)$$

or

$$P_\diamond(\Lambda_1, \Lambda_2) = \beta(\Lambda_1, \Lambda_2) \sqrt{1 - \frac{\beta(\Lambda_1, \Lambda_2)^2}{4}}. \quad (3.80)$$

Remark 3.5.9. (Ranges and Asymptotics.)

Theorem 3.5.8 shows in particular that as P_\diamond goes from 0 to 1, β goes from 0 to $\sqrt{2}$, and vice versa. Moreover, it shows that the ratio β/P_\diamond approaches 1 when they are near 0 (i.e. β and P_\diamond are asymptotically equal near 0), whereas it approaches $\sqrt{2}$ when they are near their respective maximal values. The approximation $P_\diamond \approx \beta$ near 0 is even correct to quadratic order ($P_\diamond = \beta - \beta^3/8 + O(\beta^5)$ in the limit $\beta \rightarrow 0$). \clubsuit

Remark 3.5.10. (β versus d_∞ .)

We have previously observed that P_\diamond and d_\diamond agree on isometric channels. Note that this is not the case for β and d_∞ : By Example 3.5.6, $FF(\Sigma_1, \Sigma_2)$ can range from -1 to 1 , so (by Proposition 3.5.4) $d_\infty(\Sigma_1, \Sigma_2)$ can range from 0 to 2. On the other hand, by Theorem 3.5.8 $\beta(\Sigma_1, \Sigma_2)$ only ranges from 0 to $\sqrt{2}$ as $P_\diamond(\Sigma_1, \Sigma_2)$ ranges from 0 to 1, ultimately due to the collapse of the infimum over fake fidelities to the infimum over (non-fake) fidelities. \clubsuit

Proof. The relations (3.79) and (3.80) follow from Eq. (3.78) by Proposition 3.5.4, so we need only show Eq. (3.78). Since $F(\Sigma_1, \Sigma_2) \geq FF(\Sigma_1, \Sigma_2)$ always, the inequality ‘ \geq ’ is clear so it suffices to argue that $\sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2) \leq \sup_{\Sigma_1, \Sigma_2} FF(\Sigma_1, \Sigma_2)$.

Suppose first that $\sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2) = 0$. It is sufficient to find Σ'_1, Σ'_2 such that $FF(\Sigma'_1, \Sigma'_2) = 0$. However, starting from any (compatible) Stinespring dilations Σ_1, Σ_2 , the channels $\Sigma'_1 = \Sigma_1 \otimes \psi_1$ and $\Sigma'_2 = \Sigma_2 \otimes \psi_2$ are (compatible) Stinespring dilations for any choice of pure states ψ_1 and ψ_2 on the same system. In particular, if ψ_1 and ψ_2 are chosen perfectly distinguishable ($\langle \psi_1 | \psi_2 \rangle = 0$), then

$$FF(\Sigma'_1, \Sigma'_2) = FF(\Sigma_1 \otimes \psi_1, \Sigma_2 \otimes \psi_2) = \sup_{\lambda \in \mathbb{T}} \inf_{\varrho \in \mathcal{D}(\mathcal{H})} \operatorname{Re}[\lambda \operatorname{tr}(\langle \psi_1 | \psi_2 \rangle S_1^* S_2 \varrho)] = 0, \quad (3.81)$$

as desired.¹⁴

Next, suppose that $\sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2) > 0$. We show that for any Σ_1, Σ_2 with $F(\Sigma_1, \Sigma_2) > 0$, we must have $FF(\Sigma_1, \Sigma_2) \geq 0$ and thus by Lemma 3.5.7 $FF(\Sigma_1, \Sigma_2) = F(\Sigma_1, \Sigma_2)$; this will imply that $\sup_{\Sigma_1, \Sigma_2} FF(\Sigma_1, \Sigma_2) \geq \sup_{\Sigma_1, \Sigma_2} F(\Sigma_1, \Sigma_2)$. To this end, observe that the

¹⁴In this argument lies incidentally the reason why it is possible to have $\beta(\Sigma_1, \Sigma_2) < d_\infty(\Sigma_1, \Sigma_2)$ for isometric channels Σ_1 and Σ_2 .

assumption $F(\Sigma_1, \Sigma_2) > 0$ implies by Eq. (3.74) that $0 \notin C(S_1^* S_2)$. As the set $C(S_1^* S_2) \subseteq \mathbb{C}$ is convex, it must therefore be entirely contained in the half-space defined by some straight line in \mathbb{C} through 0. This means, however, that for some $\lambda \in \mathbb{T}$, the rotated set $\lambda C(S_1^* S_2)$ is contained in the specific half-space $\{z \in \mathbb{C} \mid \operatorname{Re}(z) \geq 0\}$, and consequently we must by Eq. (3.75) have $FF(\Sigma_1, \Sigma_2) \geq 0$, as claimed. \square

3.6 Summary and Outlook

In this chapter, we have seen a general definition of topological and metric theories. We have discussed the properties of *monotonicity* (Definition 3.2.2) and *dilationality* (Definition 3.3.1) of a metric, and reformulated the latter in terms of a ‘Generalised Uhlmann Property’ (Proposition 3.3.6). This reformulation reveals in particular why monotone metrics in **CIT** (and cartesian theories) are bound to be automatically dilational, a phenomenon which might explain why the property of dilationality was only properly installed in metrics on **QIT** fairly recently ([TCR10]).

In Section 3.4, we saw how the concept of purifiability can be used to overcome the challenges that a priori surround the construction of a (monotone) dilational metric from a monotone one, and this led to the idea of *purified distances*. In particular, the *purified diamond-distance*, P_\diamond , was introduced, a metric which has to the best of my knowledge not been considered before, but whose generalisation from the purified distance for states ([TCR10, Tom12]) is fairly straightforward. In Section 3.5, we compared the metric P_\diamond to the Bures distance β of Refs. [KSW08a, KSW08b], with the main conclusion that the two metrics share their most characteristic properties but differ in their relation to the formalism of quantum information theory. We also found an explicit quantitative relation between P_\diamond and β (Theorem 3.5.8), which entails that they are asymptotically equal when they are small, but can deviate for larger values.

The most important sense in which Chapter 3 is open-ended is probably by virtue of circumstances which we shall only come to see in the two last chapters of the thesis: Namely that it is unclear how to appropriately adapt the metric theory to the setting of *causal channels* and *causal dilations* of Chapter 4, and by extension that of quantum self-testing in Chapter 5 (I will comment briefly on these issues in due time).

By itself, however, Chapter 3 leaves open the following problems:

1. Can dilational metrics be constructed by a general scheme from monotone metrics d on theories which are not necessarily purifiable? For example, does Open Problem 3.3.8 have an affirmative answer?
2. Does Theorem 3.5.2 generalise to give a non-trivial bound on \check{d} in terms of d for general monotone metrics d on purifiable theories? Can for example arguments of Refs. [KSW08a, KSW08b] (which are based on a minimax-theorem) be somehow abstracted to the general setting?

This page is intentionally left blank.

Chapter 4

Contractible Theories and Causal Dilations


§1. Introduction and Outline.

So far, we have considered channels $T : \mathbb{X} \rightarrow \mathbb{Y}$ in a theory as more or less equivalent to their underlying transformations $T : \mathcal{X} \rightarrow \mathcal{Y}$, and as such as processes which produce on a given input from the system \mathcal{X} an output in the system \mathcal{Y} . This perception, however, is too crude to adequately model a real physical device with input interface \mathbb{X} and output interface \mathbb{Y} . Indeed, in such a device it may be the case that some ports in \mathbb{Y} deliver an output already when a strict subset of the ports in \mathbb{X} have been fed with an input. We shall encode this idea in a *causal specification*, a map \mathcal{C} which associates to each output port $y \in \mathbf{ports}(\mathbb{Y})$ a set $\mathcal{C}(y) \subseteq \mathbf{ports}(\mathbb{X})$, thought of as the *causes of y*, namely the precise set of ports which require input before the output at y is available. A pair (T, \mathcal{C}) consisting of a channel T and a causal specification \mathcal{C} will be called a *causal channel*.

Causal channels force us to re-examine our understanding of dilations. If the ports in the accessible interfaces take part in an intricate causal relationship, then the ports in the hidden interfaces are involved in that relationship too. This leads to the notion of a *causal dilation* of $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$, which is simply a causal channel $(L, \mathcal{E}) : \mathbb{X} \cup \mathbb{D} \rightarrow \mathbb{Y} \cup \mathbb{E}$ whose \mathbb{Y} -marginal factors as $(T, \mathcal{C}) \parallel \mathrm{tr}_{\mathbb{D}}$ (in a sense which suitably accounts for causality). A causal dilation is meant to capture a ‘causally structured side-computation’ in the presence of (T, \mathcal{C}) ; if we understand the causal dilations of (T, \mathcal{C}) , we understand every way in which it may be immersed in its environment.

In a sense this chapter improves the causality-free theory of dilations introduced in Chapter 2. This is not to say that the theory of that chapter is rendered obsolete, merely that it is realised as the special case of a more general and stable framework. As often, greater generality does not necessarily entail greater versions of the same results, but rather entails different types of results, due to a shift in focus. In particular, having introduced the causal version of the dilational ordering (Definition 4.3.14) it will quickly become clear that its higher complexity makes it much harder to confine by dilational axioms. For example, the existence of *complete* (causal) dilations will cease to be a property of the theory in question, and instead becomes a property of the (causal) channel in question. In fact, the existence of a complete dilation will be coined as *rigidity* of the channel (Definition 4.4.3), and as we shall see in Chapter 5 this property is more or less the hallmark of quantum self-testing.

Causal Channels and Constructibility. We start in Section 4.1 by defining the concept of a *causal channel*, consisting of a channel T together with a *causal specification* \mathcal{C} . The most obvious way in which a causal specification emerges, is when the channel is represented by means of diagram like the following (each wire representing a simple interface):



The channel T is the total channel resulting from the appropriate composition of T_1, \dots, T_5 , and the causal specification \mathcal{C} is the one that associates to y_j the set of those x_i from which a directed path from left to right leads to y_j . For example, $\mathcal{C}(y_2) = \mathcal{C}(y_3) = \{x_1, x_2\}$ and $\mathcal{C}(y_4) = \emptyset$. By properties of the trashes, \mathcal{C} encodes a collection of *non-signalling conditions* to which T complies; for any subset $J \subseteq \text{ports}(\mathbb{Y})$, T is necessarily non-signalling from the interface $\mathbb{X}|_{\text{ports}(\mathbb{X}) \setminus \mathcal{C}(J)}$ (the ‘non-causes’ of J) to the interface $\mathbb{Y}|_J$. This observation tempts us to define, abstractly, a *causal channel* as any pair $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ with $T : \mathbb{X} \rightarrow \mathbb{Y}$ a channel and \mathcal{C} an additive function from subsets of $\text{ports}(\mathbb{Y})$ to subsets of $\text{ports}(\mathbb{X})$, such that T adheres to the non-signalling conditions suggested by \mathcal{C} .¹ If a causal channel can be realised using a diagram as above, we will call it *constructible*, and though the constructible causal channels will be the most important, not all imaginable causal channels are constructible (e.g. the Popescu-Rohrlich box from Example 2.1.15 is inconstructible in **QIT** when equipped with its natural causal specification). Causal channels can be composed serially and parallelly, with causal specifications composing in the obvious way, and this leads to the *category of interfaces and causal channels* (Definition 4.1.19), which will constitute our final model for physical devices.

Notions of Contraction. Whereas there is nothing inconsistent about reducing the information of a network of channels like (4.1) to the information of the causal channel (T, \mathcal{C}) , it is utterly unobvious that this reduction is reasonable. Specifically, it is unclear that we do not thereby lose the possibility of forming *contractions* of the network, that is, feeding an output to one of the inputs as long as this does not create any cycles in the network. For example, it would seem as if we should be able to feed the output at y_4 to any of the input ports x_1, x_2 or x_3 , or that of y_5 to x_2 or x_3 (provided that the corresponding systems match), but it is not trivial that the resulting causal channel can be obtained from knowledge of (T, \mathcal{C}) alone. Indeed, not only does the skeleton of the network (which we will call a *stencil*) not determine the channels T_1, T_2, \dots, T_5 uniquely, it might also be that several different stencils can be used to represent (T, \mathcal{C}) . In Section 4.2 we will see that one can define an unambiguous notion of contraction, provided that the underlying theory Θ has universal dilations (Theorem 4.2.6). This is one of the main results of this chapter,

¹Importantly, T might enjoy further non-signalling conditions than those implied by \mathcal{C} , and as such the specification \mathcal{C} cannot be derived from knowledge of T but rather formalises how causality is perceived (cf. Example 4.1.7).

and I consider it quite remarkable. In the case of **QIT**, the statement follows already by the work of Ref. [CDP09], but the proof given here is conceptually simpler and evidently has much larger scope, as it encompasses all universal theories (in particular **CIT** and all cartesian theories).

The operation we obtain from the concrete contraction of wires will be called the *standard notion of contraction*. It turns out that, similarly to the way in which causal channels can be abstracted beyond constructible ones, the standard notion of contraction can be extended to abstract *notions of contraction*, operations defined axiomatically on causal channels which swallow (‘contract’) pairs of input and output ports and yield causal channels between the remaining interfaces. This is beneficial because, as detailed in the beginning of Section 4.2.B, our construction of the standard notion of contraction will be sensitive to somewhat inessential particularities. The abstract notion is more stable, and also more widely applicable, e.g. in a thin theory a contraction is essentially a cancellation in the corresponding monoid. Finally, abstract notions of contraction will be realised to be a generalisation of *traces* in symmetric monoidal categories as introduced by Ref. [JSV96].²

A notion of contraction, abstract or standard, is not merely additional to the operations of serial and parallel composition in a theory; it actually obviates the serial composition: If $\text{--}\mathbb{X}\text{--}\boxed{(T, \mathcal{C})}\text{--}\mathbb{Y}\text{--}$ and $\text{--}\mathbb{Y}\text{--}\boxed{(S, \mathcal{D})}\text{--}\mathbb{Z}\text{--}$ are causal channels, then their serial composition can be constructed from their parallel composition $\text{--}\mathbb{X}\text{--}\boxed{(T, \mathcal{C})}\text{--}\mathbb{Y}\text{--}$ simply by contracting the interface \mathbb{Y} . This not only conceptually simplifies the complexity of operations in a theory, but is also technically simplifying in proofs.

Climax: A Hierarchy of Causal Dilations. With causal channels and contractions in place, we can in Section 4.3 elevate the theory of dilations to its full potential. *Causal dilations* of $\text{--}\mathbb{X}\text{--}\boxed{(T, \mathcal{C})}\text{--}\mathbb{Y}\text{--}$ are simply causal channels $\text{--}\mathbb{X}\text{--}\boxed{(L, \mathcal{E})}\text{--}\mathbb{Y}\text{--}$ such that

$$\text{--}\mathbb{X}\text{--}\boxed{(L, \mathcal{E})}\text{--}\mathbb{Y}\text{--} \underset{\sim \mathbb{D}}{\sim} \text{--}\mathbb{X}\text{--}\boxed{(T, \mathcal{C})}\text{--}\mathbb{Y}\text{--} \underset{\sim \mathbb{E}}{\sim} \text{--}\mathbb{D}\text{--}\boxed{\text{tr}}\text{--}\mathbb{E}\text{--}$$

, a condition which is visually similar to that of dilations in Chapter 2, but is now a condition among causal channels (Definition 4.3.1). Again, the interpretation is that the interfaces \mathbb{X} and \mathbb{Y} are *open* or *accessible* for us to probe, establishing that we interact with (T, \mathcal{C}) , and the dilation (L, \mathcal{E}) signifies a way in which the functionality described by (T, \mathcal{C}) is actually causally immersed in a larger environment, whose interfaces \mathbb{D} and \mathbb{E} are inaccessible to us. The *causal-dilational order* which we will synonymously refer to as *derivability in the environment* (Definition 4.3.14) rectifies the dilational ordering of Chapter 2. It formalises the intuition that some dilations can be derived from others by operations in the environment; this time, however, we keep track of the causal order and we have at our disposal not only serial and parallel composition, but also contractions. This finally makes it possible to treat also two-sided dilations in a satisfactory manner. Precisely, we will render a dilation $\text{--}\mathbb{X}\text{--}\boxed{(L, \mathcal{E})}\text{--}\mathbb{Y}\text{--}$ *equivalent* to the

$$\text{--}\mathbb{X}\text{--}\boxed{(L, \mathcal{E})}\text{--}\mathbb{Y}\text{--} \underset{\sim \mathbb{D}}{\sim} \text{--}\mathbb{X}\text{--}\boxed{(L, \mathcal{E})}\text{--}\mathbb{Y}\text{--} \underset{\sim \mathbb{E}}{\sim} \text{--}\mathbb{A}\text{--}\boxed{(G, \mathcal{B})}\text{--}\mathbb{B}\text{--}$$

dilations $\text{--}\mathbb{X}\text{--}\boxed{(L, \mathcal{E})}\text{--}\mathbb{Y}\text{--}$, with (G, \mathcal{B}) an arbitrary causal channel; a dilation (L', \mathcal{E}') is then *derivable* from (L, \mathcal{E}) , if a dilation equivalent to (L, \mathcal{E}) is can be contracted along some ports to yield a dilation equivalent to (L', \mathcal{E}') . The intuition is that, in the environment, it is possible to construct a network of channels which when coupled suitably to the dilation (L, \mathcal{E}) yields the dilation (L', \mathcal{E}') .

²Admittedly, I only became aware of this work a short time ago.

The definition of this relation raises an interesting question: Is it necessarily the case that the contraction of ports in a causal dilation of (T, \mathcal{C}) yields another causal dilation of (T, \mathcal{C}) ? Though this is not formally implied by the definition, a result of this kind is desirable as it consolidates the concept of causal dilations. It is also operationally relevant, since, if false, the agents controlling the environment are not free to perform operations on the hidden interfaces without the risk that these become detectable at the open interfaces. That it should actually be the case is, however, far from obvious (and as demonstrated by Example 4.3.12 it is almost false), and the most interesting result of Section 4.3 is therefore the proof that causal dilations are actually stable under contractions in the environment (Theorem 4.3.13).

Section 4.3 also contains results to the effect that causal dilations are stable under parallel compositions and contraction of the open interfaces (in combination, these two results subsume serial compositions of open interfaces); those results (Theorem 4.3.24 and Theorem 4.3.25) are of course also necessary for the concept of causal dilations to be well-behaved, but they are expectable and less subtle.

Completeness, Rigidity and Density Theorems. In stark contrast to the dilational ordering that arose in the causality-free setting of Chapter 2, axiomatic regularity of the causal-dilational ordering seems to escape into thin air. In Section 4.4 we will scratch the surface, by mostly focusing on special cases. The goal is once again to understand the large dilations w.r.t. the causal-dilational ordering, and we do this by means of *density theorems*. A class \mathbf{D} of causal dilations of (T, \mathcal{C}) is said to be *dense* if every dilation of (T, \mathcal{C}) is derivable from some dilation in \mathbf{D} . A density theorem asserts the density of a particular class of dilations. Obviously, the smaller the class, the better the theorem.

The ultimate density theorem asserts the density of a class containing just a single dilation (K, \mathcal{F}) , that is, it asserts that the dilation (K, \mathcal{F}) is a *complete* causal dilation. In general, however, causal channels will not have complete dilations, and the various dilations in the dense class will reflect genuinely different ways of implementing (T, \mathcal{C}) across the hidden interfaces. In fact, if (T, \mathcal{C}) has a complete dilation we will call it *rigid* (Definition 4.4.3), and in Chapter 5 we will see the relation of this notion to quantum self-testing.

In a cartesian theory, every causal channel turns out to be rigid (Theorem 4.4.12), but already in the theory **CIT** rigidity fails in the simplest cases – for example, we will finally see why the ‘bit refreshment’ channel from the general introduction is not rigid. Indeed, density theorems for **CIT** (Theorem 4.4.15 and Theorem 4.4.22) suggest that to study the causal-dilational ordering in **CIT** is essentially to study the various convex decompositions of channels, and in particular rigidity occurs if certain decompositions are unique (Corollary 4.4.17 and Corollary 4.4.24).

§2. Comparison to Existing Literature.

On Causality. The modelling of causality constitutes a vivid field of study in quantum foundations. Whereas much modern interest is directed towards finding a framework of ‘indefinite causal structure’ in which causality is malleable and even quantumly super-posable ([Har07, OCB12, CDPV13, Bru14]), we consider exclusively the setting of fixed causal structures, e.g. as represented by the network (4.1). The main technical obstacle in treating such networks is that it is cumbersome – boarding unfeasible – to keep track of the individual channels of the network and their intricate connections. This problem was solved in the theory **QIT** by the framework of *quantum combs* ([CDP09]), which demonstrated that a network of quantum channels can be summarised by a linear operator bearing witness only to the open ports that enter and exit the network. This linear operator is obtained from

the network by means of the Choi-Jamiołkowski isomorphism, and connecting two networks can be achieved by a binary operation among these linear operators.

The framework proposed here – namely the approach of summarising the network by a pair (T, \mathcal{C}) – was developed also with the motivation of reducing complexity. In comparison to Ref. [CDP09], it distinguishes itself by the following traits:

1. A network such as that in Eq. (4.1) is summarised not by means of a linear operator which happens to exist by virtue of the Choi-Jamiołkowski representation, but rather by its actual input-output behaviour as a channel, T , along with a specification, \mathcal{C} , of which inputs cause which outputs. This is conceptually and mathematically less awkward, and it broadens the scope to arbitrary theories.
2. The connection of two or more networks is defined not by a binary operation, but rather is derived from parallel composition along with an operation on a single network, namely *contraction*. The existences of these two operations are equivalent (since a contraction of a network can be viewed as coupling it to one consisting of identity channels), but for our purposes the approach in terms of contractions supports a cleaner presentation and method of proof.
3. The well-definedness of the contraction operation (or, equivalently, the composition of networks) is proved not using features of the Choi-Jamiołkowski representation, but rather employing universal dilations (Definition 2.4.1). As such, the approach immediately lends itself to all universal theories.

There is also a difference in spirit between the approach here and that of Ref. [CDP09]. Indeed, the main objective of the latter work is that of determining what is the class of transformations that a given network can undergo by combining it with other networks. This starts with the observation that quantum channels are the admissible transformation of states, and that networks with one ‘open slot’ (a so-called *1-comb*) are the admissible transformations of channels; the authors then proceed to define an entire hierarchy of $(n+1)$ -*combs* which act on n -*combs*. In contrast, the approach taken here is rooted in the viewpoint that all these transformations can be realised in a flat, non-hierarchical structure, in which *contraction* is a simple operation performed in a single network.

Ref. [KU17] has also introduced a framework for arguing about causality in general categorical terms; it too is based on the hierarchical viewpoint of Ref. [CDP09] (as distilled in Ref. [Per17]). Putting that difference aside, it compares to the framework introduced in this chapter by being more general in that it is able to treat indefinite causal structures as mentioned above, but less general in that it requires (among other things) transformations in the category to be determined by their action on states, and the category itself to be compact closed. In fact, this assumption is used to establish essentially an equivalent of the Choi-Jamiołkowski representation which then facilitates a treatment analogous to that of Refs. [CDP09, Per17].

Finally, all of the above comparisons pertain to the *constructible causal channels* and the *standard notion of contraction* only. As already mentioned, the framework introduced here accommodates causal channels of a more abstract character, and notions of contraction which are not necessarily physically realisable but more mathematical. It is my hope that this yields a more stable theory which might find applications elsewhere.

On the Term ‘Causal Channel’. Ref. [BGNP01] and off-spring work (e.g. Ref. [ESW02]) define the term ‘causal channel’ in a setting which can be seen as an extremely special case

of ours: Namely, a bipartite-bipartite quantum channel $\begin{array}{c} -\mathcal{X}_1- \\ -\mathcal{X}_2- \end{array} \boxed{\Lambda} \begin{array}{c} -\mathcal{Y}_1- \\ -\mathcal{Y}_2- \end{array}$ is called *causal* if it is non-signalling from x_1 to y_2 and from x_2 to y_1 . In our language, this is simply to say that Λ is compatible with the specification \mathcal{C} given by $\mathcal{C}(y_1) = \{x_1\}$ and $\mathcal{C}(y_2) = \{x_2\}$.

Contractions and Traced Monoidal Categories. Ref. [JSV96] introduced the notion of a *traced (symmetric) monoidal category*, namely a (symmetric) monoidal category equipped with an additional operation on channels generalising the trace in the symmetric monoidal category $(\mathbf{Vect}_k, \otimes, k)$. More precisely, a trace in the sense of Ref. [JSV96] maps channels $\begin{array}{c} -X- \\ -Z- \end{array} \boxed{T} \begin{array}{c} -Y- \\ -Z- \end{array}$ to channels $-X-\boxed{T'}-Y-$ by means of an operation which coheres to certain natural conditions. It turns out what is defined in this chapter as *notions of contraction* (Definition 4.2.8) can be seen a generalisation of such traces, as detailed in Remark 4.2.16. The generalisation consists in not requiring total domain for this operation, reflecting the situation that not necessarily any pair of wires is contractible just because they correspond to the same system. (For example, the output at Z might require the input at Z to be given first.) This generalisation may seem like an obvious one, but subtlety resides in the fact that those coherence conditions must now also specify the relationship between domains. (In a similar fashion, the generalisation from bounded linear operators to unbounded operators cannot be done mindlessly, but must consider what should be e.g. the domain of the operator $A + B$ given the domains of A and B .)

§3. Contributions.

The original contributions of this chapter are the following:

1. Defining the notion of a *causal channel* (Definition 4.1.3), in particular providing a general and conceptually simple way of thinking about complicated networks of channels in a theory as so-called *constructible causal channels* (Definition 4.1.15).
2. Establishing the operation of *contractions* of causal channels (Theorem 4.2.6) which offer a technically simple alternative to other frameworks ([CDP09, KU17]), and abstracting it to *notions of contraction*, which generalise traces in a symmetric monoidal category as introduced in Ref. [JSV96] (cf. Remark 4.2.16).
3. Introducing the concept of *causal dilations* (Definition 4.3.1) of a causal channel, which model the various possible environments in which the channel may be causally immersed, and defining a version of the dilational ordering for causal channel, *derivability (in the environment)*, which models the relative strength of such dilations (Definition 4.3.14). A number of stability results about causal dilations and derivability ordering is proved, most surprisingly the fact that the property of being a dilation is preserved under derivability (Theorem 4.3.13).
4. Giving a general definition of *rigidity* of a causal channel in terms of complete causal dilations (Definition 4.4.3), and establishing *density theorems* in some example theories (Section 4.4), in particular clarifying the meaning of rigidity in cartesian theories and in the classical information theory **CIT**.

§4. The Metric Aspect.

In Chapter 3 was rolled out a general theory of metric structure in theories. As with dilations, it is only reasonable to revise this theory in light of the shift from channels to causal channels. I shall not take on such a revision here, but rather leave it for future work. It seems that there there are basically two challenges to overcome.

First of all, whereas it is sensible to leave the property of *parallel invariance* of metrics unchanged in the causal setting, the *serial monotonicity* should be upgraded to monotonicity under contractions; if we take seriously the idea that a contraction is a valid operation that can be performed, and if a metric is to abstractly quantify distinguishability between two causal channels, this is only logical. It is probably easy to define such a contractually monotone metric by forming a supremum over various contractions, but it would seem difficult to calculate.³ A related idea has also been considered in Ref. [CT09], and in Ref. [CDP09] in the concept of *comb distance* (herein, a formula is derived).

Secondly, since the concept of dilations has changed, so should the concept of dilationality for a metric on causal channels. We do not have an equivalent of Proposition 3.3.6 to guide us, since we generally lack complete causal dilations, and it is not obvious how (or even in what sense) dilationality can be achieved.

³Note that e.g. the variational distance d_1 on CIT is not contractually monotone already.

4.1 Causal Channels and Constructibility

In this section, we define *causal channels* and consider some basic examples. An especially tangible and important class of causal channels is that of *constructible causal channels*, which is meant to formalise the class of causal channels which can actually be ‘built’ in the real world.

4.1.A Causal Specifications and Causal Channels

A causal channel $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ consists of two parts, a channel $T : \mathbb{X} \rightarrow \mathbb{Y}$ and a causal specification \mathcal{C} . The channel represents the input-output behaviour without any regards to causality, and the specification specifies the causal relationships between the ports of \mathbb{X} and \mathbb{Y} .

Definition 4.1.1. (Causal Specifications.)

Let \mathbb{X} and \mathbb{Y} be interfaces in Θ . A *causal specification from \mathbb{X} to \mathbb{Y}* is a map $\mathcal{C} : \mathcal{P}(\text{ports}(\mathbb{Y})) \rightarrow \mathcal{P}(\text{ports}(\mathbb{X}))$ such that

- $\mathcal{C}(\emptyset) = \emptyset$ and
- $\mathcal{C}(J_1 \cup J_2) = \mathcal{C}(J_1) \cup \mathcal{C}(J_2)$ for any $J_1, J_2 \subseteq \text{ports}(\mathbb{Y})$.

We write $\mathcal{C} : \mathbb{X} \rightarrow \mathbb{Y}$ to denote that \mathcal{C} is a causal specification from \mathbb{X} to \mathbb{Y} . ■

Remark 4.1.2. (Specifying a Specification.)

The *additivity requirement* $\mathcal{C}(J_1 \cup J_2) = \mathcal{C}(J_1) \cup \mathcal{C}(J_2)$ implies that \mathcal{C} is determined by its value on single ports, $\mathcal{C}(\{y\})$. We tend to write $\mathcal{C}(y)$ rather than $\mathcal{C}(\{y\})$, and by additivity we thus have $\mathcal{C}(J) = \bigcup_{y \in J} \mathcal{C}(y)$ for any $J \subseteq \text{ports}(\mathbb{Y})$. ✦

Given an interface \mathbb{Z} , the subsets of $\text{ports}(\mathbb{Z})$ are in natural correspondence with the sub-interfaces of \mathbb{Z} . To ease notation, we might therefore sometimes think of a causal specification as being defined on sub-interfaces rather than subsets of port names, as in the following definition:

Definition 4.1.3. (Causal Channels.)


Let $T : \mathbb{X} \rightarrow \mathbb{Y}$ be a channel in Θ , and $\mathcal{C} : \mathbb{X} \rightarrow \mathbb{Y}$ a causal specification. T and \mathcal{C} are said to be *compatible* if for any sub-interface $\mathbb{Y}_0 \subseteq \mathbb{Y}$, T is non-signalling from $\mathbb{X} \setminus \mathcal{C}(\mathbb{Y}_0)$ to \mathbb{Y}_0 , i.e. there exists a channel $T' : \mathcal{C}(\mathbb{Y}_0) \rightarrow \mathbb{Y}_0$ such that

$$\begin{array}{c} \text{---}\mathcal{C}(\mathbb{Y}_0)\text{---} \\ \text{---}\mathbb{X} \setminus \mathcal{C}(\mathbb{Y}_0)\text{---} \end{array} \boxed{T} \begin{array}{c} \text{---}\mathbb{Y}_0\text{---} \\ \text{---}\mathbb{Y} \setminus \mathbb{Y}_0\text{---} \end{array} \boxed{\text{tr}} = \begin{array}{c} \text{---}\mathcal{C}(\mathbb{Y}_0)\text{---} \\ \text{---}\mathbb{X} \setminus \mathcal{C}(\mathbb{Y}_0)\text{---} \end{array} \boxed{T'} \text{---}\mathbb{Y}_0\text{---} \boxed{\text{tr}} . \quad (4.2)$$

A *causal channel from \mathbb{X} to \mathbb{Y}* is a pair (T, \mathcal{C}) consisting of a channel $T : \mathbb{X} \rightarrow \mathbb{Y}$ and a compatible causal specification $\mathcal{C} : \mathbb{X} \rightarrow \mathbb{Y}$. ■

To verify that a channel T is compatible with a causal specification \mathcal{C} , one must in principle check $2^{|\mathbb{Y}|}$ non-signalling conditions, one for each sub-interface of \mathbb{Y} (though they are always trivial for the sub-interfaces \mathbb{I} and \mathbb{Y}). If the theory Θ has states on every system, it actually suffices to check the condition for simple sub-interfaces, reducing the number to $|\mathbb{Y}|$, but this can still be cumbersome. An important class of examples of causal channels is therefore those which are visually presented in such a way that the non-signalling conditions are obvious:

Example 4.1.4. (A Generic Causal Channel.)
 In any theory Θ , the channel T given by



for suitably composable channels T_1, \dots, T_5 is compatible with the causal specification \mathcal{C} that arises from ancestry in the network, when thinking of it as a directed (from left to right) graph. Precisely, \mathcal{C} is given by $\mathcal{C}(y_j) = \{x_1, x_2\}$ for $j = 1, 2, 3$, $\mathcal{C}(y_4) = \emptyset$ and $\mathcal{C}(y_5) = \{x_3\}$. The non-signalling conditions follow from properties of the trashes (Lemma 1.1.15), which in each case eliminate those input ports which are not ancestral to the non-trashed output ports. \blacklozenge

Example 4.1.5. (Bell-Channels.)
 Recall from Example 2.1.14 that a (*bipartite*) *Bell-channel* is a channel T of the form




We may augment it to a causal channel by defining the specification \mathcal{C} again by ancestry in the network: $\mathcal{C}(\emptyset) = \emptyset$, $\mathcal{C}(y_A) = \{x_A\}$, $\mathcal{C}(y_B) = \{x_B\}$ and $\mathcal{C}(\{y_A, y_B\}) = \{x_A, x_B\}$. \blacklozenge

It is important to realise that a channel T may enjoy non-signalling conditions which are additional to those implied by a given compatible specification \mathcal{C} . As such, the causal specification is an integral part of a causal channel (T, \mathcal{C}) and cannot be derived from T itself:

Example 4.1.6. (Primitive Causal Channels.)
 Any channel $T : \mathbb{X} \rightarrow \mathbb{Y}$ is compatible with the *primitive specification* $\mathcal{C}_0 : \mathbb{X} \rightarrow \mathbb{Y}$, according to which $\mathcal{C}_0(J) = \text{ports}(\mathbb{X})$ for all non-empty subsets $J \subseteq \text{ports}(\mathbb{Y})$. A causal channel (T, \mathcal{C}_0) is called *primitive* if it is equipped with the primitive specification \mathcal{C}_0 . Intuitively, the primitive specification expresses that the all ports of \mathbb{X} must be fed with an input before any outputs appear at \mathbb{Y} . This causal specification reflects how we implicitly thought of channels in all preceding chapters. \blacklozenge

Example 4.1.7. (One-Time Pad.)
 Let P be the channel in **CIT** given by



where κ is the state on $\{0, 1\} \times \{0, 1\}$ given by $\kappa = \frac{1}{2}\delta_0 \otimes \delta_0 + \frac{1}{2}\delta_1 \otimes \delta_1$ (i.e. κ is two copies of a uniformly random bit), and where XOR is the ‘exclusively OR’, i.e. the deterministic function which adds the inputs modulo 2. The channel P represents the so-called *one-time pad*, a cryptographic device used for encryption (see any introductory book on cryptography, e.g. [HPS08]). In the one-time pad, a key bit is chosen at random and copied; one copy of the key bit is kept in memory (the *k*-port) for later decryption, and the other copy is used to decide on whether or not to apply a bit flip on the message bit coming in at the *m*-port, thus obtaining an encrypted cipher bit at the *c*-port. Clearly, we want to equip P with the causal specification that arises from ancestry: *m* is a cause of *c*, and *k* has no causes. And indeed, this specification is guaranteed to be compatible with P : Trashing the *c*-port results in a trash on the *m*-port.

However, the entire point of the one-time pad is that trashing the *k*-port also results in a trash of the *m*-port, that is, to an individual who does not know the key bit, the cipher bit will look completely random, independently of the message bit that was to be encrypted. Thus, the channel P is non-signalling from the input port to either of its output ports. In fact, the channel P is even symmetric under permutation of the two output ports: One cannot tell which bit is the key bit, and which is the cipher bit. Asymmetry between the ports only arises in the presence of a chosen causal specification, and that specification is not derivable from the input-output behaviour of the channel P itself. \blacklozenge

What the above examples demonstrate is that causal specifications are not merely about non-signalling conditions, but rather about how causality is *perceived*.

Before proceeding, let us consider three more examples of causal channels.

Example 4.1.8. (‘Linked’ Parallel Composition.)

Let $-\mathcal{X}_1-\boxed{T_1}-\mathcal{Y}_1-$ and $-\mathcal{X}_2-\boxed{T_2}-\mathcal{Y}_2-$ be channels in Θ , and consider their parallel composition

$$-\mathcal{X}_1-\boxed{T}-\mathcal{Y}_1- \quad := \quad -\mathcal{X}_1-\boxed{T_1}-\mathcal{Y}_1- \quad (4.6)$$

$$-\mathcal{X}_2-\boxed{T}-\mathcal{Y}_2- \quad \quad \quad -\mathcal{X}_2-\boxed{T_2}-\mathcal{Y}_2-$$

Of course, we can equip T with the causal \mathcal{C} specification that arises from the ancestry on the right hand side ($\mathcal{C}(y_1) = \{x_1\}$ and $\mathcal{C}(y_2) = \{x_2\}$), but T is also compatible with another causal specification, namely \mathcal{C}' given by $\mathcal{C}'(y_1) = \{x_1\}$ and $\mathcal{C}'(y_2) = \{x_1, x_2\}$. (In general, enlarging the cause sets of a compatible specification yields another compatible specification.) According to the specification \mathcal{C}' , the output at y_1 occurs as soon as x_1 has been fed with an input, but the output at y_2 requires both x_1 and x_2 to be given an input before it shows any output.

Though (T, \mathcal{C}) is not the causal channel whose causality derives from ancestry in the network from (4.6), it does arrive from the ancestry in a different network, namely

$$-\mathcal{X}_1-\boxed{T_1}-\mathcal{Y}_1- \quad (4.7)$$

$$\quad \quad \quad \mathbf{1}-\boxed{T_2}-\mathcal{Y}_2-$$

where $\mathbf{1}$ is the trivial system. In effect, the trivial system acts to *stall* the execution of T_2 until T_1 has been applied. These two alternative depictions would not have been clear without the pedantic discussion about interfaces and channels versus systems and transformations in Section 1.3.B. \blacklozenge

Example 4.1.9. (States and Trashes as Causal Channels.)

If a causal $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ has trivial output interface, $\mathbb{Y} = \mathbb{I}$, the specification can only be the one given by $\mathcal{C}(\emptyset) = \emptyset$, and its channel T must be $\text{tr}_{\mathbb{X}}$. Similarly, when its input interface is trivial, $\mathbb{X} = \mathbb{I}$, its specification must be given by $\mathcal{C}(J) = \emptyset$ for all $J \subseteq \text{ports}(\mathbb{Y})$ and its channel T must be a state. In short, trashes and states have unique causal specifications when thought of as channels to, respectively from, the trivial interface. (Incidentally, these unique specifications are in both cases primitive.)

However, we need not think of a trash as mapping to the trivial interface, nor of a state as mapping from it. For instance, we can write a trash $\text{tr}_{\mathcal{Z}}$ as

$$-\mathcal{Z}-\boxed{\text{tr}}-\mathbf{1}- \quad (4.8)$$

and give it the (primitive) causal specification \mathcal{S} according to which $\mathcal{S}(\mathbf{1}) = \mathbf{z}$. As such, the output port becomes an *indicator* for whether or not something was trashed. Similarly, a state s on \mathcal{Y} can be represented by the channel

$$-\mathbf{1}-\boxed{s}-\mathcal{Y}-, \quad (4.9)$$

and when endowed with the (primitive) specification \mathcal{S} given by $\mathcal{S}(y) = \mathbf{1}$, the input port becomes an *activator* for the state, releasing it at the output port.

Of course, indicators and activators can act in intricate ways, indicating or activating only a subset of ports. Also, they can be mounted to any sorts of channels, not just trashes and states (indeed the stalling phenomenon from Example 4.1.8 can be seen as an example of this). ◆

Example 4.1.10. (The Popescu-Rohrlich Box as a Causal Channel.)

Recall the PR Box from Example 2.1.15, the classical channel $\begin{matrix} x_A-\{0,1\} \\ x_B-\{0,1\} \end{matrix} \boxed{P} \begin{matrix} \{0,1\}-y_A \\ \{0,1\}-y_B \end{matrix}$ determined by the probability distributions $(P^{x_A, x_B})_{x_A, x_B \in \{0,1\}}$ given by

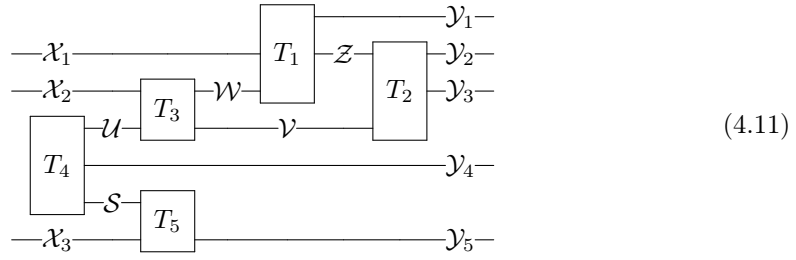
$$P^{x_A, x_B}(y_A, y_B) = \begin{cases} \frac{1}{2} & \text{for } y_A \oplus y_B = x_A \cdot x_B \\ 0 & \text{for } y_A \oplus y_B \neq x_A \cdot x_B \end{cases}, \quad (4.10)$$

with \oplus denoting addition modulo 2. We saw there that this channel was non-signalling from x_A to y_B and from x_B to y_A , and as such it is compatible with a *local* causal specification \mathcal{C} , namely the one given by $\mathcal{C}(y_i) = \{x_i\}$ for $i = A, B$. Thus, we have a causal channel (P, \mathcal{C}) . ◆


4.1.B Constructible Causal Channels

All of the above examples of causal channels, with the exception of the PR box in Example 4.1.10, were presented by means of a network of channels which defined simultaneously the channel and its causal specification. We now discuss how to formalise this idea in the concept of a *constructible* causal channel.

To exemplify, consider the depiction



from Example 4.1.4. It is comprised of two pieces of information. One of them is purely graph-theoretic, namely the *stencil*



(directed from left to right). The other piece of information rests upon the theory Θ , namely the *filling*, which assigns to each wire a simple interface in Θ and to each box a suitably compatible channel between the adjacent interfaces.

A general stencil can be described by a *directed acyclic graph* (for short, *DAG*), in which we distinguish two kinds of vertices; *boxes*, which are vertices to be filled with a channel, and *ports* (represented in (4.12) by bullets):

Definition 4.1.11. (Stencils.)

A *stencil* is a triple $(G, W_{\text{in}}, W_{\text{out}})$, where⁴

- G is a finite, non-empty DAG with no isolated vertices;
- W_{in} is a set of edges in G , each of which comes from a source in G that has no other outgoing edges;
- W_{out} is a set of edges in G , each of which go to a sink in G which has no other incoming edges.

An edge in G is called a *wire*. The edges in W_{in} are called *input wires* and the edges in W_{out} are called *output wires*. A vertex which is the source of a wire in W_{in} , or the sink of a wire in W_{out} , is called a *port*, and every other vertex in G is called a *box*. ■

The reader is encouraged to consider how these notions pan out in the example illustrated by Eq. (4.12).

⁴Recall that a *source* in a directed graph is a vertex with no incoming edges, and that a *sink* is a vertex with no outgoing edges. An *isolated vertex* is a vertex which is both a source and a sink.

By abuse of notation, we use the symbol G to represent the whole stencil $(G, \mathcal{W}_{\text{in}}, \mathcal{W}_{\text{out}})$, and we write $\mathcal{W}(G)$, $\mathcal{B}(G)$, $\mathcal{W}_{\text{in}}(G)$ and $\mathcal{W}_{\text{out}}(G)$ and to denote respectively the sets of wires, boxes, input wires and output wires in the stencil.

Let us also denote, for a box $b \in \mathcal{B}(G)$, by $\text{In}(b)$ and $\text{Out}(b)$ the set of wires which go to, respectively from, b . Fillings of a stencil are now straightforward to define:

Definition 4.1.12. (Stencil Fillings.)

Let Θ be a theory. A Θ -filling of the stencil G is a pair $\mathfrak{F} = ((\mathbb{Z}_w)_{w \in \mathcal{W}(G)}, (T_b)_{b \in \mathcal{B}(G)})$, where

1. $(\mathbb{Z}_w)_{w \in \mathcal{W}(G)}$ is a collection of simple interfaces in Θ , indexed by the wires $w \in \mathcal{W}(G)$; we require that \mathbb{Z}_w and $\mathbb{Z}_{w'}$ can only be identical if a path in G leads from w to w' , or from w' to w .
2. $(T_b)_{b \in \mathcal{B}(G)}$ is a collection of channels in Θ , indexed by the boxes $b \in \mathcal{B}(G)$, such that T_b is a channel from the interface $\bigcup_{w \in \text{In}(b)} \mathbb{Z}_w$ to the interface $\bigcup_{w \in \text{Out}(b)} \mathbb{Z}_w$.

■

Remark 4.1.13. The requirement about distinctness of the simple interfaces \mathbb{Z}_w is so as to ensure that on the one hand identical interfaces are never composed in parallel (that was explicitly forbidden), while on the other hand an input wire and output wire may correspond to the same interface. ✖

Now, it is intuitively clear that a stencil G together with a filling \mathfrak{F} defines in a natural way a total channel

$$\mathfrak{F}[G] : \bigcup_{w \in \mathcal{W}_{\text{in}}(G)} \mathbb{Z}_w \rightarrow \bigcup_{w \in \mathcal{W}_{\text{out}}(G)} \mathbb{Z}_w \quad (4.13)$$

from the interface of input wires to the interface of output wires. We will call $\mathfrak{F}[G]$ the *value of \mathfrak{F} on G* . The formal groundwork needed for this construction, however, is tedious, and the reader is referred to the original work of Ref. [JS91] for the details.

We can now give a more precise definition of what is meant by a constructible causal channel. For vertices v' and v in a directed graph G , it is customary to write $v' \rightarrow v$ if there is a path in G from v' to v . We shall write also $w' \rightarrow w$ for wires (edges) w' and w , if there is a path in G which includes the wire w' before the wire w (or if $w' = w$). As such, for any output wire $w \in \mathcal{W}_{\text{out}}(G)$, the set

$$\mathcal{A}_G(w) := \{w' \in \mathcal{W}_{\text{in}}(G) \mid w' \rightarrow w\} \quad (4.14)$$

is the set of input wires ancestral to w .

Definition 4.1.14. (Stencil-Representability.)

Let G be a stencil. We say that a causal channel $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ in Θ is *representable on G* if there exists a Θ -filling \mathfrak{F} of G such that $(T, \mathcal{C}) = (\mathfrak{F}[G], \mathcal{C}_G)$, where $\mathfrak{F}[G]$ is the value of \mathfrak{F} on G , and where \mathcal{C}_G is the causal specification given by ancestry in G , i.e. $\mathcal{C}_G(\mathbb{Z}_w) = \bigcup_{w' \in \mathcal{A}_G(w)} \mathbb{Z}_{w'}$ for $w \in \mathcal{W}_{\text{out}}(G)$, with \mathcal{A}_G as in Eq. (4.14). ■

Definition 4.1.15. (Constructible Causal Channels.)

A causal channel is called *constructible* if it is representable on some stencil. ■

It is not a priori clear that there even exist examples of causal channels which are not constructible. Observe however the following: If a channel $\begin{array}{c} \mathcal{X}_A \\ \boxed{T} \\ \mathcal{X}_B \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ is compatible with the *local* causal specification \mathcal{C} given by $\mathcal{C}(y_i) = \{x_i\}$, and if (T, \mathcal{C}) is constructible, then T must be a Bell-channel, i.e. of the form $\begin{array}{c} \mathcal{X}_A \\ \boxed{s} \\ \mathcal{X}_B \end{array} \begin{array}{c} \boxed{T_A} \\ \boxed{T_B} \\ \mathcal{Y}_B \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$. For now, I leave a formal graph-theoretic argument to the reader (since we will see in the next subsection an alternative proof technique by an *induction principle*, cf. Example 4.1.27), but the idea is basically that if (T, \mathcal{C}) is representable on G and if G is chosen to have a minimal number of boxes among all such stencils, then one can argue by contradiction that G must be either

$$\begin{array}{c} \text{---} \\ \text{---} \end{array}, \quad \begin{array}{c} \text{---} \\ \boxed{} \\ \text{---} \end{array}, \quad \begin{array}{c} \boxed{} \\ \text{---} \\ \text{---} \end{array}, \quad \text{or} \quad \begin{array}{c} \text{---} \\ \boxed{} \\ \text{---} \\ \boxed{} \\ \text{---} \end{array}, \quad (4.15)$$

and in either case T is a Bell-channel.

This observation immediately shows that the classic counterexamples to being a Bell-channel are in fact counterexamples to constructibility:

Example 4.1.16. (Inconstructibility of the Popescu-Rohrlich Box in **QIT**.)

The PR box is inconstructible in **QIT** when given the local specification from Example 4.1.10. This is because, as mentioned in Example 2.1.15, it follows from the work of Cirelson ([Cir80]) that the PR box does not have the form of a Bell-channel in **QIT**. ♦

Example 4.1.17. (Inconstructibility of the CHSH-Behaviour in **CIT**.)

As mentioned in Example 2.1.14, the work of Ref. [Bel64], as simplified in Ref. [CHSH69], gives an example of a Bell-channel in **QIT** which has classical inputs and outputs but is not a Bell-channel, i.e. is not constructible as a causal channel, in **CIT**.

Specifically, it is the causal channel given by

$$\begin{array}{c} \mathbb{C}^{X_A} \\ \text{---} \end{array} \begin{array}{c} \boxed{\Lambda_A} \\ \text{---} \end{array} \begin{array}{c} \mathbb{C}^{Y_A} \\ \text{---} \end{array}, \quad (4.16)$$

$$\begin{array}{c} \boxed{\psi} \\ \text{---} \end{array} \begin{array}{c} \mathbb{C}^2 \\ \text{---} \end{array} \begin{array}{c} \boxed{\Lambda_B} \\ \text{---} \end{array} \begin{array}{c} \mathbb{C}^{Y_B} \\ \text{---} \end{array}$$

with $X_A = X_B = \{0, 1\}$ and $Y_A = Y_B = \{+1, -1\}$, where ψ is the maximally entangled state on $\mathbb{C}^2 \otimes \mathbb{C}^2$ represented by the vector $|\psi\rangle = \frac{|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle}{\sqrt{2}}$, and where Λ_i is the ensemble of projective measurements $\Pi_i^{x_i}$ on \mathbb{C}^2 (indexed by $x_i \in X_i$) for which $\Pi_A^{x_A}(\pm 1)$ are the two projections corresponding to the orthonormal basis $(|0\rangle, |1\rangle)$ or $(\frac{|0\rangle + |1\rangle}{\sqrt{2}}, \frac{|0\rangle - |1\rangle}{\sqrt{2}})$ depending on whether $x_A = 0$ or $x_A = 1$, respectively, and for which $\Pi_B^{x_B}(\pm 1)$ are the projections corresponding to the orthonormal basis $(\frac{\cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle}{\sqrt{2}}, \frac{\sin(\pi/8)|0\rangle - \cos(\pi/8)|1\rangle}{\sqrt{2}})$ or $(\frac{\cos(\pi/8)|0\rangle - \sin(\pi/8)|1\rangle}{\sqrt{2}}, \frac{\sin(\pi/8)|0\rangle + \cos(\pi/8)|1\rangle}{\sqrt{2}})$ depending on whether $x_B = 0$ or $x_B = 1$, respectively. One can check that the resulting channel (4.16) is the classical channel $P = (P^{x_A, x_B})_{x_A, x_B \in \{0, 1\}}$ given by

$$P^{x_A, x_B}(y_A, y_B) = \frac{1}{4} + \frac{y_A y_B}{4} \frac{(-1)^{x_A \cdot x_B}}{\sqrt{2}}, \quad (4.17)$$

and it can be verified by evaluation against a convex-linear functional (known as the *CHSH game*) that this channel is not the convex combination of products of deterministic functions, hence not a Bell-channel in **CIT**. \blacklozenge

More recent work ([CS18, CS20]) implies in a similar way that some constructible causal channels in \mathbf{QIT}^∞ are inconstructible in **QIT**.

The interested reader may prove as exercise that the phenomenon of inconstructibility is not bound to occur in every theory:⁵

Theorem 4.1.18. *If Θ is a cartesian theory, every causal channel in Θ is constructible.*

4.1.C The Category of Causal Channels – Relative Constructibility

Recall that a *primitive* causal channel $(T, \mathcal{C}_0) : \mathbb{X} \rightarrow \mathbb{Y}$ is one whose specification renders every port in \mathbb{X} a cause of any port in \mathbb{Y} : $\mathcal{C}_0(J) = \text{ports}(\mathbb{X})$ for all non-empty $J \subseteq \text{ports}(\mathbb{Y})$. This notion of causality reflects how we (implicitly) thought of channels in the preceding chapters. If we think about it, what we did just above when defining constructible causal channels was again to think of each channel filling a box in the stencil as a primitive causal channel. Through an intuitive notion of how those causal channels should compose, a total causal channel thus emerged from the stencil and its filling.

It is natural and beneficial to explicitly articulate the notions of parallel and serial composition of causal channels which gives substance to this intuition. Effectively, we thereby achieve a *new* (partial) symmetric monoidal category:

Definition 4.1.19. ($\text{ICC}(\Theta)$.)

The *category of interfaces and causal channels in Θ* , denoted $\text{ICC}(\Theta)$, is the partial⁶ symmetric monoidal category which has as its objects the interfaces in Θ , and as morphisms from \mathbb{X} to \mathbb{Y} the causal channels from \mathbb{X} to \mathbb{Y} . Its serial and parallel composition is given as follows:

- The serial composition of $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ with $(S, \mathcal{D}) : \mathbb{Y} \rightarrow \mathbb{Z}$ is given by $(S \circ T, \mathcal{D}\mathcal{C})$, where $\mathcal{D}\mathcal{C} : \mathbb{X} \rightarrow \mathbb{Z}$ is the causal specification given by the functional composition $\mathcal{C} \circ \mathcal{D}$, i.e.

$$(\mathcal{D}\mathcal{C})(J) = \mathcal{C}(\mathcal{D}(J)) \quad (4.18)$$

for $J \subseteq \text{ports}(\mathbb{Z})$.⁷ The identity on the interface \mathbb{X} is given by $(\text{id}_{\mathbb{X}}, \mathcal{I}_{\mathbb{X}})$, where $\mathcal{I}_{\mathbb{X}}(J) = J$ for all $J \subseteq \text{ports}(\mathbb{X})$.

⁵It is easiest to start with a small example, e.g. by proving that a bipartite channel with local causal specification is a Bell-channel (which in the case of cartesian theories means that it factors).

⁶In the same sense as discussed towards the end of Section 1.3.B.

⁷This may look slightly confusing. The issue is ultimately that causal specifications have a ‘contravariant’ nature reversing the order in their definition, e.g. a causal specification from \mathbb{X} to \mathbb{Y} is a map from $\mathcal{P}(\text{ports}(\mathbb{Y}))$ to $\mathcal{P}(\text{ports}(\mathbb{X}))$.

- If the underlying channels are pairwise parallelly composable, then $(T_1, \mathcal{C}_1) : \mathbb{X}_1 \rightarrow \mathbb{Y}_1$ and $(T_2, \mathcal{C}_2) : \mathbb{X}_2 \rightarrow \mathbb{Y}_2$ are parallelly composable, and the parallel composition is $(T_1 \parallel T_2, \mathcal{C}_1 \parallel \mathcal{C}_2)$, where $\mathcal{C}_1 \parallel \mathcal{C}_2 : \mathbb{X}_1 \parallel \mathbb{X}_2 \rightarrow \mathbb{Y}_1 \parallel \mathbb{Y}_2$ is the causal specification given by

$$(\mathcal{C}_1 \parallel \mathcal{C}_2)(J_1 \cup J_2) = \mathcal{C}_1(J_1) \cup \mathcal{C}_2(J_2) \quad (4.19)$$

for $J_1 \subseteq \text{ports}(\mathbb{X}_1)$ and $J_2 \subseteq \text{ports}(\mathbb{X}_2)$. ■

Remark 4.1.20. One must of course check the appropriate compatibility requirements (e.g. if T is compatible with \mathcal{C} and S with \mathcal{D} , then $S \circ T$ is compatible with $\mathcal{D}\mathcal{C}$), but they are both obvious. ✦

From now on, we will work in the category $\text{ICC}(\Theta)$. In fact, the remainder of the thesis is to a great extent about revisiting and modifying the theory of dilations to the category of interfaces and causal channels, $\text{ICC}(\Theta)$, in place of the category of interfaces and channels, $\text{IC}(\Theta)$. We will use the same sort of pictorial syntax, writing e.g. $-\mathbb{X}-\boxed{(T, \mathcal{C})}-\mathbb{Y}-$ to denote the causal channel $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$.

Let us re-examine in this fresh light the concept of constructible causal channels. With the general composition of causal channels at our disposal, we can of course build networks where the constituents are not channels, but causal channels: From now on, we will thus assume a filling \mathfrak{F} of a stencil G to consist of simple interfaces $(\mathbb{Z}_w)_{w \in \mathcal{W}(G)}$ and causal channels $((T_b, \mathcal{C}_b))_{b \in \mathcal{B}(G)}$, and we will denote by $\mathfrak{F}[G]$ the total causal channel that arises from filling the stencil G according to \mathfrak{F} .⁸

Definition 4.1.21. (Relative Constructibility.)

Let \mathbf{P} be a class of causal channels in Θ . We say that a causal channel (T, \mathcal{C}) is *constructible from \mathbf{P}* if it is of the form $\mathfrak{F}[G]$ for some stencil G and some filling \mathfrak{F} , all of whose causal channels belong to \mathbf{P} . ■

The notion of (absolute) constructibility is obviously recovered as follows:

Proposition 4.1.22. (*Absolute Constructibility as Special Relative Constructibility.*)

A causal channel is constructible in the sense of Definition 4.1.15 precisely if it is constructible from primitive causal channels in the sense of Definition 4.1.21.

The notion of relative constructibility can be recast as a minimality notion. Let us denote by $\text{Const}(\mathbf{P})$ the class of causal channels constructible from a given class \mathbf{P} . Let us say that a class of causal channels \mathbf{C} is *constructibly closed* if $\text{Const}(\mathbf{C}) \subseteq \mathbf{C}$.

One can prove (by induction on the complexity of DAGs) that a class is constructibly closed if and only if it contains identities between simple interfaces and is closed under serial and parallel composition. With this equivalence in mind, it is a relatively simple exercise to show the following:

⁸The theorem of Ref. [JS91] that yielded a well-defined value $\mathfrak{F}[G]$ still applies, since what we have done is effectively to substitute the category for a different category.

Theorem 4.1.23. (*Principle of Induction.*)

The class $\text{Const}(\mathbf{P})$ is the smallest class of causal channels which contains \mathbf{P} and is constructibly closed. More precisely,⁹ $\text{Const}(\mathbf{P})$ contains \mathbf{P} and is constructibly closed, and if \mathbf{C} is any class of causal channels which contains \mathbf{P} and identities between simple interfaces, and which is closed under serial and parallel composition, then $\text{Const}(\mathbf{P}) \subseteq \mathbf{C}$.

To showcase the induction technique, we prove the following result which is often useful:

Lemma 4.1.24. (*Factorisation of Constructible Channels.*)

If $\boxed{\text{---}\mathbb{X}\text{---}(T, \mathcal{C})\text{---}\mathbb{Y}\text{---}}$ is constructible, then for any sub-interface $\mathbb{Y}_0 \subseteq \mathbb{Y}$, there exist constructible causal channels (T_1, \mathcal{C}_1) and (T_2, \mathcal{C}_2) such that

$$\boxed{\text{---}\mathbb{X}\text{---}(T, \mathcal{C})\text{---}\mathbb{Y}\text{---}} = \begin{array}{c} \text{---}\mathcal{C}(\mathbb{Y}_0)\text{---} \boxed{(T_1, \mathcal{C}_1)} \text{---}\mathbb{Y}_0\text{---} \\ \text{---}\mathbb{X} \setminus \mathcal{C}(\mathbb{Y}_0)\text{---} \boxed{(T_2, \mathcal{C}_2)} \text{---}\mathbb{Y} \setminus \mathbb{Y}_0\text{---} \end{array} \quad (4.20)$$

Remark 4.1.25. Though visually similar, this statement has nothing to do with the DiVincenzo Property (Definition 2.1.17), as clarified by the proof. It holds regardless of whether the theory Θ has the property or not, for it is ultimately a graph-theoretic statement. Rather, the DiVincenzo Property would be the statement that if $\boxed{\text{---}\mathcal{X}_1\text{---}(T, \mathcal{C})\text{---}\mathcal{Y}_1\text{---}}$ is a causal channel for which $\mathcal{C}(\{y_1\}) = \{x_1\}$, $\mathcal{C}(\{y_2\}) = \{x_1, x_2\}$, then (T, \mathcal{C}) is constructible. In other words, it would say that constructibility follows from a property of the causal specification alone. (Lemma 4.1.24 would then provide a factorisation.) \clubsuit

Proof. We prove the result by induction on constructible channels.

Let $\pi((T, \mathcal{C}))$ be the predicate, ranging over all causal channels $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ in Θ , which asserts that for all sub-interfaces $\mathbb{Y}_0 \subseteq \mathbb{Y}$ there exist constructible channels (T_1, \mathcal{C}_1) and (T_2, \mathcal{C}_2) satisfying (4.20). Given a causal channel (T, \mathcal{C}) , the statement $\pi((T, \mathcal{C}))$ is either true or false. What we wish to prove is that it is true for every constructible channel (T, \mathcal{C}) . Consider the class of causal channels

$$\mathbf{C} := \{(T, \mathcal{C}) \mid \pi((T, \mathcal{C}))\} \quad (4.21)$$

for which $\pi((T, \mathcal{C}))$ is true. By virtue of Theorem 4.1.23, we can show that \mathbf{C} contains all constructible channels by showing that \mathbf{C} contains every primitive causal channel and that \mathbf{C} is closed under serial and parallel composition (whenever these are defined).

It is easy to see that \mathbf{C} contains all primitive channels: If (T, \mathcal{C}) is primitive and $\mathbb{Y}_0 \subseteq \mathbb{Y}$, then either $\mathbb{Y}_0 = \mathbb{I}$ or $\mathcal{C}(\mathbb{Y}_0) = \mathbb{X}$; in either case, there is really nothing to show. More precisely, in the case $\mathbb{Y}_0 = \mathbb{I}$ we may take $\mathbb{H} = \mathbb{I}$, $(T_1, \mathcal{C}_1) = \text{id}_{\mathbb{I}}$ and $(T_2, \mathcal{C}_2) = (T, \mathcal{C})$, and in the case $\mathcal{C}(\mathbb{Y}_0) = \mathbb{X}$ we may take $\mathbb{H} = \mathbb{Y} \setminus \mathbb{Y}_0$, $(T_1, \mathcal{C}_1) = (T, \mathcal{C})$ and $(T_2, \mathcal{C}_2) = \text{id}_{\mathbb{Y} \setminus \mathbb{Y}_0}$ (with the identity specification $\mathcal{S}_{\mathbb{Y} \setminus \mathbb{Y}_0}$).

It is also easy to see that \mathbf{C} is closed under parallel composition. Let $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ and $(S, \mathcal{D}) : \mathbb{Z} \rightarrow \mathbb{W}$ belong to \mathbf{C} and be parallelly composable. Consider the parallel composition $(T, \mathcal{C}) \parallel (S, \mathcal{D}) : \mathbb{X} \cup \mathbb{Z} \rightarrow \mathbb{Y} \cup \mathbb{W}$. Any sub-interface \mathbb{A} of the output interface $\mathbb{Y} \cup \mathbb{W}$ is of the form $\mathbb{Y}_0 \cup \mathbb{W}_0$ for sub-interfaces $\mathbb{Y}_0 \subseteq \mathbb{Y}$ and $\mathbb{W}_0 \subseteq \mathbb{W}$, and thus

⁹The precision here is quite subtle: A standard foundation for mathematics will not abstractly prove the existence of a ‘smallest class’ of channels subject to some requirements, since it is generally too large to be a set. As such, Theorem 4.1.23 is really a so-called *theorem schema*, one theorem for each possible class \mathbf{C} .

$$(\mathcal{C} \parallel \mathcal{D})(\mathbb{A}) = (\mathcal{C} \parallel \mathcal{D})(Y_0 \cup W_0) = \mathcal{C}(Y_0) \cup \mathcal{D}(W_0). \quad (4.22)$$

Now, because (T, \mathcal{C}) and (S, \mathcal{D}) belong to \mathbf{C} we may write

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \text{---}\mathbb{Y}\text{---} \end{array} \boxed{(T, \mathcal{C})} = \begin{array}{c} \text{---}\mathcal{C}(Y_0)\text{---} \\ \text{---}\mathbb{X} \setminus \mathcal{C}(Y_0)\text{---} \end{array} \boxed{(T_1, \mathcal{C}_1)} \text{---}\mathbb{H}\text{---} \begin{array}{c} \text{---}Y_0\text{---} \\ \text{---}\mathbb{Y} \setminus Y_0\text{---} \end{array} \boxed{(T_2, \mathcal{C}_2)} \quad (4.23)$$

for constructible channels (T_1, \mathcal{C}_1) and (T_2, \mathcal{C}_2) , and

$$\begin{array}{c} \text{---}\mathbb{Z}\text{---} \\ \text{---}\mathbb{W}\text{---} \end{array} \boxed{(S, \mathcal{D})} = \begin{array}{c} \text{---}\mathcal{D}(W_0)\text{---} \\ \text{---}\mathbb{Z} \setminus \mathcal{D}(W_0)\text{---} \end{array} \boxed{(S_1, \mathcal{D}_1)} \text{---}\mathbb{K}\text{---} \begin{array}{c} \text{---}W_0\text{---} \\ \text{---}\mathbb{W} \setminus W_0\text{---} \end{array} \boxed{(S_2, \mathcal{D}_2)} \quad (4.24)$$

for constructible channels (S_1, \mathcal{D}_1) and (S_2, \mathcal{D}_2) . By parallely composing and merging the causal channels pairwise, we obtain the desired form of $(T, \mathcal{C}) \parallel (S, \mathcal{D})$.

The fact that \mathbf{C} is closed under serial composition is proved similarly, using the induction hypothesis for each constituent. The details are left as exercise. \square

We end this section by observing another factorisation result (valid for arbitrary causal channels) which is often useful, and by employing Lemma 4.1.24 to give an alternative analysis of constructible channels with local specifications.

Lemma 4.1.26. (*Extracting Trashes from a Causal Channel.*)
Every causal channel $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \text{---}\mathbb{Y}\text{---} \end{array} \boxed{(T, \mathcal{C})}$ can be written in the form

$$\begin{array}{c} \text{---}\mathbb{X}'\text{---} \\ \text{---}\mathbb{X} \setminus \mathbb{X}'\text{---} \end{array} \boxed{(T', \mathcal{C}')} \text{---}\mathbb{Y}\text{---}, \quad (4.25)$$

$$\text{---}\mathbb{X} \setminus \mathbb{X}'\text{---} \boxed{\text{tr}}$$

with $\mathbb{X}' \subseteq \mathbb{X}$ and $\mathcal{C}'(\text{ports}(\mathbb{Y})) = \text{ports}(\mathbb{X}')$.

Proof. Let \mathbb{X}' be the sub-interface of \mathbb{X} defined by $\text{ports}(\mathbb{X}') = \mathcal{C}(\text{ports}(\mathbb{Y}))$. The desired factorisation follows from the non-signalling conditions implied by \mathcal{C} (and \mathcal{C}' is given by $\mathcal{C}'(J) = \mathcal{C}(J)$ for $J \subseteq \text{ports}(\mathbb{Y})$). \square

Example 4.1.27. (Inconstructibility Revisited.)

Suppose that $\begin{array}{c} \text{---}\mathcal{X}_1\text{---} \\ \text{---}\mathcal{X}_2\text{---} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---}\mathcal{Y}_1\text{---} \\ \text{---}\mathcal{Y}_2\text{---} \end{array}$ is constructible, with \mathcal{C} the local specification given by $\mathcal{C}(y_1) = \{x_1\}$ and $\mathcal{C}(y_2) = \{x_2\}$. Then, by Lemma 4.1.24, we can write e.g.

$$\begin{array}{c} \text{---}\mathcal{X}_1\text{---} \\ \text{---}\mathcal{X}_2\text{---} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---}\mathcal{Y}_1\text{---} \\ \text{---}\mathcal{Y}_2\text{---} \end{array} = \begin{array}{c} \text{---}\mathcal{X}_1\text{---} \\ \text{---}\mathcal{X}_2\text{---} \end{array} \boxed{(T_1, \mathcal{C}_1)} \text{---}\mathbb{H}\text{---} \begin{array}{c} \text{---}\mathcal{Y}_1\text{---} \\ \text{---}\mathcal{Y}_2\text{---} \end{array} \boxed{(T_2, \mathcal{C}_2)} \quad (4.26)$$

for constructible channels (T_1, \mathcal{C}_1) and (T_2, \mathcal{C}_2) . Now, however, observe that there can be no port $h \in \text{ports}(\mathbb{H})$ for which $h \in \mathcal{C}_2(y_2)$ and $x_1 \in \mathcal{C}_1(h)$, since such a port would provide a causal link effectuating (by the definition of composition) that $x_1 \in \mathcal{C}(y_2)$. Consequently, every $h \in \text{ports}(\mathbb{H})$ either satisfies $h \notin \mathcal{C}_2(y_2)$ or $\mathcal{C}_1(h) = \emptyset$. By Lemma 4.1.26, we may assume, after possibly extracting a trash from (T_2, \mathcal{C}_2) , that there are only ports of the latter kind. Applying Lemma 4.1.24 to (T_1, \mathcal{C}_1) , we can thus write

$$\begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \text{---}\mathcal{Y}_1\text{---} \\
 \text{---}\mathbb{H}\text{---}
 \end{array}
 (T_1, \mathcal{C}_1)
 =
 =
 \begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \text{---}\mathcal{Y}_1\text{---} \\
 \text{---}\mathbb{H}\text{---}
 \end{array}
 (T'_1, \mathcal{C}'_1)
 \begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \text{---}\mathcal{Y}_1\text{---} \\
 \text{---}\mathbb{H}\text{---}
 \end{array}
 s
 \begin{array}{c}
 \text{---}\mathcal{G}\text{---} \\
 \text{---}\mathbb{H}\text{---}
 \end{array}
 , \quad (4.27)$$


which ultimately implies that (T, \mathcal{C}) is necessarily a Bell-channel. Hence, we confirm that for example the PR box considered as a causal channel (Example 4.1.10) is inconstructible in the theory **QIT**. \blacklozenge

Note, importantly, that *by the DiVincenzo Property* of **QIT**, the channel $\begin{array}{c} \text{---}\mathcal{X}_1\text{---} \\ \text{---}\mathcal{Y}_1\text{---} \\ \text{---}\mathcal{X}_2\text{---} \\ \text{---}\mathcal{Y}_2\text{---} \end{array} T$ which underlies the PR box is of the form $\begin{array}{c} \text{---}\mathcal{X}_1\text{---} \\ \text{---}\mathcal{Y}_1\text{---} \\ \text{---}\mathcal{X}_2\text{---} \\ \text{---}\mathcal{Y}_2\text{---} \end{array} T_1 \begin{array}{c} \text{---}\mathcal{Y}_1\text{---} \\ \text{---}\mathcal{Y}_2\text{---} \end{array} T_2$ for channels T_1 and T_2 ; as a causal channel, however, it is not of the form (4.26) for any causal channels (T_1, \mathcal{C}_1) and (T_2, \mathcal{C}_2) .

4.2 Notions of Contraction

The class of causal channels admits an operation additional to that of serial and parallel composition. That this should be the case is imminent when we look at stencil-representations of causal channels, but that it is the case is by no means obvious.

Consider our generic stencil, filled with causal channels:



As outlined in the previous section, the stencil G with its filling \mathfrak{F} defines a total causal channel, $(T, \mathcal{C}) := \mathfrak{F}[G]$. Now, if the system at y_4 matches the system at x_3 , it should intuitively be possible to *contract* those two wires, effectively feeding the output at y_4 to the input port x_3 . Similarly for y_3 into x_3 , or y_5 into x_2 . (On the other hand, an insertion of y_2 into x_2 is not a priori sensible, as the circuit would thereby acquire a cycle.)

Clearly, any of the total causal channels which would result from such a contracted diagram can be determined from G and \mathfrak{F} . But might it be determinable from (T, \mathcal{C}) alone?

If the answer to this question were ‘no’, we would in a sense have done a lousy job in stating the very definition of a causal channel; for the concept to be of use, it must reflect

all operational aspects of what it aims to model, and certainly contraction is such an aspect. However, it is not at all clear that the violent reduction of (\mathfrak{F}, G) to $\mathfrak{F}[G]$ should spare the life of the possibility to contract. In fact, there are not only different fillings on the same stencil which will reproduce a given causal channel (T, \mathcal{C}) , there might also be different stencils on which (T, \mathcal{C}) is representable. What we ask is ultimately that the contraction be independent of any details of the representation whatsoever.

In Section 4.2.A, we will see that if the underlying theory Θ has universal dilations (cf. Definition 2.4.1), we can introduce unambiguous contractions, definable from (T, \mathcal{C}) alone. This result is one of the technical highlights of the chapter, and one that consolidates and justifies the definition of a causal channel.

In Section 4.2.B, we will then discuss why and how one might wish to define abstract notions of contraction, also in theories where the operational narrative of connecting wires does not really make sense. (For example, *cancellation* in a monoid, i.e. the act of forming the relation $x \succeq y$ from the relation $x \star z \succeq y \star z$ can be seen as the result of contraction.) Such abstract notions of contraction might be seen as generalisation of the *traces* in symmetric monoidal categories as introduced in Ref. [JSV96].

4.2.A The Standard Notion of Contraction

Let \mathfrak{F} be a filling (with causal channels) of a stencil G . Let $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ be the resulting causal channel, i.e. the value $\mathfrak{F}[G]$. The above discussion alluded to the contraction of a single pair of ports, but we might as well contract several pairs of ports at once. It is notationally convenient that the ports in each pair we wish to contract is given the same name (their systems must match anyway) – for instance, in (4.28) we would require $y_4 = x_3$ for the contraction of y_4 with x_3 . Hence, what we will define is the notion of *contracting* (T, \mathcal{C}) *along* \mathbb{P} , where \mathbb{P} is an interface which is a sub-interface of both \mathbb{X} and \mathbb{Y} .

First, we will define what it means for a common sub-interface of \mathbb{X} and \mathbb{Y} to be *seemingly contractible*. This will be simply a matter of whether there is a stencil-representation of (T, \mathcal{C}) which renders \mathbb{P} contractible based on ancestry in the stencil, that is, which does not acquire cycles when a contraction of the wires is forced. Having done this, we will then show that in universal theories the seemingly contractible interfaces can actually be contracted in a well-defined manner.

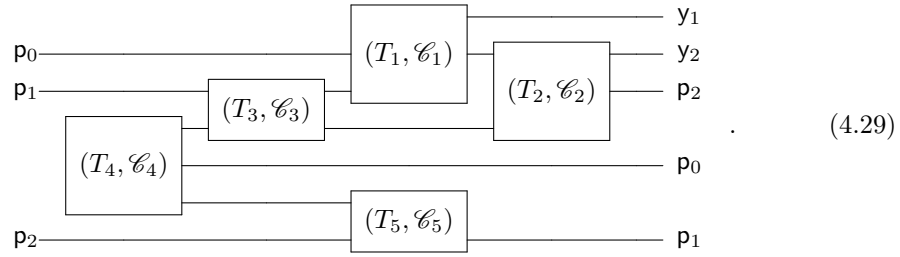
Given a stencil G and a bijection $h : P_{\text{out}} \rightarrow P_{\text{in}}$, where $P_{\text{in}} \subseteq \mathcal{W}_{\text{in}}(G)$ and $P_{\text{out}} \subseteq \mathcal{W}_{\text{out}}(G)$, let us say that G is *contractible w.r.t. h* if the directed graph that results from adjoining each wire $w \in P_{\text{out}}$ with the wire $h(w) \in P_{\text{in}}$ remains acyclic. If \mathfrak{F} is a filling of G and \mathbb{P} a common sub-interface of the input and output interfaces \mathbb{X} and \mathbb{Y} , there is a natural bijection $h_{\mathbb{P}}^{\mathfrak{F}}$ from the output wires w corresponding to $\mathbf{p} \in \text{ports}(\mathbb{P})$ (on the output side) to the input wires $w' = h_{\mathbb{P}}^{\mathfrak{F}}(w)$ corresponding to $\mathbf{p} \in \text{ports}(\mathbb{P})$ (on the input side).

Definition 4.2.1. (Seemingly Contractible Interfaces.)

Let $-\mathbb{X} \text{---} \boxed{(T, \mathcal{C})} \text{---} \mathbb{Y}$ be a causal channel, and let \mathbb{P} be a common sub-interface of \mathbb{X} and \mathbb{Y} . If (\mathfrak{F}, G) is a stencil-representation of (T, \mathcal{C}) , we say that \mathbb{P} is *contractible in (\mathfrak{F}, G)* if G is contractible w.r.t. $h_{\mathbb{P}}^{\mathfrak{F}}$. We say that \mathbb{P} is *seemingly contractible in (T, \mathcal{C})* , or that (T, \mathcal{C}) is *seemingly contractible along \mathbb{P}* , if there exists a stencil-representation (\mathfrak{F}, G) of (T, \mathcal{C}) in which \mathbb{P} is contractible. ■

Example 4.2.2. (Seeming Contractibility in a Generic Causal Channel.)

Suppose that (T, \mathcal{C}) admits the stencil-representation



In this representation, the sub-interfaces with port sets $\{p_0\}$, $\{p_1\}$, $\{p_2\}$, $\{p_0, p_1\}$ and $\{p_0, p_2\}$ are all contractible. Hence, those sub-interfaces are seemingly contractible in (T, \mathcal{C}) .

It might be that the sub-interface $\{p_1, p_2\}$ is also seemingly contractible in (T, \mathcal{C}) , but this is not witnessed by the above representation and would require the existence of another stencil-representation of (T, \mathcal{C}) , in which both p_1 and p_2 can be contracted without creating cycles. In such a representation, some of the first-mentioned sub-interfaces (e.g. $\{p_0, p_1\}$ or $\{p_0, p_2\}$) might cease to be contractible – in fact, it is unclear whether there exists a single representation of (T, \mathcal{C}) in which every seemingly contractible interface is contractible (and I do not know the answer to this question). \blacklozenge


Example 4.2.3. (Obvious Seeming Contractibility.)

The trivial interface \mathbb{I} is seemingly contractible in every causal channel (T, \mathcal{C}) . It is similarly clear that if \mathbb{P} is seemingly contractible, then so is any sub-interface $\mathbb{P}_0 \subseteq \mathbb{P}$. \blacklozenge

Example 4.2.4. (Seeming Contractibility versus the Causal Specification.)

Seeming contractibility can sometimes be excluded on the basis of the causal specification: If \mathbb{P} is seemingly contractible in (T, \mathcal{C}) , then necessarily $p \notin \mathcal{C}(p)$ for all $p \in \text{ports}(\mathbb{P})$. (Indeed, if $p \in \mathcal{C}(p)$ and $(T, \mathcal{C}) = \mathfrak{F}[G]$, then there must be a path in G leading from the input port p to the output port p ; thus a cycle would form if those two ports were contracted.) More generally, define on $\text{ports}(\mathbb{P})$ the transitive relation $<_{\mathcal{C}}^{\mathbb{P}}$ by letting $p <_{\mathcal{C}}^{\mathbb{P}} q$ if and only if there is a sequence $p_0, p_1, \dots, p_n \in \text{ports}(\mathbb{P})$, $n \geq 1$, such that $p_0 = p$, $p_n = q$ and $p_{k-1} \in \mathcal{C}(p_k)$ for all $k = 1, \dots, n$; if \mathbb{P} is seemingly contractible in (T, \mathcal{C}) , then $<_{\mathcal{C}}^{\mathbb{P}}$ must be irreflexive on $\text{ports}(\mathbb{P})$, i.e. there can be no $p \in \text{ports}(\mathbb{P})$ with $p <_{\mathcal{C}}^{\mathbb{P}} p$.

Since intuitively the relationship $p \not<_{\mathcal{C}}^{\mathbb{P}} p$ signifies that the input p is not in the causal past of the output p , one might naively think that an interface \mathbb{P} is seemingly contractible in (T, \mathcal{C}) if and only if $<_{\mathcal{C}}^{\mathbb{P}}$ is irreflexive on $\text{ports}(\mathbb{P})$. This fails, however, even for the simple interfaces: If a simple interface with port p is seemingly contractible, then, by merging components in a stencil-representation that witnesses this, we see that (T, \mathcal{C}) admits a stencil-representation of the specific form



Thus, for example, if we consider any of our inconstructible channels from Section 4.1.B, e.g. the PR box $\begin{array}{c} x_A - \{0, 1\} \\ x_B - \{0, 1\} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \{0, 1\} - y_A \\ \{0, 1\} - y_B \end{array}$ and if we pair the ports ‘oppositely’ (letting

$\mathfrak{p} := x_A = y_B$ and $\mathfrak{q} := x_B = y_A$), then we have e.g. $\mathfrak{p} \notin \mathcal{C}(\mathfrak{p})$, but the interface with port \mathfrak{p} is not seemingly contractible, since, as we saw in Example 4.1.10, the PR box admits no representation of the form (4.30).

Note, however, that if (T, \mathcal{C}) is constructible, then \mathcal{C} alone determines seeming contractibility of \mathbb{P} . Indeed, if $<_{\mathcal{C}}^{\mathbb{P}}$ is irreflexive on $\text{ports}(\mathbb{P})$, then \mathbb{P} must be contractible in any stencil-representation of (T, \mathcal{C}) whose filling consists of primitive causal channels, since there is in such a representation a path from x to y if and only if $x \in \mathcal{C}(y)$. ◆

Now, let us consider the problem of actually forming contractions of seemingly contractible interfaces, independently of the representations. A simple instance of this problem is the following: Suppose that

$$\begin{array}{c} \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T_1, \mathcal{C}_1) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T_2, \mathcal{C}_2) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{c} \mathfrak{p} \\ \mathfrak{p} \end{array} = \begin{array}{c} \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T'_1, \mathcal{C}'_1) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T'_2, \mathcal{C}'_2) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{c} \mathfrak{p} \\ \mathfrak{p} \end{array} \quad (4.31)$$

Is it necessarily the case that

$$\begin{array}{c} \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T_1, \mathcal{C}_1) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T_2, \mathcal{C}_2) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{c} \mathfrak{p} \\ \mathfrak{p} \end{array} = \begin{array}{c} \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T'_1, \mathcal{C}'_1) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{|c|} \hline (T'_2, \mathcal{C}'_2) \\ \hline \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \vdots \\ \text{---} \\ \text{---} \\ \vdots \end{array} \begin{array}{c} \mathfrak{p} \\ \mathfrak{p} \end{array} \quad ? \quad (4.32)$$

(Here, the wire-pairs corresponding to \mathfrak{p} have been contracted while all others remain as before.) In fact, we can ask an even simpler question: Disregarding the causal specifications altogether, is it even the case that the identity

$$\begin{array}{c} \text{---} \mathcal{X} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T_1 \\ \hline \end{array} \begin{array}{c} \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T_2 \\ \hline \end{array} = \begin{array}{c} \text{---} \mathcal{X} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z}' \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T'_1 \\ \hline \end{array} \begin{array}{c} \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z}' \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T'_2 \\ \hline \end{array} \quad (4.33)$$

between channels implies the identity

$$\begin{array}{c} \text{---} \mathcal{X} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T_1 \\ \hline \end{array} \begin{array}{c} \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T_2 \\ \hline \end{array} = \begin{array}{c} \text{---} \mathcal{X} \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z}' \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T'_1 \\ \hline \end{array} \begin{array}{c} \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Z}' \text{---} \\ \text{---} \mathcal{W} \text{---} \\ \text{---} \mathcal{Y} \text{---} \end{array} \begin{array}{|c|} \hline T'_2 \\ \hline \end{array} \quad ? \quad (4.34)$$

The latter identity asserts the equality of the serial compositions $T_2 \circ T_1$ and $T'_2 \circ T'_1$, whereas the former asserts the equality of a ‘partial’ serial composition, as if interrupted in the midst of composing. What we are asking is thus *whether it is possible to finish a half-hearted serial composition*.

It turns out that it is enough to solve that problem: If we can answer the latter question affirmatively, then we obtain independence of the stencil-representations all the way up to general contractions, essentially because any contraction can be decomposed into a sequence

of contractions of this kind. The most interesting part is thus the result that the interrupted serial compositions can be finished without knowing the individual components.

To show this, an old acquaintance comes to rescue:

Lemma 4.2.5. (*Interrupted Serial Compositions can be Completed in Universal Theories.*)

Suppose that Θ is a universal theory. If $\begin{array}{c} -\mathcal{X}- \\ \boxed{T_1} \\ -\mathcal{Z}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}$, $\begin{array}{c} -\mathcal{Z}- \\ \boxed{T_2} \\ -\mathcal{Y}- \end{array}$ and $\begin{array}{c} -\mathcal{X}- \\ \boxed{T'_1} \\ -\mathcal{Z}'- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}$, $\begin{array}{c} -\mathcal{Z}'- \\ \boxed{T'_2} \\ -\mathcal{Y}- \end{array}$ are pairs of channels such that

$$\begin{array}{c} -\mathcal{X}- \\ \boxed{T_1} \\ -\mathcal{Z}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{T'_1} \\ -\mathcal{Z}'- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}, \quad (4.35)$$

then

$$\begin{array}{c} -\mathcal{X}- \\ \boxed{T_1} \\ -\mathcal{Z}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{T'_1} \\ -\mathcal{Z}'- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}. \quad (4.36)$$

Proof. By trashing \mathcal{Y} in Eq. (4.35), we see that

$$\begin{array}{c} -\mathcal{X}- \\ \boxed{T_1} \\ -\mathcal{Z}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ \text{tr} \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{T'_1} \\ -\mathcal{Z}'- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ \text{tr} \end{array}. \quad (4.37)$$

Call this channel $\begin{array}{c} -\mathcal{X}- \\ \boxed{T} \\ -\mathcal{W}- \end{array}$, and let $\begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}$ be a universal dilation of T . Evidently, both T_1 and T'_1 are one-sided dilations of T , so, by completeness of U , we find G and G' with

$$\begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{T_1} \\ -\mathcal{Z}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} \quad \text{and} \quad \begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{T'_1} \\ -\mathcal{Z}'- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}. \quad (4.38)$$

Plugging this into Eq. (4.35) yields

$$\begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}, \quad (4.39)$$

so universality of U implies that $\begin{array}{c} \sim_{\mathcal{E}_0} \\ \boxed{G} \\ -\mathcal{W}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} \sim_{\mathcal{E}_0} \\ \boxed{G'} \\ -\mathcal{W}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}$. By Eq. (4.38), we must then have

$$\begin{aligned} \begin{array}{c} -\mathcal{X}- \\ \boxed{T_1} \\ -\mathcal{Z}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} &= \begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} \\ &= \begin{array}{c} -\mathcal{X}- \\ \boxed{U} \\ \sim_{\mathcal{E}_0} \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array} = \begin{array}{c} -\mathcal{X}- \\ \boxed{T'_1} \\ -\mathcal{Z}'- \end{array} \begin{array}{c} \mathcal{W}- \\ \\ -\mathcal{Y}- \end{array}, \end{aligned} \quad (4.40)$$

as desired. \square

Let us now argue that Lemma 4.2.5 implies that contractions are generally independent of stencil-representations.

The first thing to realise is that Lemma 4.2.5 entails the implication of Eq. (4.32) by Eq. (4.31). It follows clearly from the lemma that the channels on each side will be the same, so only the causal specifications need to be accounted for. However, we can think of causal specifications as *transformations* in a theory Θ' in which the *systems* are interfaces, with serial and parallel composition given as in Definition 4.1.19. This theory can be regarded a sub-theory of \mathbf{Sets}^* since causal specifications are ultimately functions between sets,¹⁰ and as the theory \mathbf{Sets}^* is cartesian Lemma 4.2.5 itself applies to show that an interrupted serial composition of causal specifications can be completed, which exactly yields the desired.

Next, we must realise that the implication of Eq. (4.32) by Eq. (4.31) is enough to show general independence of the stencil-representation. More precisely, given a stencil-representation (\mathfrak{F}, G) of $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ in which $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$ is contractible, let us denote by $(G)_{\mathbb{P}}$ the stencil that arises when contracting G along the wires corresponding \mathbb{P} ; the boxes in $(G)_{\mathbb{P}}$ are the same as those in G so we can consider \mathfrak{F} as a filling of $(G)_{\mathbb{P}}$, and the value $\mathfrak{F}[(G)_{\mathbb{P}}]$ is precisely the causal channel from $\mathbb{X} \setminus \mathbb{P}$ to $\mathbb{Y} \setminus \mathbb{P}$ that corresponds to contracting \mathbb{P} in the representation (\mathfrak{F}, G) . What we must show is that $\mathfrak{F}[(G)_{\mathbb{P}}] = \mathfrak{F}'[(G')_{\mathbb{P}}]$ whenever (\mathfrak{F}, G) and (\mathfrak{F}', G') are stencil-representations of (T, \mathcal{C}) in which \mathbb{P} is contractible. This can be done by an induction argument. Let $P(n)$, with $n \in \mathbb{N}_0$, be the statement that for every causal channel (T, \mathcal{C}) and every seemingly contractible interface \mathbb{P} of size $|\mathbb{P}| = n$, we have $\mathfrak{F}[(G)_{\mathbb{P}}] = \mathfrak{F}'[(G')_{\mathbb{P}}]$ whenever (\mathfrak{F}, G) and (\mathfrak{F}', G') are stencil-representations of (T, \mathcal{C}) in which \mathbb{P} is contractible. Then, $P(0)$ is true by definition of what a stencil-representation is. And if $P(n)$ is true, then $P(n+1)$ follows. Indeed, in any given stencil-representation the total contraction can be executed by first contracting a single port, say $\mathfrak{p} \in \mathbf{ports}(\mathbb{P})$, and then the remaining n ports. But by merging channels, the contraction of a single pair of ports always takes the form of going from (4.31) to (4.32), so for any two stencil-representations (\mathfrak{F}, G) and (\mathfrak{F}', G') , we have $\mathfrak{F}[(G)_{\mathfrak{p}}] = \mathfrak{F}'[(G')_{\mathfrak{p}}] =: (S, \mathcal{D})$ and the induction hypothesis then applies to the two stencil-representations $(\mathfrak{F}, (G)_{\mathfrak{p}})$ and $(\mathfrak{F}', (G')_{\mathfrak{p}})$ of (S, \mathcal{D}) and yields the desired.

All in all, we have proved the following:

Theorem 4.2.6. *(The Standard Notion of Contraction.)*

Let Θ be a universal theory. If $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ is a causal channel in Θ which is seemingly contractible along $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$, then $\mathfrak{F}[(G)_{\mathbb{P}}] = \mathfrak{F}'[(G')_{\mathbb{P}}]$ for all stencil-representations (\mathfrak{F}, G) and (\mathfrak{F}', G') of (T, \mathcal{C}) in which \mathbb{P} is contractible.

4.2.B General Notions of Contraction

In order to use contractions effectively, it is desirable to have a more abstract reformulation of the concept.

Ideally, we would like to posit the existence of *contraction maps* $\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}}$, which render certain causal channels $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ *contractible along* \mathbb{P} , and map them to causal channels $\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}}((T, \mathcal{C})) : \mathbb{X} \setminus \mathbb{P} \rightarrow \mathbb{Y} \setminus \mathbb{P}$, their *contractions along* \mathbb{P} . In the case of the standard contraction, ‘contractibility along \mathbb{P} ’ would mean simply seeming contractibility along \mathbb{P} , and the ‘contraction’ would be simply the contraction in any stencil-representation. Abstractly, this interpretation must be enforced by subjecting the maps $\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}}$ to certain

¹⁰Note that though $\mathbf{ports}(\mathbb{Z})$ may be empty, the set $\mathcal{P}(\mathbf{ports}(\mathbb{Z}))$ is always non-empty, and $\mathcal{P}(\mathbf{ports}(\mathbb{Z}_1 \cup \mathbb{Z}_2)) \cong \mathcal{P}(\mathbf{ports}(\mathbb{Z}_1)) \times \mathcal{P}(\mathbf{ports}(\mathbb{Z}_2))$ when \mathbb{Z}_1 and \mathbb{Z}_2 are parallelly composable (i.e. disjoint).

coherence conditions, regulating how contractibility of given causal channels affects that of others, and how their contractions relate to each other.

Such a reformulation is not only desirable on the account of swifter applicability, but also because the standard notion of contraction is somewhat fickle. Consider for example the bipartite causal channels $\begin{array}{c} -\mathcal{X}_A \\ \boxed{(T, \mathcal{C})} \\ -\mathcal{X}_B \end{array} \begin{array}{c} \mathcal{Y}_A \\ \\ \mathcal{Y}_B \end{array}$ in **QIT**, with the local causal specification given by $\mathcal{C}(y_i) = \{x_i\}$. Suppose moreover that the output interface labelled by A matches the input interface labelled by B, i.e. that $y_A = x_B$ and $\mathcal{Y}_A = \mathcal{X}_B$. As we saw in Example 4.2.4, this common sub-interface is seemingly contractible if and only if (T, \mathcal{C}) factors as

$\begin{array}{c} -\mathcal{X}_A \\ \boxed{(T_1, \mathcal{C}_1)} \\ -\mathcal{X}_B \end{array} \begin{array}{c} \mathcal{Y}_A \\ \\ \mathcal{Y}_B \end{array}$. However, as discussed earlier, some causal channels in **QIT** admit such a factorisation with (T_1, \mathcal{C}_1) and (T_2, \mathcal{C}_2) in **QIT**[∞], even though they have no such factorisation in **QIT**. In other words, by slightly enlarging the theory, some channels become constructible, and (in the present case) indeed contractible,¹¹ even though they were not covered by the standard notion of contraction in **QIT**.

Finally, abstraction is of general interest as it often allows us to see more clearly what are the important features of a given concept. As such, a rather mathematical example of notions of contraction will be *cancellations* in a thin theory whose corresponding monoid satisfies a cancellation law.

It turns out that a reformulation of the precise content of Theorem 4.2.6 in terms of contraction maps might not be within scope. The reason is that one of the properties that one would like an abstract notion of contraction to have (and which seems necessary to recover the physical contraction of wires, cf. the proof of Proposition 4.2.11 below) is ‘transitivity’ of contractibility, namely: If \mathbb{P}_1 is contractible in $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ and if \mathbb{P}_2 is contractible in the resulting contraction, $\mathbf{C}_{\mathbb{P}_1}^{\mathbb{X} \rightarrow \mathbb{Y}}((T, \mathcal{C}))$, then the total interface $\mathbb{P}_1 \cup \mathbb{P}_2$ is contractible in (T, \mathcal{C}) . That this property should hold for seeming contractibility is however not obvious, and in fact I do not know whether it is true or false;¹² it is stated as an open problem below.

Instead, what we shall do in providing an abstract reformulation of Theorem 4.2.6 is to restrict its content to constructible causal channels, since the transitivity property can be proven for those. More general notions of contraction might work for larger collections of causal channels:

Definition 4.2.7. (Schemes of Causal Channels.)

Let Θ be a theory. A *scheme of causal channels* in Θ is a class \mathbf{E} of causal channels in Θ which is closed under serial and parallel compositions, and which contains every constructible causal channel in Θ . ■

Physically, a scheme \mathbf{E} of causal channels reflects a vision of which causal channels in Θ we imagine occurring in the world. Mathematically, a scheme is just an aggregate for defining notions of contraction. The largest scheme is $\mathbf{E} = \text{ICC}(\Theta)$, the class of all causal channels, and the smallest scheme is $\mathbf{E} = \text{Cons}(\Theta)$, the class of constructible causal channels; the latter is the most important example.

We will often call a causal channel belonging to \mathbf{E} simply an *\mathbf{E} -channel*. For interfaces \mathbb{X}, \mathbb{Y} in Θ , let us denote by $\mathbf{E}(\mathbb{X}, \mathbb{Y})$ the class of \mathbf{E} -channels from \mathbb{X} to \mathbb{Y} .

¹¹Note that the contraction will belong to **QIT** even if T_1 and T_2 do not, since the systems \mathcal{X}_i and \mathcal{Y}_i are finite-dimensional.

¹²It seems related to the problem raised in Example 4.2.2 about whether there always exists a stencil-representation (\mathfrak{F}, G) of (T, \mathcal{C}) such that if \mathbb{P} is seemingly contractible in (T, \mathcal{C}) then \mathbb{P} is contractible in the specific representation (\mathfrak{F}, G) .

Definition 4.2.8. (Notions of Contraction.)

Let \mathbf{E} be a scheme of causal channels in Θ . A *notion of contraction in \mathbf{E}* is a collection \mathfrak{C} of partially defined *contraction maps*, indexed by interfaces \mathbb{X}, \mathbb{Y} and common sub-interfaces $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$,

$$\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}} : \text{Dom}(\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}}) \subseteq \mathbf{E}(\mathbb{X}, \mathbb{Y}) \rightarrow \mathbf{E}(\mathbb{X} \setminus \mathbb{P}, \mathbb{Y} \setminus \mathbb{P}), \quad (4.41)$$

for which the causal channels in the domain $\text{Dom}(\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}})$ are called *contractible along \mathbb{P}* , and which are subject to the following five conditions (abbreviating $\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}}$ by $\mathfrak{C}_{\mathbb{P}}$ when \mathbb{X} and \mathbb{Y} are implicit):

1. **Soundness.** If $\boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-}$ and $\boxed{-\mathbb{Y}- (S, \mathcal{D}) -\mathbb{Z}-}$ are parallelly composable \mathbf{E} -channels, then their parallel composition is contractible along \mathbb{Y} , and the contraction equals the serial composition,

$$\mathfrak{C}_{\mathbb{Y}} \left(\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{Y}- (S, \mathcal{D}) -\mathbb{Z}-} \end{array} \right) = \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}- (S, \mathcal{D}) -\mathbb{Z}-}. \quad (4.42)$$

2. **Coherence of Nested Contraction.** An \mathbf{E} -channel $\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}_1- (T, \mathcal{E}) -\mathbb{P}_1-} \\ \boxed{-\mathbb{P}_2- (T, \mathcal{E}) -\mathbb{P}_2-} \end{array}$ is contractible along $\mathbb{P}_1 \cup \mathbb{P}_2$ if and only if it is contractible along \mathbb{P}_1 and the resulting contraction is contractible along \mathbb{P}_2 . Moreover, in this case the successive contraction coincides with the total contraction,

$$\mathfrak{C}_{\mathbb{P}_2} \left(\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}_2- \mathfrak{C}_{\mathbb{P}_1}((T, \mathcal{E})) -\mathbb{P}_2-} \end{array} \right) = \mathfrak{C}_{\mathbb{P}_1 \cup \mathbb{P}_2} \left(\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}_1- (T, \mathcal{E}) -\mathbb{P}_1-} \\ \boxed{-\mathbb{P}_2- (T, \mathcal{E}) -\mathbb{P}_2-} \end{array} \right). \quad (4.43)$$

3. **Freeness of Non-Contracted Interfaces.** If $\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}- (T, \mathcal{E}) -\mathbb{P}-} \end{array}$ is contractible along \mathbb{P} , then for any \mathbf{E} -channels $\boxed{-\mathbb{X}'- (S_1, \mathcal{D}_1) -\mathbb{X}-}$ and $\boxed{-\mathbb{Y}- (S_2, \mathcal{D}_2) -\mathbb{Y}'-}$ the causal channel $\begin{array}{c} \boxed{-\mathbb{X}'- (S_1, \mathcal{D}_1) -\mathbb{X}-} \\ \boxed{-\mathbb{P}- (T, \mathcal{E}) -\mathbb{P}-} \\ \boxed{-\mathbb{Y}- (S_2, \mathcal{D}_2) -\mathbb{Y}'-} \end{array}$ is contractible along \mathbb{P} , and contraction commutes with the processing of the non-contracted interfaces,

$$\mathfrak{C}_{\mathbb{P}} \left(\begin{array}{c} \boxed{-\mathbb{X}'- (S_1, \mathcal{D}_1) -\mathbb{X}-} \\ \boxed{-\mathbb{P}- (T, \mathcal{E}) -\mathbb{P}-} \\ \boxed{-\mathbb{Y}- (S_2, \mathcal{D}_2) -\mathbb{Y}'-} \end{array} \right) = \boxed{-\mathbb{X}'- (S_1, \mathcal{D}_1) -\mathbb{X}-} \mathfrak{C}_{\mathbb{P}}((T, \mathcal{E})) \boxed{-\mathbb{Y}- (S_2, \mathcal{D}_2) -\mathbb{Y}'-}. \quad (4.44)$$

4. **Freeness of Parallel Composition.** If $\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}- (T, \mathcal{E}) -\mathbb{P}-} \end{array}$ and $\boxed{-\mathbb{Z}- (S, \mathcal{D}) -\mathbb{W}-}$ are parallelly composable \mathbf{E} -channels, then $\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}- (T, \mathcal{E}) -\mathbb{P}-} \end{array}$ is contractible along \mathbb{P} if and only if $\begin{array}{c} \boxed{-\mathbb{X}- (T, \mathcal{E}) -\mathbb{Y}-} \\ \boxed{-\mathbb{P}- (T, \mathcal{E}) -\mathbb{P}-} \end{array}$ is contractible along \mathbb{P} , and in that case contraction commutes with the parallel composition,

$$\mathfrak{C}_{\mathbb{P}} \left(\begin{array}{c} \begin{array}{ccc} \text{---}\mathbb{X}\text{---} & \boxed{(T, \mathcal{C})} & \text{---}\mathbb{Y}\text{---} \\ \text{---}\mathbb{P}\text{---} & & \text{---}\mathbb{P}\text{---} \end{array} \\ \begin{array}{ccc} & & \\ & & \\ \text{---}\mathbb{Z}\text{---} & \boxed{(S, \mathcal{D})} & \text{---}\mathbb{W}\text{---} \end{array} \end{array} \right) = \begin{array}{ccc} \text{---}\mathbb{X}\text{---} & \boxed{\mathfrak{C}_{\mathbb{P}}((T, \mathcal{C}))} & \text{---}\mathbb{Y}\text{---} \\ & & \\ \text{---}\mathbb{Z}\text{---} & \boxed{(S, \mathcal{D})} & \text{---}\mathbb{W}\text{---} \end{array}. \quad (4.45)$$

5. **Well-Foundedness.** If $\text{---}\mathbb{X}\text{---}\boxed{(T, \mathcal{C})}\text{---}\mathbb{Y}\text{---}$ is contractible along $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$, then $\mathbb{p} \notin \mathcal{C}(\mathbb{p})$ for all $\mathbb{p} \in \text{ports}(\mathbb{P})$. ■

To digest this definition, several observations are in order.

Remark 4.2.9. (On the Significance of Well-Foundedness.)

In comparison to the remaining conditions, Well-Foundedness might seem a strange condition to highlight. However, we will use it at a crucial point in the proof of Theorem 4.3.13. ✦

Remark 4.2.10. (On the Statement and Significance of Soundness.)

The conditions have not been stated in their weakest form under their mutual presence. For example, we could have restricted the statement of the Soundness condition to the case where both (T, \mathcal{C}) and (S, \mathcal{D}) are identity channels between simple interfaces. Since a general identity channel $(\text{id}_{\mathbb{Y}}, \mathcal{I}_{\mathbb{Y}})$ can be realised as a parallel composition of identity channels between simple interfaces, Freeness of Parallel Composition along with Coherence of Nested Contraction would then extend the scope to this case, and by Freeness of Non-Contracted Interfaces the case of arbitrary (T, \mathcal{C}) and (S, \mathcal{D}) would follow.

In fact, using a suitable combination of conditions it follows even more generally that a channel of the form $\begin{array}{ccc} \text{---}\mathbb{X}\text{---} & \boxed{(T_1, \mathcal{C}_1)} & \text{---}\mathbb{P}\text{---} \\ & \text{---}\mathbb{H}\text{---} & \\ & \text{---}\mathbb{P}\text{---} & \boxed{(T_2, \mathcal{C}_2)} \text{---}\mathbb{Y}\text{---} \end{array}$ is contractible along \mathbb{P} with contraction given by $\begin{array}{ccc} \text{---}\mathbb{X}\text{---} & \boxed{(T_1, \mathcal{C}_1)} & \text{---}\mathbb{Y}\text{---} \\ & \text{---}\mathbb{H}\text{---} & \\ & \text{---}\mathbb{P}\text{---} & \boxed{(T_2, \mathcal{C}_2)} \end{array}$, and this requirement would thus have been yet another equivalent way of stating the Soundness condition. ✦

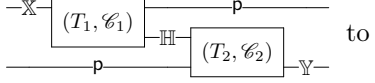
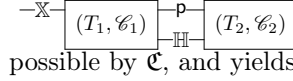
Whatever its guise, the point of the Soundness condition is to anchor any notion of contraction in terms of actually contracting wires in stencil-representations, as expressed by the following:

Proposition 4.2.11. (*Abstract Notions of Contractions Generalise the Standard Notion.*)
 Let \mathfrak{C} be a notion of contraction in \mathbf{E} . Suppose that $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ is a causal channel which has a stencil-representation (\mathfrak{F}, G) using a filling with \mathbf{E} -channels, and that $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$ is contractible in the representation (\mathfrak{F}, G) in the sense of Definition 4.2.1. Then, (T, \mathcal{C}) is contractible along \mathbb{P} according to \mathfrak{C} (i.e. $(T, \mathcal{C}) \in \text{Dom}(\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}})$) and the contraction $\mathfrak{C}_{\mathbb{P}}((T, \mathcal{C}))$ is the causal channel that arises from contracting \mathbb{P} in the stencil-representation (\mathfrak{F}, G) .

Proof. The statement is proved by induction on $n = |\text{ports}(\mathbb{P})|$.

For $n = 0$, i.e. $\mathbb{P} = \mathbb{I}$, what we must show is simply that \mathbb{I} is contractible in (T, \mathcal{C}) according to \mathfrak{C} , with $\mathfrak{C}_{\mathbb{I}}((T, \mathcal{C})) = (T, \mathcal{C})$. But by Soundness, the causal channel $\text{id}_{\mathbb{I}}$ is contractible along \mathbb{I} with contraction $\text{id}_{\mathbb{I}}$, and by Freeness of Parallel Composition the desired then follows.

As for the induction step, suppose the statement is always true when $|\text{ports}(\mathbb{P})| = n$. If $|\text{ports}(\mathbb{P}')| = n + 1$, then by the induction hypothesis and the ‘if’-direction in Coherence of

Nested Contraction it suffices to show that a single port $p \in \text{ports}(\mathbb{P}')$ is contractible according to \mathfrak{C} , and that the contraction is given by contracting p in any stencil-representation. In other words, the induction step is really the case $n = 1$. But in the case $n = 1$, any contraction in a stencil-representation takes the form of going from  to , and by Remark 4.2.10 such a contraction is indeed rendered possible by \mathfrak{C} , and yields the correct contraction. □

An immediate corollary of Proposition 4.2.11 is that contractions of seemingly contractible interfaces are independent of the chosen representation. However, as mentioned above, it is not clear that this statement can replace Theorem 4.2.6. The problem is essentially that ‘the standard notion of contraction’ might not literally be a notion of contraction in the sense of Definition 4.2.8. More precisely, if in a universal theory Θ we take $\text{Dom}(\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}})$ to be the class of all causal channels $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ in which \mathbb{P} is seemingly contractible, and if we define for that class the contraction $\mathfrak{C}_{\mathbb{P}}((T, \mathcal{C})) : \mathbb{X} \setminus \mathbb{P} \rightarrow \mathbb{Y} \setminus \mathbb{P}$ as the contraction in any valid representation, then, whereas the conditions 1., 3., 4., 5. and the ‘only if’-direction in 2. can be rather easily proved, I do not know a proof in condition 2. of the ‘if’-direction:

Open Problem 4.2.12. (*Existence of Notions of Contraction.*)

Suppose that Θ is universal. Does there exist a notion of contraction in $\mathbf{E} = \text{ICC}(\Theta)$, the scheme of all causal channels?

What we can obtain is a notion of contraction for constructible channels, and with this we will content ourselves:

Theorem 4.2.13. (*The Standard Notion of Contraction in $\text{Cons}(\Theta)$.*)

Suppose that Θ is a universal theory. Then there exists a notion of contraction in the scheme $\text{Cons}(\Theta)$ of constructible causal channels in Θ .

Proof. We simply define $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ to be an element of the domain $\text{Dom}(\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}})$ (i.e. ‘contractible along \mathbb{P} ’) precisely if (T, \mathcal{C}) is seemingly contractible along \mathbb{P} in the sense of Definition 4.2.1. For $(T, \mathcal{C}) \in \text{Dom}(\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}})$, we define the image $\mathfrak{C}_{\mathbb{P}}^{\mathbb{X} \rightarrow \mathbb{Y}}((T, \mathcal{C}))$ as the contraction in any stencil-representation; by Theorem 4.2.6, this is independent of the chosen representation. We have already seen in Example 4.2.4 that 5. (Well-Foundedness) holds, and 1. (Soundness) and 3. (Freeness of Non-Contracted Interfaces) are elementary to verify. The ‘only if’-directions in 2. (Coherence of Nested Contraction) and 4. (Freeness of Parallel Composition) are clear, and it is elementary that in this case the relevant contractions coincide. The ‘if’-direction in 4. follows by observing that a stencil-representation of the parallel composition $(T, \mathcal{C}) \parallel (S, \mathcal{D})$ which witnesses seeming contractibility of \mathbb{P} factors to witness seeming contractibility of (T, \mathcal{C}) .

Thus, only the ‘if’-direction in 2. remains. So far, we have not used a single time that the considered channels are constructible, but we need that now.

Suppose that \mathbb{P}_1 is seemingly contractible in (T, \mathcal{C}) , and that \mathbb{P}_2 is seemingly contractible in $\mathfrak{C}_{\mathbb{P}_1}((T, \mathcal{C}))$. We must show that $\mathbb{P}_1 \cup \mathbb{P}_2$ is seemingly contractible in (T, \mathcal{C}) . To this end, observe that by constructibility we can pick a stencil-representation (\mathfrak{F}, G) of (T, \mathcal{C}) whose filling has only primitive causal channels. By the observation in Example 4.2.4, such a representation correctly decides contractibility of sub-interfaces, so \mathbb{P}_1 is contractible in (\mathfrak{F}, G) and contracting \mathbb{P}_1 yields the representation $(\mathfrak{F}, (G)_{\mathbb{P}_1})$ of $\mathfrak{C}_{\mathbb{P}_1}((T, \mathcal{C}))$. This representation also has only primitive causal channels; thus, again, it correctly decides contractibility of

sub-interfaces. In particular, since \mathbb{P}_2 is seemingly contractible in $\mathfrak{C}_{\mathbb{P}_1}((T, \mathcal{C}))$, it must be contractible in the representation $(\mathfrak{F}, (G)_{\mathbb{P}_1})$; but by definition of what this means – namely, that $(G)_{\mathbb{P}_1}$ does not acquire cycles when contracting \mathbb{P}_2 – we see that all of $\mathbb{P}_1 \cup \mathbb{P}_2$ can be contracted in (\mathfrak{F}, G) without creating cycles. Consequently, $\mathbb{P}_1 \cup \mathbb{P}_2$ is seemingly contractible in (T, \mathcal{C}) , as desired. \square

In general, it is natural to sum up a notion of contraction along with its scheme and underlying theory:

Definition 4.2.14. (Contractible Theories.)

A *contractible theory* is a triple $(\Theta, \mathbf{E}, \mathfrak{C})$, where Θ is a theory, \mathbf{E} is a scheme of causal channels in Θ , and \mathfrak{C} is a notion of contraction in \mathbf{E} . \blacksquare

A point which is both conceptually and technically simplifying is the following:

Remark 4.2.15. (Serial Composition as a Derived Notion.)

If $(\Theta, \mathbf{E}, \mathfrak{C})$ is a contractible theory, then the parallel composition in Θ together with \mathfrak{C} actually *defines* the serial composition in Θ , by the Soundness condition for \mathfrak{C} . As such, we can think of serial composition as a derived notion, and we shall often do so in the following. \clubsuit

It is important to stress that even though schemes were introduced above mainly to define notions of contraction, they are relevant even if the open problem 4.2.12 has an affirmative answer. Indeed, if \mathfrak{C} is a notion of contraction in all of $\text{ICC}(\Theta)$, it makes sense to talk about schemes \mathbf{E} which are closed under the contraction \mathfrak{C} as models of *those causal channels which are physically implementable*. This will be significant in particular when we introduce the causal-dilational ordering in Section 4.3.

Let me conclude the subsection by mentioning a relation between abstract notions of contraction and so-called *traces* in symmetric monoidal categories:

Remark 4.2.16. (Relation to Traces in Symmetric Monoidal Categories.)

Generalising the ordinary concept of trace in the symmetric monoidal category $(\mathbf{Vect}_k, \otimes, k)$, Ref. [JSV96] introduced the concept of (*abstract*) *traces* in general (symmetric) monoidal categories, operations which map transformations $T : \mathcal{X} \parallel \mathcal{Z} \rightarrow \mathcal{Y} \parallel \mathcal{Z}$ in the category to transformations $T' : \mathcal{X} \rightarrow \mathcal{Y}$ (thus ‘swallowing’ the object \mathcal{Z}), subjected to certain coherence conditions. Though I was not aware of this work when first developing the theory presented here, there are extremely close connections between abstract traces and abstract notions of contraction. The connection is probably most easily seen by means of the simplification presented in Ref. [Has09] of the original conditions for traces.

Specifically, the *Yanking* condition of Ref. [Has09] corresponds to the Soundness condition in Definition 4.2.8, whereas the *Tightening* and *Superposing* conditions correspond to the Freeness of Non-Contracted Interfaces and Freeness of Parallel Composition, respectively. The *Vanishing* condition corresponds to Coherence under Nested Contraction. The Well-Foundedness condition has no equivalent, since causality is absent in the framework of traced categories. Likewise, any system is considered contractible, so the statements in Definition 4.2.8 about contractibility have no equivalents in that framework either (as we will see in Section 4.3, those conditions have been carefully chosen). In fact, a trace in a symmetric monoidal category can be seen as a notion of contraction for which causality (and the condition of Well-Foundedness) is completely disregarded and for which every channel $T : \mathbb{X} \rightarrow \mathbb{Y}$ is contractible along any common sub-interface $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$. \clubsuit

4.3 Causal Dilations

Now that we have introduced the concept of causal channels and notions of contraction, we are ready to introduce a causal version of the theory of Chapter 2. Throughout, $(\Theta, \mathbf{E}, \mathfrak{C})$ will denote a fixed contractible theory. The narrative is that \mathbf{E} defines the collection of causal channels which we imagine to be physically possible (the reader is free to think of the example $\mathbf{E} = \text{Cons}(\Theta)$, the constructible channels, since this will anyway be the only example we ultimately care about). Whenever we speak of ‘contractibility’ we will mean contractibility according to \mathfrak{C} , and implicitly assume that the involved channels are \mathbf{E} -channels.¹³

In Section 4.3.A, we define the concept of *causal dilations*, consider a lot of examples, and prove three general stability results about causal dilations, most significantly Theorem 4.3.13 which asserts that causal dilations are stable under contractions in the environment.

Then, in Section 4.3.B, we introduce the causal version of the dilational ordering; this *causal-dilational ordering* will depend on the scheme \mathbf{E} , since it reflects that one dilation can be derived from another using \mathbf{E} -channels in the environment. We then show some basic *composability* results, regarding how the causal-dilational ordering interplays with various compositions.

4.3.A Causal Dilations – Definition and Basic Properties

The concept of a causal dilation can in principle be very compactly defined: It is simply the notion of dilation obtained by replacing channels by causal channels.

More explicitly, recall that the trash channel $\text{tr}_{\mathbb{Z}}$ can be equipped with a unique causal specification, and that when we write $\text{id}_{\mathbb{Z}}$ as a causal channel, we mean really $(\text{id}_{\mathbb{Z}}, \mathcal{I}_{\mathbb{Z}})$, where $\mathcal{I}_{\mathbb{Z}}$ is the specification given by $\mathcal{I}_{\mathbb{Z}}(\mathbb{J}) = \mathbb{J}$.

Definition 4.3.1. (Causal Dilations.)

Let $(T, \mathcal{C}) : \mathbb{X} \rightarrow \mathbb{Y}$ be a causal channel in Θ . A (*causal*) *dilation* of (T, \mathcal{C}) is a causal channel $(L, \mathcal{E}) : \mathbb{X} \amalg \mathbb{D} \rightarrow \mathbb{Y} \amalg \mathbb{E}$ such that $(\text{id}_{\mathbb{Y}} \amalg \text{tr}_{\mathbb{E}}) \circ (L, \mathcal{E}) = (T, \mathcal{C}) \amalg \text{tr}_{\mathbb{D}}$, with composition in the sense of Definition 4.1.19. In pictures,

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{\text{tr}} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(T, \mathcal{C})} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{\text{tr}} \end{array} . \quad (4.46)$$

■

Remark 4.3.2. (Terminology.)

We will often simply say that (L, \mathcal{E}) is a *dilation* of (T, \mathcal{C}) (rather than a *causal dilation*), since by including the causal specifications there is no risk of confusion. ✖

It is instructive to write out explicitly what the condition (4.46) says about the channels and specifications individually. It is simply the following:

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{L} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{\text{tr}} \end{array} = \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{T} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{\text{tr}} \end{array} \quad \text{and} \quad \mathcal{E}|_{\text{ports}(\mathbb{Y})} = \mathcal{C} ; \quad (4.47)$$

¹³As such, it would often be really nice if we could take $\mathbf{E} = \text{ICC}(\Theta)$, i.e. if there were an affirmative answer to the open problem 4.2.12.

in other words, L is a dilation of T and $\mathcal{E}(J) = \mathcal{C}(J)$ for all $J \subseteq \text{ports}(\mathbb{Y})$. The requirement on the specifications expresses that in the causal channel (L, \mathcal{E}) , every set of ports in the accessible output interface \mathbb{Y} has the precise same causes as they do in the causal channel (T, \mathcal{C}) . As such, the specification \mathcal{E} is completely determined by its restriction to $\text{ports}(\mathbb{E})$, i.e. by the cause sets $\mathcal{E}(J) \subseteq \text{ports}(\mathbb{X} \cup \mathbb{D})$ for $J \subseteq \text{ports}(\mathbb{E})$. In particular, though there may exist intricate causal relationships between the outputs in \mathbb{E} and the inputs in $\mathbb{X} \cup \mathbb{D}$, no port in \mathbb{Y} can have causes in \mathbb{D} .¹⁴

Example 4.3.3. (Primitive Dilations.)

If $\text{---}\mathbb{X}\text{---}\boxed{(T, \mathcal{C})}\text{---}\mathbb{Y}\text{---}$ is any causal channel and L is a dilation of T as a channel, then the causal channel $\begin{array}{c} \sim\mathbb{D}\sim \\ \text{---}\mathbb{X}\text{---} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ whose specification \mathcal{E} is determined by $\mathcal{E}(\mathbb{E}_0) = \mathbb{X} \cup \mathbb{D}$ for any non-trivial sub-interface $\mathbb{E}_0 \subseteq \mathbb{E}$, is a causal dilation of (T, \mathcal{C}) . Intuitively, it corresponds to thinking of the hidden outputs as requiring all inputs in $\mathbb{X} \cup \mathbb{D}$ to be fed. We will call such a dilation a *primitive causal dilation*, as it is the natural extension of primitive specifications in the sense of Example 4.3.3 to the realm of dilations. (Note, however, that (L, \mathcal{E}) itself need not be a primitive causal channel; it is primitive if and only if (T, \mathcal{C}) is.) \blacklozenge

Example 4.3.4. (All Causal Dilations of a Trash.)

Consider the channel $\text{---}\mathbb{X}\text{---}\boxed{\text{tr}}\text{---}\mathbb{I}\text{---}$, equipped with its unique causal specification. We can characterise all its causal dilations. Evidently, they are simply the causal channels $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{E}\sim \\ \text{---}\mathbb{E}\text{---} \end{array}$, with \mathbb{D} and \mathbb{E} arbitrary. These can be written in the form

$$\text{---}\mathbb{X}\text{---} \boxed{(\text{id}_{\mathbb{X}}, \mathcal{I}_{\mathbb{X}})} \begin{array}{c} \sim\mathbb{X}\sim \\ \sim\mathbb{D}\sim \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{E}\sim \\ \text{---}\mathbb{E}\text{---} \end{array}, \quad (4.48)$$

and though this observation is trivial, it suggests that the particular causal dilation $\text{---}\mathbb{X}\text{---} \boxed{(\text{id}_{\mathbb{X}}, \mathcal{I}_{\mathbb{X}})} \begin{array}{c} \sim\mathbb{X}\sim \\ \sim\mathbb{D}\sim \end{array}$ should be rendered a *complete causal dilation*, in an extension of the completeness notion to causal dilations.

The reader is encouraged to consider similarly what are all the causal dilations of $\text{---}\mathbb{X}\text{---}\boxed{\text{tr}}\text{---}\mathbb{I}\text{---}$ if we replace the trivial output interface \mathbb{I} by a non-trivial interface in which some of the ports act as ‘indicators’ of its use, as in Example 4.1.9. \blacklozenge

Example 4.3.5. (All Causal Dilations of a State.)

Consider a state $\boxed{s}\text{---}\mathbb{Y}\text{---}$, equipped with its unique causal specification. If $\boxed{s}\text{---}\mathbb{Y}\text{---}$ as a channel in Θ has a complete one-sided dilation $\begin{array}{c} \text{---}\mathbb{Y}\text{---} \\ \boxed{c} \\ \sim\mathbb{E}_0\sim \end{array}$, then every dilation (disregarding

causality) is of the form $\begin{array}{c} \text{---}\mathbb{Y}\text{---} \\ \boxed{c} \\ \sim\mathbb{D}\sim \end{array} \boxed{G} \begin{array}{c} \sim\mathbb{E}\sim \\ \text{---}\mathbb{E}\text{---} \end{array}$ for some channel G . Now, if (L, \mathcal{E}) is a causal dilation of s then necessarily $\mathcal{E}(\mathbb{Y}) = \mathbb{I}$ (i.e. no port in \mathbb{Y} has any causes whatsoever), so it is easy to see that (L, \mathcal{E}) must in fact be of the form $\begin{array}{c} \text{---}\mathbb{Y}\text{---} \\ \boxed{c} \\ \sim\mathbb{D}\sim \end{array} \boxed{(G, \mathcal{B})} \begin{array}{c} \sim\mathbb{E}\sim \\ \text{---}\mathbb{E}\text{---} \end{array}$ for some causal channel (G, \mathcal{B}) .

As such, we have a ‘complete’ causal dilation, like in the previous example. \blacklozenge

¹⁴One could certainly speculate whether this restriction is fair, but it turns out that the dilation concept which results from its absence is not only much more complicated, but also ill-behaved, cf. Example 4.3.12. The requirement might be physically motivated on the same ground that we require non-signalling from \mathbb{D} to \mathbb{Y} in the concept of a dilation: An implementer of a channel who wants to give the impression that we are interacting with (T, \mathcal{C}) through the accessible interface could not exploit functionality which would ruin this impression.

Example 4.3.6. (Causal-Dilational Purity?)


Recall that a channel is called dilationally pure if every dilation is obtained by parallel composition (Definition 2.2.9). One might be tempted to similarly call dilationally pure a causal channel $\boxed{(T, \mathcal{E})}$ with the property that every causal dilation is of the form $\boxed{(T, \mathcal{E})}$. This, however, is quickly realised to be a rather dull notion: As soon as $\mathbb{X} \neq \mathbb{I}$, there exists a dilation which extracts as side-information *the fact that some port in \mathbb{X} was fed with an input*, i.e. we have the dilation $\boxed{(T, \mathcal{E}')}$ with $\mathcal{E}'(1) = \text{ports}(\mathbb{X})$, and this dilation does not factor by parallel composition. (If \mathbb{X} has several ports, there are other non-trivial dilations as well, since we may have indicators for the different subsets of ports in \mathbb{X} .) As such, the ‘dilationally pure’ causal channels would be exactly those causal channels with $\mathbb{X} = \mathbb{I}$ and for which the underlying channel is a dilationally pure state. \blacklozenge

Example 4.3.7. (Causal Dilations of the Bit Refreshment.)

Consider the ‘bit refreshment’ from the general introduction to the thesis, i.e. the causal channel $\boxed{(T, \mathcal{E}_0)}$ in **CIT** with primitive specification \mathcal{E}_0 and with $T(b) = r$ for any input b , where $r \in \text{St}(\{0, 1\})$ is the uniformly random bit. Pictorially, we can represent (T, \mathcal{E}_0) as


$$\boxed{\text{tr}} \text{---} \boxed{r} \text{---} \text{---} \tag{4.49}$$

abbreviating the system $\{0, 1\}$ as ‘2’, and with the understanding that each component is given its primitive specification. One causal dilation of this channel is



each component with its primitive specification, one identity thus stalling the other. It corresponds to a scenario in which the agents controlling the environment draw a random bit, and upon our input to the open interface gives us a copy of that bit, while storing our input in their memory. Precisely, the specification \mathcal{E} of this dilation is such that the two upper output ports have the input as a cause, whereas the lower output port has no causes. This reflects the fact that the random bit can be drawn in advance of seeing our input.

Another causal dilation is given by



(again equipping each component with its primitive specification), corresponding to a scenario in which the agents draw a random bit and use it to decide whether or not to give us back as output our original input, or to flip it.

Now, disregarding causality altogether, both of the channels (4.50) and (4.51) are easily seen to be complete dilations of T , and as such they are equivalent in the dilational ordering of Section 2.2. *However*, what we shall do in a moment is to introduce a causal version of this dilational ordering, and it will not be possible to go from one causal dilation to the other in a way that respects the causality: As causal dilations, (4.50) and (4.51) are simply different, formalising the intuition that the side-information in one dilation (pre-existing knowledge of which bit will be given as output) is information about something entirely different than the side-information in the other (pre-existing knowledge of whether or not the input will be flipped). \blacklozenge

Example 4.3.8. (Acausal Side-Information.)

Let $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ be a causal dilation of $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(T, \mathcal{C})} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$. A sub-interface $\mathbb{E}_0 \subseteq \mathbb{E}$ is said to be *acausal* if $\mathcal{E}(\mathbb{E}_0) \subseteq \mathbb{D}$, or, equivalently, if $\mathcal{E}(\mathbb{E}_0) \cap \mathbb{X} = \mathbb{I}$. The output at an acausal interface \mathbb{E}_0 represents the information that can be available at the hidden interface before the open interface \mathbb{X} of the channel has been used. For instance, in Example 4.3.7 the copies in the environment of the random bits were acausal.

We will say a dilation (L, \mathcal{E}) *has no acausal side-information* if \mathbb{I} is the only acausal sub-interface of \mathbb{E} , i.e. if $\mathcal{E}(e) \cap \text{ports}(\mathbb{X}) \neq \emptyset$ for all $e \in \text{ports}(\mathbb{E})$. In such a dilation, no hidden outputs are available before the accessible interface has been used.

On the other hand, a dilation (L, \mathcal{E}) for which all of \mathbb{E} is acausal is called simply an *acausal dilation*. Every trivial dilation $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(T, \mathcal{C})} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ is acausal, but there may be other

acausal dilations. For example, if $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{T} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ can be represented as $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{p} \boxed{\tilde{T}} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ for some channel \tilde{T} and some state p (e.g. in **CIT** or **QIT**, if T admits a convex decomposition with weights described by p), then for any dilation $\begin{array}{c} \text{---}\mathbb{H}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{t} \end{array}$ of $\begin{array}{c} \text{---}\mathbb{H}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{p} \end{array}$ the causal channel

$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{t} \boxed{(\tilde{T}, \tilde{\mathcal{C}})} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$, with $\tilde{\mathcal{C}}$ given by $\tilde{\mathcal{C}}(\mathbb{Y}_0) = \mathcal{C}(\mathbb{Y}_0) \cup \mathbb{H}$ for $\mathbb{Y}_0 \subseteq \mathbb{Y}$ non-trivial, is an acausal dilation of (T, \mathcal{C}) , assuming that \tilde{T} is compatible with $\tilde{\mathcal{C}}$. In general (for example, if \tilde{T} and p constitute a non-trivial convex decomposition of T) this dilation does not factor as a trivial dilation.

As we will see in Chapter 5, quantum self-testing ensures that all acausal dilations of certain channels are trivial, and this has the interpretation that any randomness produced by those channels is *fresh*, independent of pre-existing randomness (Proposition 5.3.2). By the preceding considerations, it also implies that the channel at the open interface has only trivial convex decompositions, i.e. is extremal (Proposition 5.3.4). \blacklozenge

Example 4.3.9. (Causal Dilations of a Quantum Measurement?)

Consider in **QIT** the decoherence channel $\begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \sim\mathbb{C}^2\sim \\ \boxed{\Delta} \\ \sim\mathbb{C}^2\sim \\ \text{---}\mathbb{C}^2\text{---} \end{array}$ given by $\Delta(A) = |0\rangle\langle 0| A |0\rangle\langle 0| + |1\rangle\langle 1| A |1\rangle\langle 1|$ for $A \in \text{End}(\mathbb{C}^2)$. The channel Δ models a measurement in the computational basis ([NC02]). Let us equip Δ with the primitive causal specification \mathcal{C}_0 (this is the only compatible specification anyway). As in Example 4.3.3, we obtain of course for any dilation

$\begin{array}{c} \text{---}\mathbb{E}\text{---} \\ \sim\mathbb{C}^2\sim \\ \boxed{\Phi} \\ \sim\mathbb{C}^2\sim \\ \text{---}\mathbb{C}^2\text{---} \end{array}$ of Δ as a channel, a causal dilation by simply equipping Φ with its primitive specification.

There is, however, another dilation, namely one which is acausal: Based on the convex decomposition $\Delta = \frac{1}{2}\text{id}_{\mathbb{C}^2} + \frac{1}{2}Z$, where Z is the channel given by conjugation by the Pauli

z -unitary $\sigma_z := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$, we can write $-\mathbb{C}^2-\boxed{\Delta}-\mathbb{C}^2-$ as $\begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \boxed{\tau}\text{---}\mathbb{C}^2\text{---} \end{array} \boxed{\Gamma} \text{---}\mathbb{C}^2\text{---}$, where $\tau := \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|$ is the fully mixed state on \mathbb{C}^2 , and where Γ measures the lower system to determine whether to apply $\text{id}_{\mathbb{C}^2}$ or Z to the upper system. But then, letting $K : \mathbb{C}^2 \rightarrow \mathbb{C}^2 \otimes \mathbb{C}^2$ denote the embedding in **QIT** of the classical copy-channel, i.e. the quantum channel given by $K(A) = \langle 0|A|0\rangle|0\rangle\langle 0| \otimes |0\rangle\langle 0| + \langle 1|A|1\rangle|1\rangle\langle 1| \otimes |1\rangle\langle 1|$, the state $\begin{array}{c} \boxed{\tau} \\ \boxed{K} \end{array}$ is a dilation of τ (corresponding to a copy), and following Example 4.3.8 we thus obtain the causal dilation

$$\begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \boxed{\tau} \\ \boxed{K} \\ \boxed{\Gamma} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \quad (4.52)$$

of $-\mathbb{C}^2-\boxed{\Delta}-\mathbb{C}^2-$, where each component is equipped with its primitive specification. In equations, the channel (4.52) is given by $A \mapsto \frac{1}{2}A \otimes |0\rangle\langle 0| + \frac{1}{2}\sigma_z A \sigma_z^* \otimes |1\rangle\langle 1|$.

This dilation is strange in the context of measurements since it is acausal, and thus represents side-information available before the execution of the measurement; the dilation formalises the curious circumstance that a quantum measurement can – mathematically – be implemented by tossing a fair coin and using the outcome to decide whether to apply a Pauli z -conjugation or do nothing to the system at hand. There are other such strange causal dilations, indeed one for each convex decomposition of Δ (for example, we also have $\Delta = \frac{1}{2}X + \frac{1}{2}Y$, where X is Y are conjugations by the Pauli x - and y -unitaries, respectively.)

It would seem that there is no sensible way in which these strange dilations really reflect the nature of quantum measurements, and I do not know how to interpret them other than by refuting them as valid dilations when thinking of Δ as a measurement. As we will see in Chapter 5, the strangeness is not purely philosophical, but constitutes a mathematical nuisance when establishing the precise connection between quantum self-testing and the theory of causal dilations.

◆

We now show three stability results about causal dilations, two of which are simple to prove and the third of which is highly subtle, representing perhaps the most surprising result about causal dilations at all.

The first result is immediate from the definition, like it was in the causality-free setting:

Proposition 4.3.10. *(Stability of Causal Dilations under Parallel Composition.)*

Suppose that $\begin{array}{c} \sim\mathbb{D}_1\sim \\ \text{---}\mathbb{X}_1\text{---} \end{array} \boxed{(L_1, \mathcal{E}_1)} \begin{array}{c} \sim\mathbb{E}_1\sim \\ \text{---}\mathbb{Y}_1\text{---} \end{array}$ is a causal dilation of $\text{---}\mathbb{X}_1\text{---}\boxed{(T_1, \mathcal{C}_1)}\text{---}\mathbb{Y}_1\text{---}$ and $\begin{array}{c} \sim\mathbb{D}_2\sim \\ \text{---}\mathbb{X}_2\text{---} \end{array} \boxed{(L_2, \mathcal{E}_2)} \begin{array}{c} \sim\mathbb{E}_2\sim \\ \text{---}\mathbb{Y}_2\text{---} \end{array}$ is a causal dilation of $\text{---}\mathbb{X}_2\text{---}\boxed{(T_2, \mathcal{C}_2)}\text{---}\mathbb{Y}_2\text{---}$. Then,

$$\begin{array}{c} \text{---}\mathbb{X}_1\text{---} \\ \sim\mathbb{D}_1\sim \\ \boxed{(L_1, \mathcal{E}_1)} \\ \sim\mathbb{D}_2\sim \\ \text{---}\mathbb{X}_2\text{---} \end{array} \begin{array}{c} \text{---}\mathbb{Y}_1\text{---} \\ \sim\mathbb{E}_1\sim \\ \boxed{(L_2, \mathcal{E}_2)} \\ \sim\mathbb{E}_2\sim \\ \text{---}\mathbb{Y}_2\text{---} \end{array} \quad (4.53)$$

is a causal dilation of $\begin{array}{c} \text{---}\mathbb{X}_1\text{---} \\ \boxed{(T_1, \mathcal{C}_1)} \\ \text{---}\mathbb{Y}_1\text{---} \\ \text{---}\mathbb{X}_2\text{---} \\ \boxed{(T_2, \mathcal{C}_2)} \\ \text{---}\mathbb{Y}_2\text{---} \end{array}$.

The next obvious result to prove would be that also the serial composition of causal dilations is a causal dilation, and this is immediate from the definition too. However, as discussed earlier, serial composition can be derived from parallel composition and contraction, so given Proposition 4.3.10 it is stronger to state a stability result about contractions in the accessible interface:

Proposition 4.3.11. (*Stability of Causal Dilations under Accessible Contractions.*)

Suppose that $\begin{array}{c} \sim \mathbb{D} \sim \\ \text{---} \mathbb{X} \text{---} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim \mathbb{E} \sim \\ \text{---} \mathbb{Y} \text{---} \end{array}$ is a causal dilation of $\begin{array}{c} \text{---} \mathbb{X} \text{---} \\ \text{---} \mathbb{Y} \text{---} \end{array} \boxed{(T, \mathcal{C})}$, and that \mathbb{P} is a common sub-interface of \mathbb{X} and \mathbb{Y} . Then, \mathbb{P} is contractible in (T, \mathcal{C}) if and only if it is contractible in (L, \mathcal{E}) , and in that case the contraction

$$\begin{array}{c} \sim \mathbb{D} \sim \\ \text{---} \mathbb{X} \setminus \mathbb{P} \text{---} \end{array} \boxed{\mathfrak{C}_{\mathbb{P}}((L, \mathcal{E}))} \begin{array}{c} \sim \mathbb{E} \sim \\ \text{---} \mathbb{Y} \setminus \mathbb{P} \text{---} \end{array} \quad (4.54)$$

is a causal dilation of the contraction $\begin{array}{c} \text{---} \mathbb{X} \setminus \mathbb{P} \text{---} \\ \text{---} \mathbb{Y} \setminus \mathbb{P} \text{---} \end{array} \boxed{\mathfrak{C}_{\mathbb{P}}((T, \mathcal{C}))}$.

Proof. Consider the causal channel

$$\begin{array}{c} \sim \mathbb{E} \sim \\ \text{---} \mathbb{X} \text{---} \end{array} \boxed{\text{tr}} \begin{array}{c} \sim \mathbb{D} \sim \\ \text{---} \mathbb{X} \text{---} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim \mathbb{E} \sim \\ \text{---} \mathbb{Y} \text{---} \end{array} . \quad (4.55)$$

By Soundness it is contractible along \mathbb{E} , and its contraction along \mathbb{E} is given by

$$\begin{array}{c} \sim \mathbb{D} \sim \\ \text{---} \mathbb{X} \text{---} \end{array} \boxed{\text{tr}} \begin{array}{c} \text{---} \mathbb{X} \text{---} \\ \text{---} \mathbb{Y} \text{---} \end{array} \boxed{(T, \mathcal{C})} , \quad (4.56)$$

since (L, \mathcal{E}) dilates (T, \mathcal{C}) . Now, by Freeness of Parallel Composition, \mathbb{P} is contractible in (T, \mathcal{C}) if and only if it is contractible in the channel (4.56). But by Nested Contraction, \mathbb{P} is contractible in (4.56) if and only if it is contractible in (4.55). And this again is the case if and only if \mathbb{P} is contractible in (L, \mathcal{E}) . Moreover, in the affirmative case, Nested Contraction implies that the order of contraction of along \mathbb{P} and \mathbb{E} can be interchanged, so

$$\begin{array}{c} \sim \mathbb{E} \sim \\ \text{---} \mathbb{X} \setminus \mathbb{P} \text{---} \end{array} \boxed{\text{tr}} \begin{array}{c} \sim \mathbb{D} \sim \\ \text{---} \mathbb{X} \setminus \mathbb{P} \text{---} \end{array} \boxed{\mathfrak{C}_{\mathbb{P}}((L, \mathcal{E}))} \begin{array}{c} \sim \mathbb{E} \sim \\ \text{---} \mathbb{Y} \setminus \mathbb{P} \text{---} \end{array} \quad (4.57)$$

contracts along \mathbb{E} to yield

$$\begin{array}{c} \sim \mathbb{D} \sim \\ \text{---} \mathbb{X} \setminus \mathbb{P} \text{---} \end{array} \boxed{\text{tr}} \begin{array}{c} \text{---} \mathbb{X} \setminus \mathbb{P} \text{---} \\ \text{---} \mathbb{Y} \setminus \mathbb{P} \text{---} \end{array} \boxed{\mathfrak{C}_{\mathbb{P}}((T, \mathcal{C}))} , \quad (4.58)$$

which by Soundness is precisely to say that $\mathfrak{C}_{\mathbb{P}}((L, \mathcal{E}))$ is a dilation of $\mathfrak{C}_{\mathbb{P}}((T, \mathcal{C}))$. \square

Theorem 4.3.13. (*Hidden Contractions of Causal Dilations are Causal Dilations.*)

Let $\text{-}\mathbb{X}\text{-}\boxed{(T, \mathcal{E})}\text{-}\mathbb{Y}\text{-}$ be a causal channel. Suppose that $\begin{array}{c} \sim\mathbb{D} \\ \text{-}\mathbb{X} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{E} \\ \text{-}\mathbb{Y} \end{array}$ is a causal dilation of $\text{-}\mathbb{X}\text{-}\boxed{(T, \mathcal{E})}\text{-}\mathbb{Y}\text{-}$, and that $\mathbb{Q} \subseteq \mathbb{D} \cap \mathbb{E}$ is contractible. Then the contraction

$$\begin{array}{c} \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{\mathfrak{C}_{\mathbb{Q}}((L, \mathcal{E}))} \begin{array}{c} \sim\mathbb{E} \setminus \mathbb{Q} \\ \text{-}\mathbb{Y} \end{array} \quad (4.61)$$

is a causal dilation of $\text{-}\mathbb{X}\text{-}\boxed{(T, \mathcal{E})}\text{-}\mathbb{Y}\text{-}$.

Proof. By induction and Coherence of Nested Contraction, it suffices to prove the statement in the case where \mathbb{Q} is a simple interface, i.e. has only a single port q .

Under this assumption, consider the causal channel

$$\begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{E} \setminus \mathbb{Q} \text{ tr} \\ \text{-}\mathbb{Y} \end{array} . \quad (4.62)$$

By Freeness of Non-Contracted Interfaces, this causal channel is contractible along \mathbb{Q} , and its contraction is

$$\mathfrak{C}_{\mathbb{Q}} \left(\begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{E} \setminus \mathbb{Q} \text{ tr} \\ \text{-}\mathbb{Y} \end{array} \right) = \begin{array}{c} \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{\mathfrak{C}_{\mathbb{Q}}((L, \mathcal{E}))} \begin{array}{c} \sim\mathbb{E} \setminus \mathbb{Q} \text{ tr} \\ \text{-}\mathbb{Y} \end{array} . \quad (4.63)$$

Notice that the right hand side is precisely the thing to consider if we wish to show that $\mathfrak{C}_{\mathbb{Q}}((L, \mathcal{E}))$ is a dilation of (T, \mathcal{E}) . Hence, let us compute this contraction in a different way.

Observe the following about the causal specification of the channel (4.62): The input port $q \in \text{ports}(\mathbb{Q})$ is not a cause of \mathbb{Y} , since $q \notin \mathcal{E}(\mathbb{Y})$ by virtue of (L, \mathcal{E}) being a causal dilation of (T, \mathcal{E}) ; on the other hand, q is not a cause of the output port q either, since this would contradict Well-Foundedness by virtue of \mathbb{Q} being contractible. By additivity of causal specifications, we thus conclude that the input interface \mathbb{Q} is not a cause of $\mathbb{Y} \cup \mathbb{Q}$, i.e. not a cause of any outputs whatsoever. But this implies (by Lemma 4.1.26) that (4.62) must factor as

$$\begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{E} \setminus \mathbb{Q} \text{ tr} \\ \text{-}\mathbb{Y} \end{array} = \begin{array}{c} \sim\mathbb{Q} \text{ tr} \\ \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{(S, \mathcal{D})} \begin{array}{c} \sim\mathbb{Q} \\ \text{-}\mathbb{Y} \end{array} \quad (4.64)$$

for some causal channel (S, \mathcal{D}) . This identity allows us to determine the contraction (4.63) differently.

By Soundness we see from (4.64) that

$$\mathfrak{C}_{\mathbb{Q}} \left(\begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \sim\mathbb{Q} \\ \sim\mathbb{E} \setminus \mathbb{Q} \text{ tr} \\ \text{-}\mathbb{Y} \end{array} \right) = \begin{array}{c} \sim\mathbb{D} \setminus \mathbb{Q} \\ \text{-}\mathbb{X} \end{array} \boxed{(S, \mathcal{D})} \begin{array}{c} \sim\mathbb{Q} \text{ tr} \\ \text{-}\mathbb{Y} \end{array} . \quad (4.65)$$

We also see from (4.64) that

$$\begin{array}{c}
\begin{array}{c} \text{tr} \\ \text{---} \text{Q} \text{---} \\ \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{(S, \mathcal{D})} \begin{array}{c} \text{---} \text{Q} \text{---} \\ \text{---} \text{tr} \end{array} \\
= \begin{array}{c} \text{---} \text{Q} \text{---} \\ \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---} \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{E} \setminus \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{Y} \text{---} \end{array} \\
= \begin{array}{c} \text{---} \text{Q} \text{---} \\ \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---} \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{Y} \text{---} \end{array} , \\
\end{array} \tag{4.66}$$

as (L, \mathcal{E}) is a dilation of (T, \mathcal{C}) . Hence, $\begin{array}{c} \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{(S, \mathcal{D})} \begin{array}{c} \text{---} \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{Y} \text{---} \end{array} = \begin{array}{c} \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---} \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{Y} \text{---} \end{array}$ by normality of the theory. Combining this with Eq. (4.65) and Eq. (4.63), we finally conclude that

$$\begin{array}{c} \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{\mathfrak{C}_{\mathbb{Q}}((L, \mathcal{E}))} \begin{array}{c} \text{---} \text{E} \setminus \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{Y} \text{---} \end{array} = \begin{array}{c} \sim \mathbb{D} \setminus \text{Q} \text{---} \\ \text{---} \text{X} \text{---} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---} \text{Q} \text{---} \\ \text{---} \text{tr} \\ \text{---} \text{Y} \text{---} \end{array} , \tag{4.67}$$

which shows that $\mathfrak{C}_{\mathbb{Q}}((L, \mathcal{E}))$ is a dilation of (T, \mathcal{C}) , as desired. \square

4.3.B The Causal-Dilational Ordering

Now that we have seen a number of examples of causal dilations, we are ready to introduce the causal version of *derivability*, namely the dilational ordering of Chapter 2. As outlined in the introduction to the chapter, this order is meant to formalise the idea that some causal dilations can be ‘constructed in the environment’ from others. Since we now have contractions in the bag of possible ‘constructions’ which can take place in the environment, we can treat two-sided dilations properly, but must also slightly rethink what a sensible notion of derivability should be.

Given a dilation $\begin{array}{c} \text{---} \text{X} \text{---} \\ \sim \mathbb{D} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---} \text{Y} \text{---} \\ \sim \mathbb{E} \end{array}$ of $\begin{array}{c} \text{---} \text{X} \text{---} \\ \sim \mathbb{D} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---} \text{Y} \text{---} \\ \sim \mathbb{E} \end{array}$, we would like to render it equivalent to any dilation of the form $\begin{array}{c} \text{---} \text{X} \text{---} \\ \sim \mathbb{D} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---} \text{Y} \text{---} \\ \sim \mathbb{E} \end{array}$.¹⁵ We would also like to render a dilation $\begin{array}{c} \sim \mathbb{A} \text{---} \\ \sim \mathbb{B} \end{array} \boxed{(G, \mathcal{B})} \begin{array}{c} \text{---} \text{Y} \text{---} \\ \sim \mathbb{E} \end{array}$ (L', \mathcal{E}') derivable from a dilation (L, \mathcal{E}) , if (L', \mathcal{E}') results from the contraction of a sub-interface of the hidden interface in (L, \mathcal{E}) . These two ideas are combined in Definition 4.3.14 below.

As we want to speak of contraction, we need to restrict attention to \mathbf{E} -channels from now on. (As mentioned earlier, the reader is free to consider for concreteness the case where \mathbf{E} is the class of constructible causal channels in Θ .) Given an \mathbf{E} -channel (T, \mathcal{C}) , let us denote by $\text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$ the class of causal dilations of (T, \mathcal{C}) which are \mathbf{E} -channels.

¹⁵That the latter should be deemed derivable from the former is clear, since the latter dilation represents the independent execution of a causal channel (G, \mathcal{B}) in the environment. That the former should be derivable from the latter, however, is ultimately a choice which reflects a convention. For example, it is not obvious that $\begin{array}{c} \text{---} \text{X} \text{---} \\ \sim \mathbb{D} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---} \text{Y} \text{---} \\ \sim \mathbb{E} \end{array}$ can be ‘constructed’ from $\begin{array}{c} \text{---} \text{X} \text{---} \\ \sim \mathbb{D} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---} \text{Y} \text{---} \\ \sim \mathbb{E} \end{array}$; if the theory Θ has states, we can construct (L, \mathcal{E}) by inserting a state into \mathbb{A} , but if there are no states on \mathbb{A} there is no obvious ‘construction’ in the environment which yields (L, \mathcal{E}) .

Definition 4.3.14. (The Causal-Dilational Ordering.)

Let $\text{-X-} \boxed{(T, \mathcal{E})} \text{-Y-}$ be an \mathbf{E} -channel. We denote by $\succeq_{(T, \mathcal{E})}^{\mathbf{E}}$ the relation on $\text{CausDil}^{\mathbf{E}}((T, \mathcal{E}))$ given by

$$\begin{array}{c} \text{-X-} \\ \sim_{\mathbb{D}} \sim \\ \boxed{(L, \mathcal{E})} \\ \sim_{\mathbb{E}} \sim \\ \text{-Y-} \end{array} \succeq_{(T, \mathcal{E})}^{\mathbf{E}} \begin{array}{c} \text{-X-} \\ \sim_{\mathbb{D}'} \sim \\ \boxed{(L', \mathcal{E}')} \\ \sim_{\mathbb{E}'} \sim \\ \text{-Y-} \end{array} \quad (4.68)$$

if and only if there exists, possibly after renaming the ports in \mathbb{D} and \mathbb{E} , \mathbf{E} -channels $\sim_{\mathbb{A}} \sim \boxed{(G, \mathcal{B})} \sim_{\mathbb{B}} \sim$ and $\sim_{\mathbb{A}'} \sim \boxed{(G', \mathcal{B}')} \sim_{\mathbb{B}'} \sim$ such that

$$\mathfrak{C}_{\mathbb{Q}} \left(\begin{array}{c} \text{-X-} \\ \sim_{\mathbb{D}} \sim \\ \boxed{(L, \mathcal{E})} \\ \sim_{\mathbb{E}} \sim \\ \text{-Y-} \\ \sim_{\mathbb{A}} \sim \\ \boxed{(G, \mathcal{B})} \\ \sim_{\mathbb{B}} \sim \end{array} \right) = \begin{array}{c} \text{-X-} \\ \sim_{\mathbb{D}'} \sim \\ \boxed{(L', \mathcal{E}')} \\ \sim_{\mathbb{E}'} \sim \\ \text{-Y-} \\ \sim_{\mathbb{A}'} \sim \\ \boxed{(G', \mathcal{B}')} \\ \sim_{\mathbb{B}'} \sim \end{array} \quad (4.69)$$

for some contractible common sub-interface \mathbb{Q} of $\mathbb{D} \cup \mathbb{A}$ and $\mathbb{E} \cup \mathbb{B}$. We say in this case that (L', \mathcal{E}') is derivable from (L, \mathcal{E}) . ■

Before exemplifying derivability, some observations are in order:

Proposition 4.3.15. (Derivability is a Pre-Order.)

The derivability relation $\succeq_{(T, \mathcal{E})}^{\mathbf{E}}$ is a pre-order (i.e. is reflexive and transitive) on the class $\text{CausDil}^{\mathbf{E}}((T, \mathcal{E}))$.

Proof. The relation is reflexive by taking $(G, \mathcal{B}) = (G', \mathcal{B}')$ and $\mathbb{Q} = \mathbb{I}$ in the condition (4.69). As for transitivity, suppose that $(L, \mathcal{E}) \succeq (L', \mathcal{E}')$ and $(L', \mathcal{E}') \succeq (L'', \mathcal{E}'')$. If (G_1, \mathcal{B}_1) , works on the left-hand side of (4.69) to realise the relation $(L, \mathcal{E}) \succeq (L', \mathcal{E}')$, and if (G_2, \mathcal{B}_2) works on the left-hand side of (4.69) to realise the relation $(L', \mathcal{E}') \succeq (L'', \mathcal{E}'')$, then the parallel composition $(G, \mathcal{B}) := (G_1, \mathcal{B}_1) \parallel (G_2, \mathcal{B}_2)$ works to realise the relation $(L, \mathcal{E}) \succeq (L'', \mathcal{E}'')$, using Freeness of Parallel Composition and Coherence of Nested Contraction. (The intuition is obvious: If (L', \mathcal{E}') can be constructed from (L, \mathcal{E}) and (L'', \mathcal{E}'') from (L', \mathcal{E}') , then by successive construction we can reach (L'', \mathcal{E}'') from (L, \mathcal{E}) .) The details are left as exercise. □

Remark 4.3.16. (On the ‘Renaming of Ports’ in Definition 4.3.14.)

The reader may wonder why it the renaming of ports cannot simply be absorbed into the action of (G, \mathcal{B}) ; this is ultimately due to our choice of formally defining contractibility for common sub-interfaces (rather than, say, by means of merely a bijective correspondence between sub-interfaces), and it is illustrated in Example 4.3.19. ✦

Remark 4.3.17. (On the Significance of (G', \mathcal{B}') in Condition (4.69).)

We will shortly give a simplifying reformulation of the condition of derivability (Lemma 4.3.21). Let us already now observe, however, that we may without loss of generality assume that

$$\mathbb{B}' = \mathbb{I} \text{ and that } (G', \mathcal{B}') = \text{tr}_{\mathbb{A}'}; \text{ indeed, } \begin{array}{c} \text{-X-} \\ \sim_{\mathbb{D}'} \sim \\ \boxed{(L', \mathcal{E}')} \\ \sim_{\mathbb{E}'} \sim \\ \text{-Y-} \end{array} \text{ is derivable from } \begin{array}{c} \text{-X-} \\ \sim_{\mathbb{D}'} \sim \\ \boxed{(L', \mathcal{E}')} \\ \sim_{\mathbb{E}'} \sim \\ \text{-Y-} \\ \sim_{\mathbb{A}'} \sim \\ \boxed{\text{tr}} \\ \sim_{\mathbb{B}'} \sim \end{array},$$

by parallelly composing the latter with $\sim_{\mathbb{B}'} \sim \boxed{\text{tr}} \sim_{\mathbb{B}'} \sim$ and contracting along \mathbb{B}' , using Soundness. The formulation (4.69) was merely chosen for the sake of symmetric appearance. ✦

Remark 4.3.18. (Notation – On the Dependence of $\succeq_{(T, \mathcal{C})}^{\mathbf{E}}$ on (T, \mathcal{C}) and \mathbf{E} .)

Like the dilational ordering \succeq_T in the causality-free case, the relation $\succeq_{(T, \mathcal{C})}^{\mathbf{E}}$ depends only on (T, \mathcal{C}) through the dependence on the open interfaces \mathbb{X} and \mathbb{Y} (cf. Remark 2.2.3). On the other hand, there is a strong dependence of $\succeq_{(T, \mathcal{C})}^{\mathbf{E}}$ on \mathbf{E} , since the channel (G, \mathcal{B}) used to construct (L', \mathcal{E}') in the environment must be an \mathbf{E} -channel. For example, in **QIT**, if \mathbf{E} is the constructible class, any Bell-channel can be used for (G, \mathcal{B}) , whereas the inconstructible PR box cannot.

In practice, we will abbreviate $\succeq_{(T, \mathcal{C})}^{\mathbf{E}}$ by $\succeq^{\mathbf{E}}$, or even by \succeq if \mathbf{E} is clear from the context. \spadesuit

The canonical examples of derivability are given by serial composition in the environment, and by contraction in the environment:

Example 4.3.19. (Derivation by Serial Composition.)

Let $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ be a dilation of $\text{---}\mathbb{X}\text{---}\boxed{(T, \mathcal{C})}\text{---}\mathbb{Y}\text{---}$. For any causal channel $\sim\mathbb{E}\sim\boxed{(G, \mathcal{B})}\sim\mathbb{E}'\sim$, the dilation $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \sim\mathbb{E}\sim\boxed{(G, \mathcal{B})}\sim\mathbb{E}'\sim \end{array}$ can be derived from (L, \mathcal{E}) , as witnessed by the parallel composition

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \sim\mathbb{E}\sim\boxed{(G, \mathcal{B})}\sim\mathbb{E}'\sim \end{array} \quad (4.70)$$

which by Soundness is contractible along \mathbb{E} with the desired contraction. Note, however, that we might need to intermittently rename the ports in \mathbb{E} so as to make them differ from those of \mathbb{E}' , \mathbb{D} and \mathbb{X} , in order for the channels to be parallelly composable in the first place (according to our requirement of distinct port names in a parallel composition).

This example serves to illustrate how the causal-dilational ordering generalises the one of Chapter 2 (note that (G, \mathcal{B}) could easily have had an additional input interface \mathbb{D}'). \blacklozenge

Example 4.3.20. (Derivation by Contraction.)

Let $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \text{---}\mathbb{Y}\text{---} \end{array}$ be a dilation of $\text{---}\mathbb{X}\text{---}\boxed{(T, \mathcal{C})}\text{---}\mathbb{Y}\text{---}$, and let $\begin{array}{c} \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \\ \boxed{(M, \mathcal{F})} \\ \sim\mathbb{K}\sim \\ \text{---}\mathbb{Z}\text{---} \end{array}$ be a dilation of $\text{---}\mathbb{Y}\text{---}\boxed{(S, \mathcal{D})}\text{---}\mathbb{Z}\text{---}$ whose hidden input interface matches the hidden output interface of (L, \mathcal{E}) . Of course, the causal channel

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \sim\mathbb{E}\sim\boxed{(M, \mathcal{F})}\sim\mathbb{K}\sim \\ \text{---}\mathbb{Z}\text{---} \end{array} \quad (4.71)$$

is a causal dilation of $\text{---}\mathbb{X}\text{---}\boxed{(T, \mathcal{C})}\text{---}\boxed{(S, \mathcal{D})}\text{---}\mathbb{Z}\text{---}$. (A better two-dimensional drawing of the channel (4.71) would have the \mathbb{E} -wires prolonged, crossing each other, so as to display all the input interfaces to the far left, and all the output interfaces to the far right.) The interface \mathbb{E} is contractible, and the contraction is given by

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \sim\mathbb{E}\sim \\ \sim\mathbb{E}\sim\boxed{(M, \mathcal{F})}\sim\mathbb{K}\sim \end{array} ; \quad (4.72)$$

consequently, (4.72) is dilation derivable from the dilation (4.71).

This derivation could not have been realised by ‘building clever circuits around’ the original dilation; at some point a contraction is needed, and ordinary serial and parallel composition will not accomplish it. A similar example can be based on the parallel composition of two dilations, rather than a serial. These two examples illustrate why we could not properly have treated a dilational ordering among two-sided dilations in Chapter 2 before introducing the formalism of causal channels and contractions.

◆

We shall now restate the definition of derivability in simpler terms, effectively demonstrating that Example 4.3.19 and Example 4.3.20 are indicative of what general derivability looks like, namely, a combination of serial composition and contraction of the inaccessible interface. We have the following:

Lemma 4.3.21. (*Simplified Statement of Derivability.*)

Let $\begin{array}{c} \text{---X---} \\ \sim\mathbb{D}\sim \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---Y---} \\ \sim\mathbb{E}\sim \end{array}$ and $\begin{array}{c} \text{---X---} \\ \sim\mathbb{D}'\sim \end{array} \boxed{(L', \mathcal{E}')} \begin{array}{c} \text{---Y---} \\ \sim\mathbb{E}'\sim \end{array}$ be causal dilations of (T, \mathcal{C}) , and assume that the naming of ports is such that the indicated interfaces are sufficiently disjoint (for example, all pairwise disjoint). Then, $(L, \mathcal{E}) \succeq^{\mathbf{E}} (L', \mathcal{E}')$ if and only if there exists an \mathbf{E} -channel (G, \mathcal{B}) such that

$$\mathfrak{C}_{\mathbb{D}} \left(\begin{array}{c} \text{---X---} \\ \sim\mathbb{D}\sim \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---Y---} \\ \sim\mathbb{E}\sim \end{array} \begin{array}{c} \text{---} \\ \sim\mathbb{D}'\sim \\ \sim\mathbb{A}'\sim \end{array} \boxed{(G, \mathcal{B})} \begin{array}{c} \text{---} \\ \sim\mathbb{D}\sim \\ \sim\mathbb{E}'\sim \end{array} \end{array} \right) = \begin{array}{c} \text{---X---} \\ \sim\mathbb{D}'\sim \end{array} \boxed{(L', \mathcal{E}')} \begin{array}{c} \text{---Y---} \\ \sim\mathbb{E}'\sim \end{array} \begin{array}{c} \text{---} \\ \sim\mathbb{A}'\sim \end{array} \boxed{\text{tr}} \quad , \quad (4.73)$$

with \mathbb{D} contractible. (If Θ has a state on every system, we may additionally take $\mathbb{A}' = \mathbb{I}$.)

Proof. The ‘if’-direction is clear. For the ‘only if’-direction, assume (G, \mathcal{B}) and (G', \mathcal{B}') are given such that (4.69) holds. Recall from Remark 4.3.17 that we may assume without loss of generality that $\mathbb{B}' = \mathbb{I}$ and $(G', \mathcal{B}') = \text{tr}_{\mathbb{A}'}$, since we may always (possibly after renaming inaccessible ports) parallelly compose with $\text{tr}_{\mathbb{B}'}$ and contract along \mathbb{B}' using Soundness. (If the theory Θ has states on every system, we may here similarly assume that $\mathbb{A}' = \mathbb{I}$, so that altogether the channel (G', \mathcal{B}') does not appear at all on the right hand side of Eq. (4.69).)

We thus have a causal channel (G, \mathcal{B}) and a contractible sub-interface \mathbb{Q} , such that

$$\mathfrak{C}_{\mathbb{Q}} \left(\begin{array}{c} \text{---X---} \\ \sim\mathbb{D}\sim \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---Y---} \\ \sim\mathbb{E}\sim \end{array} \begin{array}{c} \text{---} \\ \sim\mathbb{A}\sim \end{array} \boxed{(G, \mathcal{B})} \begin{array}{c} \text{---} \\ \sim\mathbb{B}\sim \end{array} \end{array} \right) = \begin{array}{c} \text{---X---} \\ \sim\mathbb{D}'\sim \end{array} \boxed{(L', \mathcal{E}')} \begin{array}{c} \text{---Y---} \\ \sim\mathbb{E}'\sim \end{array} \begin{array}{c} \text{---} \\ \sim\mathbb{A}'\sim \end{array} \boxed{\text{tr}_{\mathbb{A}'}} \quad (4.74)$$

(the naming assumption means that a potential renaming of hidden ports can be absorbed into (G, \mathcal{B})). Now, observe that we can always enlarge the contracted interface \mathbb{Q} so as to ensure that $\mathbb{Q} \supseteq \mathbb{D} \cup \mathbb{E}$, by replacing (G, \mathcal{B}) with $(G, \mathcal{B}) \parallel \text{id}_{(\mathbb{D} \cup \mathbb{E}) \setminus \mathbb{Q}}$ and redefining \mathbb{Q} as $\mathbb{Q} \cup [(\mathbb{D} \cup \mathbb{E}) \setminus \mathbb{Q}]$; indeed, contracting the ports in $(\mathbb{D} \cup \mathbb{E}) \setminus \mathbb{Q}$ with identities makes no difference, by Soundness. Under this modification, we actually have

$$\mathfrak{C}_{\mathbb{Q}} \left(\begin{array}{c} \text{---} \mathbb{X} \text{---} \boxed{(L, \mathcal{E})} \text{---} \mathbb{Y} \text{---} \\ \sim \mathbb{D} \sim \sim \mathbb{E} \sim \\ \sim \mathbb{E} \sim \sim \mathbb{D} \sim \\ \sim \mathbb{Q}_A \sim \sim \mathbb{Q}_B \sim \\ \sim \mathbb{A} \setminus \mathbb{Q}_A \sim \sim \mathbb{B} \setminus \mathbb{Q}_B \sim \end{array} \right) = \begin{array}{c} \text{---} \mathbb{X} \text{---} \boxed{(L', \mathcal{E}')} \text{---} \mathbb{Y} \text{---} \\ \sim \mathbb{D}' \sim \sim \mathbb{E}' \sim \\ \sim \mathbb{A}' \sim \boxed{\text{tr}_{\mathbb{A}'}} \end{array}, \quad (4.75)$$

where $\mathbb{Q}_A = \mathbb{Q} \cap \mathbb{A}$ and $\mathbb{Q}_B = \mathbb{Q} \cap \mathbb{B}$, and where $\mathbb{Q} = \mathbb{D} \cup \mathbb{E} \cup \mathbb{Q}_A = \mathbb{D} \cup \mathbb{E} \cup \mathbb{Q}_B$. This last identity implies that $\mathbb{Q}_A = \mathbb{Q}_B =: \mathbb{Q}_0$, and by Nested Contraction we can start by contracting the interface \mathbb{Q}_0 and may thus assume without loss of generality that actually $\mathbb{Q}_A = \mathbb{Q}_B = \mathbb{I}$, and thus $\mathbb{Q} = \mathbb{D} \cup \mathbb{E}$. But then, by comparing the remaining interfaces on each side of (4.75), we must have $\mathbb{A} = \mathbb{A} \setminus \mathbb{Q}_A = \mathbb{D}' \cup \mathbb{A}'$ and $\mathbb{B} = \mathbb{B} \setminus \mathbb{Q}_B = \mathbb{E}'$, so that Eq. (4.75) reads

$$\mathfrak{C}_{\mathbb{D} \cup \mathbb{E}} \left(\begin{array}{c} \text{---} \mathbb{X} \text{---} \boxed{(L, \mathcal{E})} \text{---} \mathbb{Y} \text{---} \\ \sim \mathbb{D} \sim \sim \mathbb{E} \sim \\ \sim \mathbb{E} \sim \sim \mathbb{D} \sim \\ \sim \mathbb{D}' \sim \sim \mathbb{E}' \sim \\ \sim \mathbb{A}' \sim \end{array} \right) = \begin{array}{c} \text{---} \mathbb{X} \text{---} \boxed{(L', \mathcal{E}')} \text{---} \mathbb{Y} \text{---} \\ \sim \mathbb{D}' \sim \sim \mathbb{E}' \sim \\ \sim \mathbb{A}' \sim \boxed{\text{tr}_{\mathbb{A}'}} \end{array}. \quad (4.76)$$

It now follows by Nested Contraction that we can start in Eq. (4.76) by contracting \mathbb{E} , and by Soundness that this contraction is given by serial composition, so we conclude Eq. (4.73) as desired. \square

Remark 4.3.22. (On Interpretation.)

Intuitively, Lemma 4.3.21 is not surprising. It merely says that any construction in the environment can be realised by adjoining all hidden outputs and inputs of (L, \mathcal{E}) to a network in the environment (one of them by a serial composition, the other necessarily by a contraction). The statement could equally well have been formulated using a serial composition on the \mathbb{D} -interface followed by a contraction of the \mathbb{E} -interface. \blacktimes

Remark 4.3.23. (Derivations from One-Sided Dilations.)

Lemma 4.3.21 implies that any dilation derivable from a one-sided dilation can be obtained by a serial composition with a causal channel (G, \mathcal{B}) . As such, when restricting to one-sided dilations the causal-dilational ordering looks similar to the dilational ordering of Chapter 2. \blacktimes

We end this subsection by observing two results to the effect that derivability interacts sensibly with parallel composition and contraction in the accessible interface (thus subsuming serial composition). Both results are easy to prove, boiling down to axiomatic gymnastics of contractions, similar to that of earlier proofs. They can be seen as companions to Proposition 4.3.10 and Proposition 4.3.11.

Theorem 4.3.24. (Parallel Composability of $\triangleright^{\mathbb{E}}$.)

Let $\begin{array}{c} \sim \mathbb{D}_j \sim \\ \sim \mathbb{X}_j \sim \end{array} \boxed{(L_j, \mathcal{E}_j)} \begin{array}{c} \sim \mathbb{E}_j \sim \\ \sim \mathbb{Y}_j \sim \end{array}$ and $\begin{array}{c} \sim \mathbb{D}'_j \sim \\ \sim \mathbb{X}'_j \sim \end{array} \boxed{(L'_j, \mathcal{E}'_j)} \begin{array}{c} \sim \mathbb{E}'_j \sim \\ \sim \mathbb{Y}'_j \sim \end{array}$ be causal dilations of $\text{---} \mathbb{X}_j \text{---} \boxed{(T_j, \mathcal{C}_j)} \text{---} \mathbb{Y}_j \text{---}$ for $j = 1, 2$, such that

$$\begin{array}{c} \text{---}\mathbb{X}_1\text{---} \\ \sim\mathbb{D}_1\sim \\ \boxed{(L_1, \mathcal{E}_1)} \\ \text{---}\mathbb{Y}_1\text{---} \\ \sim\mathbb{E}_1\sim \end{array} \succeq^{\mathbf{E}} \begin{array}{c} \text{---}\mathbb{X}_1\text{---} \\ \sim\mathbb{D}'_1\sim \\ \boxed{(L'_1, \mathcal{E}'_1)} \\ \text{---}\mathbb{Y}_1\text{---} \\ \sim\mathbb{E}'_1\sim \end{array} \quad \text{and} \quad \begin{array}{c} \text{---}\mathbb{X}_2\text{---} \\ \sim\mathbb{D}_2\sim \\ \boxed{(L_2, \mathcal{E}_2)} \\ \text{---}\mathbb{Y}_2\text{---} \\ \sim\mathbb{E}_2\sim \end{array} \succeq^{\mathbf{E}} \begin{array}{c} \text{---}\mathbb{X}_2\text{---} \\ \sim\mathbb{D}'_2\sim \\ \boxed{(L'_2, \mathcal{E}'_2)} \\ \text{---}\mathbb{Y}_2\text{---} \\ \sim\mathbb{E}'_2\sim \end{array} . \quad (4.77)$$

Then, it holds that

$$\begin{array}{c} \text{---}\mathbb{X}_1\text{---} \\ \sim\mathbb{D}_1\sim \\ \boxed{(L_1, \mathcal{E}_1)} \\ \text{---}\mathbb{Y}_1\text{---} \\ \sim\mathbb{E}_1\sim \\ \sim\mathbb{D}_2\sim \\ \boxed{(L_2, \mathcal{E}_2)} \\ \text{---}\mathbb{Y}_2\text{---} \\ \sim\mathbb{E}_2\sim \\ \text{---}\mathbb{X}_2\text{---} \end{array} \succeq^{\mathbf{E}} \begin{array}{c} \text{---}\mathbb{X}'_1\text{---} \\ \sim\mathbb{D}'_1\sim \\ \boxed{(L'_1, \mathcal{E}'_1)} \\ \text{---}\mathbb{Y}'_1\text{---} \\ \sim\mathbb{E}'_1\sim \\ \sim\mathbb{D}'_2\sim \\ \boxed{(L'_2, \mathcal{E}'_2)} \\ \text{---}\mathbb{Y}'_2\text{---} \\ \sim\mathbb{E}'_2\sim \\ \text{---}\mathbb{X}'_2\text{---} \end{array} . \quad (4.78)$$

Proof. The proof uses Coherence of Nested Contraction and Freeness of Parallel Composition and is left as a straightforward exercise. \square

Theorem 4.3.25. (Contractual Composability of $\succeq^{\mathbf{E}}$.)

Let $\begin{array}{c} \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \text{---}\mathbb{X}\text{---} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \end{array}$ and $\begin{array}{c} \sim\mathbb{D}'\sim \\ \boxed{(L', \mathcal{E}')} \\ \text{---}\mathbb{X}\text{---} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}'\sim \end{array}$ be causal dilations of $\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \boxed{(T, \mathcal{C})} \\ \text{---}\mathbb{Y}\text{---} \end{array}$, such that

$$\begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{(L, \mathcal{E})} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}\sim \end{array} \succeq^{\mathbf{E}} \begin{array}{c} \text{---}\mathbb{X}\text{---} \\ \sim\mathbb{D}'\sim \\ \boxed{(L', \mathcal{E}')} \\ \text{---}\mathbb{Y}\text{---} \\ \sim\mathbb{E}'\sim \end{array} . \quad (4.79)$$

Then, a sub-interface $\mathbb{P} \subseteq \mathbb{X} \cap \mathbb{Y}$ is contractible in (L, \mathcal{E}) if and only if it is contractible in (L', \mathcal{E}') , and in that case

$$\begin{array}{c} \text{---}\mathbb{X} \setminus \mathbb{P}\text{---} \\ \sim\mathbb{D}\sim \\ \boxed{\mathfrak{C}_{\mathbb{P}}((L, \mathcal{E}))} \\ \text{---}\mathbb{Y} \setminus \mathbb{P}\text{---} \\ \sim\mathbb{E}\sim \end{array} \succeq^{\mathbf{E}} \begin{array}{c} \text{---}\mathbb{X} \setminus \mathbb{P}\text{---} \\ \sim\mathbb{D}'\sim \\ \boxed{\mathfrak{C}_{\mathbb{P}}((L', \mathcal{E}'))} \\ \text{---}\mathbb{Y} \setminus \mathbb{P}\text{---} \\ \sim\mathbb{E}'\sim \end{array} . \quad (4.80)$$

Remark 4.3.26. Observe that, as usual, Theorem 4.3.25 about contractions subsumes serial composition, given Theorem 4.3.24 about parallel composition. \boxtimes

Proof. Contractions now occur on two distinct levels; on the explicitly stated level of contraction in the accessible interfaces \mathbb{X} and \mathbb{Y} , and on the level of contractions in the environment implicit within the definition of the pre-order $\succeq^{\mathbf{E}}$.

If (G, \mathcal{B}) realises the relation $(L, \mathcal{E}) \succeq^{\mathbf{E}} (L', \mathcal{E}')$, then the desired follows by Freeness under Parallel Composition and Coherence of Nested Contraction. Again, the details are left as exercise. (The statement about equivalence of contractibility follows since contractibility of \mathbb{P} in either dilation is equivalent to contractibility of \mathbb{P} in (T, \mathcal{C}) , by Proposition 4.3.11.) \square

4.4 Density Theorems and Rigidity

Now that we have understood the basic properties of causal dilations and the pre-order of derivability, it is natural to start executing the same program as for dilations in Chapter 2, namely, that of uncovering the structure of causal dilations of a given causal channel (T, \mathcal{C}) .

This time, however, the problem at hand is much more complex, and we will not be as successful in doing this. In this section, I will only make some initial observations about the

structure of causal dilations, and almost exclusively in specific examples. Whether a more systematic theory is possible is unclear, and it is probably the most important problem left open from the work of this thesis.

Again, we work in a fixed contractible theory $(\Theta, \mathfrak{C}, \mathbf{E})$. First, let us transfer to the setting of causal dilations the idea of *completeness*, which was one of the most powerful concepts of Chapter 2:

Definition 4.4.1. (Causal Completeness.)

Let $\text{--}\mathbb{X}\text{--}\overline{(T, \mathcal{C})}\text{--}\mathbb{Y}\text{--}$ be an \mathbf{E} -channel, and let $\mathbf{D} \subseteq \text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$ be a class of dilations of (T, \mathcal{C}) . A dilation $(K, \mathcal{F}) \in \mathbf{D}$ is called (*causally*) *complete for \mathbf{D}* , if $(K, \mathcal{F}) \succeq^{\mathbf{E}} (L, \mathcal{E})$ for all $(L, \mathcal{E}) \in \mathbf{D}$. A dilation $(K, \mathcal{F}) \in \text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$ is called simply (*causally*) *complete* if it is complete for $\text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$. ■

Example 4.4.2. By Example 4.3.4, both trashes and states have complete causal dilations. ◆

In Chapter 2, completeness was ubiquitous and its absence confined to obscure examples. *Causal* completeness, on the other hand, turns out to be a scarce resource, at least in the information theories **CIT** and **QIT**. The reality seems to be that causal completeness should no longer be regarded as a property of a (contractible) theory, cf. the idea of ‘complete theories’ in Chapter 2, but rather as a property of certain causal channels in the theory:

Definition 4.4.3. (Rigidity.)

An \mathbf{E} -channel $\text{--}\mathbb{X}\text{--}\overline{(T, \mathcal{C})}\text{--}\mathbb{Y}\text{--}$ is said to be *rigid* if it has a complete causal dilation. More generally, we call (T, \mathcal{C}) *rigid relative to $\mathbf{D} \subseteq \text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$* if (T, \mathcal{C}) has a causal dilation which is complete for \mathbf{D} . ■

Rigidity signifies that by locking the input-output behaviour (T, \mathcal{C}) between the accessible interfaces, there is a single ‘worst case’ causal dilation, formalising the strongest possible causal side-computations the environment can perform during our interaction with the causal channel (T, \mathcal{C}) . At this point, it is not clear that this property has anything to do with rigidity as it is understood in quantum self-testing (except, perhaps, for similarity in spirit), but the relationship will be uncovered in Chapter 5. Essentially, ‘rigidity’ as used in the field of quantum self-testing will correspond to rigidity relative to a certain class of dilations, but not relative to all; that some dilations have to be disregarded slightly disturbs the simplicity, but it is intimately related to the fact that measurement channels have strange dilations (Example 4.3.9).

How should we tackle the problem of clarifying the structure of dilations if we cannot generally rely on completeness?

As a substitute for complete (i.e. $\succeq^{\mathbf{E}}$ -largest) dilations of (T, \mathcal{C}) , one might look for $\succeq^{\mathbf{E}}$ -maximal¹⁶ dilations, and try to prove results to the effect that any dilation is derivable from some maximal dilation. Whereas this could indeed be successful, pre-orders in general need not exhibit such universal boundedness by maximal elements.¹⁷

As such, the best we can do is to identify a *dense* class of dilations, that is, a class¹⁸ \mathbf{D} of dilations such that for any dilation (L', \mathcal{E}') , there exists $(L, \mathcal{E}) \in \mathbf{D}$ with $(L, \mathcal{E}) \succeq^{\mathbf{E}} (L', \mathcal{E}')$.

¹⁶See the Preliminary section of the thesis for basic terminology of pre-orders.

¹⁷It seems to me hard to find principles which in this context facilitate an application of Zorn’s Lemma.

¹⁸This class has conceptually nothing to do with the class in the definition of relative completeness and rigidity, though we denoted it too by the symbol ‘ \mathbf{D} ’. Observe, however, that if \mathbf{D} is dense in $\text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$ and (K, \mathcal{F}) is complete for \mathbf{D} , then (K, \mathcal{F}) is in fact complete for all of $\text{CausDil}^{\mathbf{E}}((T, \mathcal{C}))$.

The smaller we can make the dense class, the better we have understood the large elements in the causal-dilational ordering, that is, the better we have understood the abilities of a potentially adversarial environment of the channel (T, \mathcal{C}) .

We will use the term *density theorem* about a result which identifies a certain class of dilations as dense. A *completeness theorem* would be an extreme density theorem, asserting that a class $\{(K, \mathcal{F})\}$ consisting of a single dilation is dense.

To make any sort of progress, we will from now on restrict attention to constructible channels, i.e. we will fix $\mathbf{E} = \text{Cons}(\Theta)$ and take \mathfrak{C} to be the standard notion of contraction. This is not only a mathematical convenience, but is also, at least in the case of the information theories **CIT** and **QIT**, physically reasonable.¹⁹ We will write simply \triangleright in place of $\triangleright_{\text{Cons}(\Theta)}$.

First, we show a general density theorem which holds by virtue of the constructibility restriction. The most important case of this result is that of Bell-channels. Then, we observe that causal channels in cartesian theories are always rigid. Finally, we examine density and rigidity for Bell-channels in **CIT**.

4.4.A ... in General

In approaching a general density theorem, the first observation to make is that, by Lemma 4.1.24, every constructible dilation of a causal channel $-_{\mathbb{X}}\text{--}\boxed{(T, \mathcal{C})}\text{--}_{\mathbb{Y}}$ is of the form

$$\begin{array}{c}
 \text{---}\mathbb{X}\text{---} \\
 \boxed{(L, \mathcal{E})} \\
 \text{---}\mathbb{E}\text{---} \\
 \text{---}\mathbb{D}'\text{---} \\
 \boxed{(G, \mathcal{B})} \\
 \text{---}\mathbb{E}'\text{---} \\
 \text{---}\mathbb{Y}\text{---}
 \end{array} \tag{4.81}$$

with (L, \mathcal{E}) and (G, \mathcal{B}) constructible. This implies that the class of one-sided constructible dilations of (T, \mathcal{C}) is dense. As such, the density question is reduced to finding a class which is dense *within* the class of one-sided dilations.

A first step in this direction is to clarify the complexity of the various possible hidden output interfaces \mathbb{E} . The ports $\mathbf{e} \in \text{ports}(\mathbb{E})$ distinguish themselves by their sets of causes $\mathcal{E}(\mathbf{e}) \subseteq \text{ports}(\mathbb{X})$. These cause-sets encode which ports in the open interface \mathbb{X} have to be fed with an input before side-information becomes available to the environment.

Example 4.4.4. (One-sided Dilations of a General Causal Channel.)

Let $-_{\mathbb{X}}\text{--}\boxed{(L, \mathcal{E})}\text{--}_{\mathbb{Y}}$ be a one-sided dilation of $-_{\mathbb{X}}\text{--}\boxed{(T, \mathcal{C})}\text{--}_{\mathbb{Y}}$. If two ports $\mathbf{e}, \mathbf{e}' \in \text{ports}(\mathbb{E})$ have the same causes, $\mathcal{E}(\mathbf{e}) = \mathcal{E}(\mathbf{e}')$, then they can be merged in the environment to a single port, and this merging can be reversed also by a causal channel in the environment. In other words, (L, \mathcal{E}) is \triangleright -equivalent to a one-sided dilation for which all hidden ports have distinct sets of causes. Consequently, a dense class of dilations is given by those one-sided dilations with (at most) $2^{|\mathbb{X}|}$ hidden ports, one for each possible set of causes in the open interface \mathbb{X} . \blacklozenge

A priori, the size of \mathbb{E} might not be further reducible than suggested by the exponential bound $2^{|\mathbb{X}|}$. But sometimes we can reduce it by virtue of (constructibility and) details of the causal specification \mathcal{C} . The following example not only serves to illustrate this point,

¹⁹Otherwise, even a causal channel as simple as the bit refreshment channel has as among its dilation the PR box, whose physical existence is questionable ([vD13]).

but will also be extremely important for establishing the relation to quantum self-testing in Chapter 5:

Example 4.4.5. (Constructible Dilations of a Bipartite Bell-Channel.)

Consider a bipartite Bell-channel, i.e. a constructible causal channel $\begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} -\mathcal{Y}_A- \\ -\mathcal{Y}_B- \end{array}$ with $\mathcal{C}(y_A) = \{x_A\}$ and $\mathcal{C}(y_B) = \{x_B\}$. By Example 4.4.4, any constructible dilation of (T, \mathcal{C}) is derivable from a one-sided constructible causal dilation with four hidden ports, corresponding to the cause-sets $\emptyset, \{x_A\}, \{x_B\}$ and $\{x_A, x_B\}$. It turns out that already the dilations with three types of hidden output ports – corresponding to cause-sets $\emptyset, \{x_A\}$ and $\{x_B\}$ – suffice to form a dense class. (As such, side-information with cause-set $\{x_A, x_B\}$ can always be derived from side-information with cause-sets $\{x_A\}$ and $\{x_B\}$.) The argument is essentially graph-theoretic:

Define the *stencil-complexity* of a constructible channel to be the minimal number of boxes required in a stencil-representation whose filling uses only primitive causal channels. Given any constructible dilation (L', \mathcal{E}') of (T, \mathcal{C}) , let $(L, \mathcal{E}) \succeq (L', \mathcal{E}')$ be a one-sided constructible dilation with smallest possible stencil-complexity. Let (\mathfrak{F}, G) be a stencil-representation of (L, \mathcal{E}) witnessing this stencil-complexity. We aim to show that (L, \mathcal{E}) has no hidden port with cause set $\{x_A, x_B\}$. This will prove the desired, since the remaining ports can then be merged according to causes.

To arrive at a contradiction, assume that for some $e \in \mathbb{E}$ we have $\mathcal{E}(e) = \{x_A, x_B\}$. Then in the stencil G , the port e must be the child²⁰ of some box, say b . Moreover, both x_A and x_B must be ancestral to b , and therefore x_A and x_B are ancestral to any descendant of b . In particular, any output port that descends from b has cause-set $\{x_A, x_B\}$. But by virtue of the specification \mathcal{C} , this implies that b cannot have y_A or y_B as descendants, so the ports which descend from b must correspond to a sub-interface \mathbb{E}_0 of \mathbb{E} . Pictorially, the situation now looks as follows:

$$\begin{array}{c} -\mathbb{X}- \\ \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} -\mathbb{Y}- \\ \sim \mathbb{E} \end{array} = \begin{array}{c} -\mathbb{X}- \\ \end{array} \boxed{(L_0, \mathcal{E}_0)} \begin{array}{c} -\mathbb{Y}- \\ \sim \mathbb{E} \setminus \mathbb{E}_0 \\ \sim \mathbb{H} \end{array} \boxed{T_b} \begin{array}{c} \sim \mathbb{E}_0 \end{array}, \quad (4.82)$$

where $\boxed{T_b}$ denotes the box b filled with its primitive channel T_b , and where (L_0, \mathcal{E}_0) abbreviates the remaining part of the representation (\mathfrak{F}, G) , that results from disregarding the box b . Now, however, if we remove the box b we are left with a one-sided constructible dilation $(L_0, \mathcal{E}_0) \succeq (L, \mathcal{E})$ of strictly lower stencil-complexity than (L, \mathcal{E}) , contradicting the choice of (L, \mathcal{E}) . Consequently, there cannot be any ports $e \in \text{ports}(\mathbb{E})$ with $\mathcal{E}(e) = \{x_A, x_B\}$, and the desired follows: Any constructible dilation of (T, \mathcal{C}) derives from a constructible

$$\begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} -\mathcal{Y}_A- \\ \sim \mathcal{E}_A \\ \sim \mathcal{E}_0 \\ \sim \mathcal{E}_B \\ -\mathcal{Y}_B- \end{array}, \text{ where } \mathcal{E}(e_0) = \emptyset, \mathcal{E}(e_A) = \{x_A\} \text{ and } \mathcal{E}(e_B) = \{x_B\}.$$

◆

It is clear that Example 4.4.5 can be easily generalised to the case of a multipartite Bell-channel, i.e. a constructible causal channel $\begin{array}{c} -\mathbb{X}- \\ \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} -\mathbb{Y}- \\ \end{array}$ with $\text{ports}(\mathbb{X}) = \{x_1, \dots, x_n\}$ and $\text{ports}(\mathbb{Y}) = \{y_1, \dots, y_n\}$, and for which the specification \mathcal{C} is given by $\mathcal{C}(y_j) = \{x_j\}$ for $j = 1, \dots, n$. As dense class of dilations we thus obtain those one-sided constructible

²⁰Recall that a vertex v in a directed graph is a *child* of the vertex v' if there exists an edge from v' to v .

dilations with $n + 1 = |\mathbb{X}| + 1$ hidden ports, one port for each $j = 1, \dots, n$ with cause-set $\{x_j\}$, and one port which is acausal i.e. has cause-set \emptyset . This is a tremendous improvement over the $2^{|\mathbb{X}|}$ ports that would a priori be expected.

In fact, the graph-theoretic argument we gave can even be generalised beyond Bell-channels without much difficulty.

Let $-\mathbb{X}-\boxed{(T, \mathcal{C})}-\mathbb{Y}-$ be an arbitrary constructible causal channel, but let us suppose for simplicity that $\mathbb{Y} \neq \mathbb{I}$ and $\mathcal{C}(\mathbb{Y}) = \mathbb{X}$. (Minding Lemma 4.1.26, this is to say that no trash can be factored out of (T, \mathcal{C}) .) Let us call a set of ports $P \subseteq \text{ports}(\mathbb{X})$ a *minorant of \mathcal{C} given y* if $P \subseteq \mathcal{C}(y)$. Let

$$\text{Min}_{\mathcal{C}} := \{P \subseteq \text{ports}(\mathbb{X}) \mid \exists y \in \text{ports}(\mathbb{Y}) : P \subseteq \mathcal{C}(y)\} \quad (4.83)$$

be the collection of minorants of \mathcal{C} (given some y). If \mathcal{C} is a primitive causal specification ($\mathcal{C}(y) = \text{ports}(\mathbb{X})$ for all $y \in \text{ports}(\mathbb{Y})$), then $\text{Min}_{\mathcal{C}}$ is the collection of all subsets of $\text{ports}(\mathbb{X})$, so $|\text{Min}_{\mathcal{C}}| = 2^{|\mathbb{X}|}$. At the other extreme, if $\mathcal{C}(y)$ contains just a single port $x_y \in \text{ports}(\mathbb{X})$ for every $y \in \text{ports}(\mathbb{Y})$ (like for the Bell-channels), then $\text{Min}_{\mathcal{C}}$ contains precisely the singletons $\{x_y\}$ along with the empty set \emptyset , so $|\text{Min}_{\mathcal{C}}| = |\mathbb{X}| + 1$ (since $y \mapsto x_y$ is surjective by the requirement $\mathcal{C}(\mathbb{Y}) = \mathbb{X}$).

We have the following:

Theorem 4.4.6. (*Density of Minorantal One-Sided Dilations.*)

Let $-\mathbb{X}-\boxed{(T, \mathcal{C})}-\mathbb{Y}-$ be a constructible causal channel with $\mathbb{Y} \neq \mathbb{I}$ and $\mathcal{C}(\mathbb{Y}) = \mathbb{X}$. Then the class of constructible one-sided dilations $-\mathbb{X}-\boxed{(L, \mathcal{E})}-\mathbb{E}-$ for which \mathbb{E} contains (at most) $|\text{Min}_{\mathcal{C}}|$ ports, one port e with $\mathcal{E}(e) = P$ for each minorant $P \in \text{Min}_{\mathcal{C}}$, is a dense class of dilations.

Remark 4.4.7. For the case $\mathbb{Y} = \mathbb{I}$ we already have a density theorem in the form of Example 4.3.4, indeed, the trashes have complete causal dilations. For $\mathbb{Y} \neq \mathbb{I}$ but $\mathcal{C}(\mathbb{Y}) \subsetneq \mathbb{X}$, one can factor out a maximal trash from (T, \mathcal{C}) and prove that the dense class of Theorem 4.4.6 can be combined in parallel with those complete dilations and yield a dense class. \blackstar

Proof. The proof closely follows that given in Example 4.4.5.

For any constructible dilation (L', \mathcal{E}') of (T, \mathcal{C}) , let $(L, \mathcal{E}) \supseteq (L', \mathcal{E}')$ be a one-sided constructible dilation with minimal stencil-complexity, and let (\mathfrak{F}, G) be a stencil-representation witnessing this. We prove that $\mathcal{E}(e) \in \text{Min}_{\mathcal{C}}$ for every $e \in \text{ports}(\mathbb{E})$, and this is enough.

Any port $e \in \text{ports}(\mathbb{E})$ is the child of some box b in the stencil G (if it were the child of an input port x , then the wire $x - e$ would be disconnected from the rest of G , so (T, \mathcal{C}) , which appears by trashing all the ports in \mathbb{E} , would have tr_x as a factor, contradicting $\mathcal{C}(\mathbb{X}) = \mathbb{Y}$). By the exact same argument as given in Example 4.4.5, b cannot only have descendants in $\text{ports}(\mathbb{E})$, since this would allow us to remove the box b and obtain a constructible dilation $(L_0, \mathcal{E}_0) \supseteq (L, \mathcal{E})$ of strictly lower stencil-complexity. Hence, b must have some $y \in \text{ports}(\mathbb{Y})$ as descendant. But then every ancestor of b is also an ancestor of y , so every ancestor of e is an ancestor of y . In particular, the cause-set $\mathcal{E}(e)$ is contained in the cause-set $\mathcal{E}(y) = \mathcal{C}(y)$, which is precisely to say that $\mathcal{E}(e) \in \text{Min}_{\mathcal{C}}$. \square

Theorem 4.4.6 is neat, in particular as it applies to Bell-channels in the form of Example 4.4.5, but it does not really bring us much closer to what causal dilations look like. The point is, however, that in certain cases we can now use the knowledge of the cause-sets $\mathcal{E}(e)$ to desiccate the dilations (L, \mathcal{E}) even further, by giving a concrete stencil-representation:

Theorem 4.4.8. (General Density Theorem for Bell-Channels, Part I.)

Let $\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \boxed{(T, \mathcal{E})} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ be a bipartite Bell-channel. Then, the causal dilations of the form

$$\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \begin{array}{c} \boxed{L_A} \\ \boxed{L_B} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \begin{array}{c} \mathcal{E}_A \\ \mathcal{E}_B \\ \mathcal{E}_0 \end{array} \quad , \quad (4.84)$$

where each component is given its primitive specification, constitute a dense class for constructible dilations of (T, \mathcal{E}) .

Remark 4.4.9. (General Bell-Channels.)

It should be clear how to generalise this statement to multipartite Bell-channels. In the case of a unipartite Bell-channel (that is, in the case of a primitive causal channel $-\mathcal{X} \boxed{(T, \mathcal{E})} \mathcal{Y}$),

the dense class is comprised by causal dilations of the form $\begin{array}{c} -\mathcal{X} \\ t \end{array} \boxed{L} \mathcal{Y} \begin{array}{c} \mathcal{E} \\ \mathcal{E}_0 \end{array}$. \boxtimes

Proof. We know from Example 4.4.5 that the constructible dilations

$$\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \begin{array}{c} \mathcal{E}_A \\ \mathcal{E}_0 \\ \mathcal{E}_B \end{array} \quad ,$$

where $\mathcal{E}(e_0) = \emptyset$, $\mathcal{E}(e_A) = \{x_A\}$ and $\mathcal{E}(e_B) = \{x_B\}$, form a dense class. However, using

Lemma 4.1.24 such a causal channel must be of the form

$$\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \begin{array}{c} \boxed{(S_A, \mathcal{D}_A)} \\ \boxed{(L_B, \mathcal{F}_B)} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \begin{array}{c} \mathcal{E}_A \\ \mathcal{E}_0 \\ \mathcal{E}_B \end{array}$$

for constructible (S_A, \mathcal{D}_A) and (L_B, \mathcal{F}_B) . Since $x_A \notin \mathcal{E}(\{y_B, e_B\})$, we may assume as in Example 4.1.27 without loss of generality that $\mathcal{D}_A(\text{ports}(\mathbb{H})) = \emptyset$ (if $\text{ports}(\mathbb{H})$ had \mathcal{D}_A -causes, these would propagate to $\{y_B, e_B\}$, unless trashed in (L_B, \mathcal{F}_B)). Using Lemma 4.1.24 again,

(S_A, \mathcal{D}_A) must therefore be of the form $\begin{array}{c} -\mathcal{X}_A \\ t \end{array} \boxed{(L_A, \mathcal{F}_A)} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{E}_A \end{array}$ for some state t and some

constructible (L_A, \mathcal{F}_A) . Altogether, (L, \mathcal{E}) is now of the form

$$\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \begin{array}{c} \boxed{(L_A, \mathcal{F}_A)} \\ \boxed{(L_B, \mathcal{F}_B)} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \begin{array}{c} \mathcal{E}_A \\ \mathcal{E}_0 \\ \mathcal{E}_B \end{array} \quad , \quad \text{with}$$

$x_i \in \mathcal{F}_i(y_i)$ and $x_i \in \mathcal{F}_i(e_i)$. Now, however, the total causal specification that results from this is precisely the same as when \mathcal{F}_A and \mathcal{F}_B are substituted by primitive specifications. Hence, we can make this substitution and achieve the desired. \square

The significance of Theorem 4.4.8 is the following: A channel of the form (4.84) is a dilation of (T, \mathcal{E}) if and only if t is a dilation a state s and the channels L_i are dilations of channels T_i , such that

$$\begin{array}{c}
-\mathcal{X}_A- \\
\boxed{s} \\
-\mathcal{X}_B-
\end{array}
\begin{array}{c}
\boxed{T_A} \\
\boxed{T_B}
\end{array}
\begin{array}{c}
-\mathcal{Y}_A- \\
-\mathcal{Y}_B-
\end{array}
=
\begin{array}{c}
-\mathcal{X}_A- \\
-\mathcal{X}_B-
\end{array}
\boxed{(T, \mathcal{E})}
\begin{array}{c}
-\mathcal{Y}_A- \\
-\mathcal{Y}_B-
\end{array}
. \quad (4.85)$$

Though it was of course already clear that (T, \mathcal{E}) must have this form (since it is a Bell-channel), it was not clear that any causal dilation of (T, \mathcal{E}) must derive from a dilation obtained by simply dilating every component in a triple of possible components. That result also brings us for the first time promisingly close to the concepts of quantum self-testing: The various triples (T_A, T_B, s) will correspond essentially to *quantum strategies* which produce a given input-output behaviour.

Now, in a general theory it seems that we cannot achieve a better density theorem (in the case of Bell-channels) than Theorem 4.4.8. In specific theories, however, certain features often make an improvement possible. As such, it is essential to have a criterion which expresses when two dilations of the form (4.84) are related in the pre-order \succeq . Indeed, if one is derivable from another, then the weaker one can be removed from the class without losing density. We have the following:

Theorem 4.4.10. (General Density Theorem for Bell-Channels, Part II.)

Let $\begin{array}{c} -\mathcal{X}_A- \\ \boxed{(T, \mathcal{E})} \\ -\mathcal{X}_B- \end{array}$ be a bipartite Bell-channel. Then, two causal dilations of the form

$$(4.84) \quad \begin{array}{c} -\mathcal{X}_A- \\ \boxed{t} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} \boxed{L_A} \\ \boxed{L_B} \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}_0 \\ -\mathcal{Y}_B- \end{array} \succeq \begin{array}{c} -\mathcal{X}_A- \\ \boxed{t'} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} \boxed{L'_A} \\ \boxed{L'_B} \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}'_0 \\ -\mathcal{Y}_B- \end{array} \quad \text{if and only if there exist channels } F, G_A$$

and G_B such that

$$\begin{array}{c} -\mathcal{X}_A- \\ \boxed{t} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} \boxed{L_A} \\ \boxed{L_B} \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}_0 \\ -\mathcal{Y}_B- \end{array} \begin{array}{c} \boxed{G_A} \\ \boxed{F} \\ \boxed{G_B} \end{array} \begin{array}{c} -\mathcal{Y}'_A- \\ \mathcal{E}'_0 \\ -\mathcal{Y}'_B- \end{array} = \begin{array}{c} -\mathcal{X}_A- \\ \boxed{t'} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} \boxed{L'_A} \\ \boxed{L'_B} \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}'_0 \\ -\mathcal{Y}_B- \end{array}
. \quad (4.86)$$

Remark 4.4.11. (General Bell-channels.)

Again, it should be clear how to generalise the statement to multipartite Bell-channels. In the case of a unipartite Bell-channel (i.e. of a primitive causal channel $-\mathcal{X}-\boxed{(T, \mathcal{E})}-\mathcal{Y}-$),

we have $\begin{array}{c} -\mathcal{X}_A- \\ \boxed{t} \\ \mathcal{E}_0 \end{array} \boxed{L} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E} \end{array} \succeq \begin{array}{c} -\mathcal{X}_A- \\ \boxed{t'} \\ \mathcal{E}'_0 \end{array} \boxed{L'} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}' \end{array}$ if and only if there exist channels G and F such that

$$\begin{array}{c} -\mathcal{X}_A- \\ \boxed{t} \\ \mathcal{E}_0 \end{array} \boxed{L} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E} \end{array} \begin{array}{c} \boxed{G} \\ \boxed{F} \end{array} \begin{array}{c} -\mathcal{Y}'_A- \\ \mathcal{E}' \end{array} = \begin{array}{c} -\mathcal{X}_A- \\ \boxed{t'} \\ \mathcal{E}'_0 \end{array} \boxed{L'} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}' \end{array}
. \quad (4.87)$$

✠

Proof. By Lemma 4.3.21, the derivability relation holds if and only if there exists a constructible causal channel (G, \mathcal{B}) such that

$$\begin{array}{c}
 \begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{Y}_B\text{---} \\
 \begin{array}{|c|} \hline L_A \\ \hline \end{array} \\
 \text{---}\mathcal{E}_A\text{---} \\
 \text{---}\mathcal{E}_0\text{---} \\
 \begin{array}{|c|} \hline L_B \\ \hline \end{array} \\
 \text{---}\mathcal{E}_B\text{---} \\
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{Y}_B\text{---} \\
 \begin{array}{|c|} \hline t \\ \hline \end{array}
 \end{array}
 \begin{array}{|c|} \hline (G, \mathcal{B}) \\ \hline \end{array}
 \begin{array}{c}
 \text{---}\mathcal{E}'_A\text{---} \\
 \text{---}\mathcal{E}'_0\text{---} \\
 \text{---}\mathcal{E}'_B\text{---}
 \end{array}
 =
 \begin{array}{c}
 \begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{Y}_B\text{---} \\
 \begin{array}{|c|} \hline L'_A \\ \hline \end{array} \\
 \text{---}\mathcal{E}'_A\text{---} \\
 \text{---}\mathcal{E}'_0\text{---} \\
 \begin{array}{|c|} \hline L'_B \\ \hline \end{array} \\
 \text{---}\mathcal{E}'_B\text{---} \\
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{Y}_B\text{---} \\
 \begin{array}{|c|} \hline t' \\ \hline \end{array}
 \end{array}
 \end{array}
 . \quad (4.88)
 \end{array}$$

Since the causal specifications on each side must match, we conclude that $\mathcal{B}(e'_0) \subseteq \{e_0\}$, $\mathcal{B}(e'_A) \subseteq \{e_0, e_A\}$ and $\mathcal{B}(e'_B) \subseteq \{e_0, e_B\}$. Now, however, it can be demonstrated by an argument slightly more general than in the proof of Theorem 4.4.8 (when arguing about the structure of (L, \mathcal{E})) that if a constructible (G, \mathcal{B}) has a specification with those properties, then it is necessarily of the form represented in Eq. (4.86). (Importantly, we rely in this proof on the fact that we fixed the scheme \mathbf{E} of allowed causal channels to be the constructible causal channels.) \square

4.4.B ... in Cartesian Theories

In cartesian theories, the above observations are not particularly useful, except perhaps in the case of Bell-channels. However, it is quite easy to see that, regardless of whether the channel (T, \mathcal{C}) is a Bell-channel or not, it has a causally complete dilation:

Theorem 4.4.12. (*Completeness and Rigidity in Cartesian Theories.*)

Assume that Θ is a cartesian theory, and let $\text{---}\mathcal{X}\text{---}\boxed{(T, \mathcal{C})}\text{---}\mathcal{Y}\text{---}$ be any (constructible) causal channel in Θ . Then (T, \mathcal{C}) is rigid, and a complete causal dilation is given by

$$\begin{array}{c}
 \begin{array}{c}
 \text{---}\mathcal{X}_1\text{---} \\
 \vdots \\
 \text{---}\mathcal{X}_n\text{---}
 \end{array}
 \begin{array}{|c|} \hline (T, \mathcal{C}) \\ \hline \end{array}
 \text{---}\mathcal{Y}\text{---}
 \end{array}
 , \quad (4.89)$$

where the individual copy channels are given their primitive causal specifications. (Observe in this regard that the copy channels are formally also boxes in the stencil.)

Remark 4.4.13. Note that a restriction to constructible channels in Θ is void, since by Theorem 4.1.18 all causal channels in a cartesian theory are constructible. \blacktimes

Proof. It suffices to show that any one-sided causal dilation of (T, \mathcal{C}) can be derived from the dilation (4.89). If however (L, \mathcal{E}) is a one-sided dilation, then the underlying channel L can be derived by completeness of copying for channels (Theorem 2.3.6), so only the causal specification \mathcal{E} remains to be accounted for. But we can argue abstractly as earlier, when we discussed the standard notion of contraction, by noting that the collection of causal specifications themselves form a cartesian theory themselves, that \mathcal{E} is as such a dilation of \mathcal{C} , and that the causal specification of the channel (4.89) is a complete dilation of the specification \mathcal{C} . More down-to-earth arguments are also possible, and left as exercise. \square

4.4.C ... in CIT

In the theory **CIT** things are much more complicated than in cartesian theories, and here the general discussion of Section 4.4.A can be put to use. We will restrict attention entirely to Bell-channels.

Let us start by considering the case of a unipartite Bell-channel. Put differently, this is really to study the causal dilations of a channel between simple interfaces, $-X-\boxed{(T, \mathcal{C})}-Y-$, where \mathcal{C} is the primitive specification given by $\mathcal{C}(y) = \{x\}$. (Note that this scenario includes the ‘bit refreshment’ channel of Example 4.3.7.) The analysis of this simple case is already rather involved, culminating with Theorem 4.4.15.

By Theorem 4.4.8, any constructible dilation of (T, \mathcal{C}) is derivable from a dilation of the form

$$\begin{array}{c} \text{---} X \text{---} \\ \text{---} R \text{---} \end{array} \boxed{L} \begin{array}{c} \text{---} Y \text{---} \\ \text{---} E \text{---} \end{array} , \quad (4.90)$$

where the state t and the channel L are given their primitive specifications. Due to special features of **CIT**, this density statement can be improved; we shall now argue that effectively any randomness involved in L can be pushed to E_0 as acausal side-information, thus moving up higher in the causal-dilational ordering and thereby narrowing the dense class:


Every classical channel is a convex combination of deterministic channels (functions). In particular, this holds for the channel $-X-\boxed{T}-Y-$ $:=$ $-X-\boxed{L}-Y-$ $\sim_{E \sim \text{tr}}$; but this is to say that

$$\begin{array}{c} \text{---} X \text{---} \\ \text{---} R \text{---} \end{array} \boxed{T} \begin{array}{c} \text{---} Y \text{---} \end{array} = \begin{array}{c} \text{---} X \text{---} \\ \text{---} R \text{---} \\ \boxed{q} \text{---} H \end{array} \boxed{f} \begin{array}{c} \text{---} Y \text{---} \end{array} \quad (4.91)$$


for some deterministic channel (function) f and some state q on some system H . Forgetting for a moment about causality, L is evidently a dilation of the channel (4.91), so by completeness and localisability of **CIT** (Theorem 2.3.21 and Theorem 2.3.17) the channel L can be derived (in the dilational ordering of Chapter 2) from the dilation given by copying q and the inputs to f , that is, there exists a channel G such that

$$\begin{array}{c} \text{---} X \text{---} \\ \text{---} R \text{---} \end{array} \boxed{L} \begin{array}{c} \text{---} Y \text{---} \\ \text{---} E \text{---} \end{array} = \begin{array}{c} \text{---} X \text{---} \\ \text{---} R \text{---} \\ \boxed{q} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{f} \\ \boxed{G} \end{array} \begin{array}{c} \text{---} Y \text{---} \\ \text{---} E \text{---} \end{array} . \quad (4.92)$$

(Clearly, we need not include two copies of q ; one suffices.) Turning causality back on, we see that by giving each component on the right hand side its primitive specification, L indeed gains the primitive specification, so in summary we can rewrite (4.90) as




with each component given its primitive specification. Now, by simply removing G we obviously rise in the causal-dilational ordering – in fact, we generically rise strictly, since, importantly, the copy of q transits from contributing to the side-information available after the input to x (as stalled by G) to acausal side-information (cf. Example 4.3.8). A similar thing can be said about the copy of the R -system, and by merging the three acausal hidden outputs in the environment, we conclude in summary that the causal dilation Eq. (4.90) is derivable from a causal dilation of the form



so such dilations form a dense class.

One further improvement is possible by noting that since \tilde{t} is a dilation of its R -marginal, it can be derived from a copy of that marginal by completeness of copying (Theorem 2.3.21). Hence, we have altogether proved the following density theorem for constructible dilations:

Theorem 4.4.14. *(Preliminary Density Theorem for Unipartite Bell-Channels in CIT.)*
 Let $-X-\boxed{(T, \mathcal{C})}-Y-$ be a primitive causal channel in **CIT** between simple interfaces. Then the class of causal dilations of the form



where p is a state on some system R and f a deterministic channel (function), and where each component is given its primitive causal specification, is a dense class of dilations.

Theorem 4.4.10 tells us when two dilations of the form (4.95) are related in the causal-dilational ordering, and will allow us to investigate whether a further thinning of this dense class is possible. But another question merits attention first: How can we interpret the dilations (4.95)?

This question is actually easy to answer. The dilations (4.95) are indexed by pairs (f, p) , and to understand which pairs give rise to a dilation we simply need to observe that the

channel (4.95) constitutes a dilation of (T, \mathcal{C}) precisely when $\begin{array}{c} -X- \\ \boxed{p} \end{array} \begin{array}{c} \boxed{f} \\ -Y- \end{array} = -X- \boxed{T} -Y-$. If we consider the function $f : X \times R \rightarrow Y$ equivalently as a family of functions $f_r := f(\cdot, r) : X \rightarrow Y$ indexed by $r \in R$, then this requirement is the condition

$$\sum_{r \in R} p(r) f_r = T, \quad (4.96)$$

in other words, *the dilations (4.95) correspond to convex decompositions of T into deterministic functions.*

In this light, Theorem 4.4.14 is perhaps not so surprising – what it says, basically, is that any causal dilation of T can be thought of as deriving from one in which the agent in the environment draws a random $r \in R$ according to the distribution p , keeps a copy of r in memory while using another copy to choose a function f_r to apply to our input, copies our input $x \in X$ to memory and gives us the value $f_r(x)$ as output. If we think about it, this exactly fits the scheme of the two causal dilations of the bit refreshment channel from Example 4.3.7; the two dilations correspond to the convex decompositions $\frac{1}{2}f_0 + \frac{1}{2}f_1$, where $f_k : \{0, 1\} \rightarrow \{0, 1\}$ is the function that is constantly k , respectively $\frac{1}{2}\text{id}_{\{0,1\}} + \frac{1}{2}\text{NOT}$, where $\text{NOT} : \{0, 1\} \rightarrow \{0, 1\}$ is the bit flip.

This understanding of the nature of the dilations also gives a hint on how to thin the dense class further. First of all, we may clearly assume that p has full rank on R , i.e. $p(r) > 0$ for all $r \in R$. Secondly, any convex decomposition (f, p) of T can intuitively be reduced to one for which $f_r \neq f_{r'}$ whenever $r \neq r'$; this comes about by merging any instances of r, r' which violate this. Formally, we replace R by the set of equivalence classes R/\sim under the equivalence relation $r \sim r' \Leftrightarrow f_r = f_{r'}$, we replace $(f_r)_{r \in R}$ by $(\tilde{f}_{[r]})_{[r] \in R/\sim}$, well-defined by $\tilde{f}_{[r]} = f_r$, and we replace p by $\tilde{p} : R/\sim \rightarrow [0, 1]$ given by $\tilde{p}([r]) = \sum_{r' \in [r]} p(r')$; the original dilation (4.95) can be derived from the new dilation given by (\tilde{f}, \tilde{p}) , since we can draw a value of r in the environment conditional on its equivalence class. In summary, the dense class of Theorem 4.4.14 can be thinned to those dilations (4.95) for which the pair (f, p) correspond to a proper convex decomposition of T , that is, one for which $p(r) > 0$ for all r and $f_r \neq f_{r'}$ for $r \neq r'$. That dense class essentially cannot be thinned any further, however, for what we have done is to identify a class of inequivalent maximal causal dilations:

Theorem 4.4.15. (*Ultimate Density Theorem for Unipartite Bell-Channels in CIT.*)

Let $-X- \boxed{(T, \mathcal{C})} -Y-$ be a primitive causal channel in **CIT** between simple interfaces. Then the class of causal dilations of the form

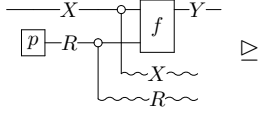
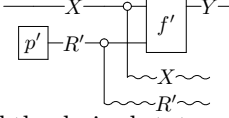
$$\begin{array}{c} -X- \\ \boxed{p} \end{array} \begin{array}{c} \boxed{f} \\ -Y- \end{array}, \quad (4.97)$$

where p is a full-rank state on some system R , and where $f : X \times R \rightarrow Y$ is a function with $f_r \neq f_{r'}$ for $r \neq r'$, is a dense class of dilations. Moreover, each such dilation is \triangleright -maximal, and the dilations given by (f, p) and (f', p') are \triangleright -equivalent if and only they are related by a relabelling of the elements in R , i.e. there exists a bijection $h : R \rightarrow R'$ such that $p(r) = p'(h(r))$ and $f_r = f'_{h(r)}$ for all $r \in R$.

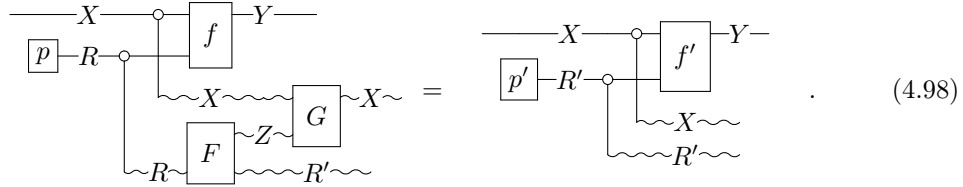
Remark 4.4.16. (The Dilational Ordering versus the Causal-Dilational Ordering.)

It is worth observing that if we forget about causality, then the dilations (4.97) are all complete dilations of the underlying channel T . In particular, all these dilations are equivalent under the dilational ordering of Chapter 2, that is, there always exist channels relating them in the environment. The significance of the causal-dilational ordering, which renders some of them inequivalent, is that those channels in the environment are required to take a specific form so as to preserve causality. \clubsuit

Proof. We have already argued that the class is dense. It remains to show maximality of each dilation, and the criterion for equivalence.

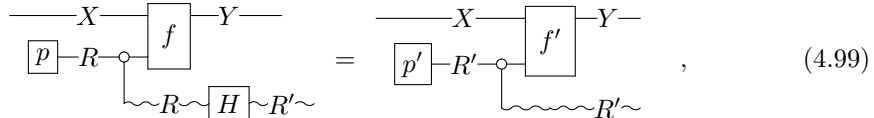
Assume we are given two comparable dilations within the class, say  \succeq . We show that the relabelling criterion is satisfied. This will prove all the desired statements (since the relabelling condition is clearly sufficient for equivalence).

First, note that Theorem 4.4.10 implies the existence of channels F and G such that



$$\text{Diagram of } f \text{ with } p \text{ and } R \text{ inputs} = \text{Diagram of } f' \text{ with } p' \text{ and } R' \text{ inputs} \text{ via } F \text{ and } G. \quad (4.98)$$

By trashing the hidden copy of X , we see that



$$\text{Diagram of } f \text{ with } p \text{ and } R \text{ inputs} = \text{Diagram of } f' \text{ with } p' \text{ and } R' \text{ inputs via } H. \quad (4.99)$$

where H is the marginal of F . Let us prove that this identity can only hold if $H =: h$ is in fact a deterministic bijection from R to R' ; this implies the desired.

Given $r \in R$ and $r' \in R'$, let us write $\mathbb{P}(H(r) = r')$ for the ‘probability that H maps r to r' ’, that is, for the quantities defined by $H(\delta_r) = \sum_{r' \in R'} \mathbb{P}(H(r) = r') \delta_{r'}$. Algebraically, Eq. (4.99) is then the identity $\sum_{r' \in R'} \sum_{r \in R} p(r) \mathbb{P}(H(r) = r') f_r \otimes \delta_{r'} = \sum_{r' \in R'} p'(r') f'_{r'} \otimes \delta_{r'}$.

Comparing terms, this identity is equivalent the identities

$$f'_{r'} = \sum_{r \in R} \frac{p(r)}{p'(r')} \mathbb{P}(H(r) = r') f_r \quad \text{for all } r' \in R'. \quad (4.100)$$

(These have the following interpretation: $f'_{r'}$ is the probabilistic mixture of the f_r ’s, with weights determined by the conditional distribution of r given $r' = H(r)$.) But now the desired follows: For any fixed $r' \in R'$, extremality of the channel $f'_{r'}$ and the fact that $p(r) > 0$ for all $r \in R$ implies by Eq. (4.100) that $\mathbb{P}(H(r) = r') \in \{0, 1\}$ for all $r \in R$,

and that there must be a unique (since the f_r 's are all distinct) r with $P(H(r) = r') = 1$. This is precisely to say that $H =: h$ is deterministic, injective and (since r' was arbitrary) surjective. \square

Theorem 4.4.15 ends our analysis of the unipartite case, and the conclusion is clear: Every causal dilation of (T, \mathcal{C}) is derivable from a maximal causal dilation, and the maximal causal dilations correspond precisely to the distinct convex decompositions of T into deterministic functions. As such, *the study of causal dilations of (T, \mathcal{C}) is the study of convex decompositions of T .*

The following consequence is immediate:

Corollary 4.4.17. (*Rigidity of Unipartite Bell-Channels in CIT.*)

A primitive causal channel between simple interfaces $-X-\boxed{(T, \mathcal{C})}-Y-$ in CIT is rigid if and only if T admits a unique convex decomposition into functions from X to Y .

Example 4.4.18. (Extremal implies Rigid.)

If $-X-\boxed{(T, \mathcal{C})}-Y-$ is deterministic, it is rigid. \blacklozenge

Example 4.4.19. (Rigid does not imply Extremal.)

If $|X| = 1$, then $-X-\boxed{(T, \mathcal{C})}-Y-$ corresponds to a state on Y , which is released upon giving an input ('pushing a button'). Any state in CIT has a unique convex decomposition into deterministic functions (pure states), hence every such causal channel is rigid. In particular, extremality of T is not necessary for rigidity. \blacklozenge

Example 4.4.20. (Non-Rigidity of Bit Refreshment.)

The bit refreshment dilations from Example 4.3.7 are maximal and inequivalent, as they correspond to the distinct convex decompositions $\frac{1}{2}f_0 + \frac{1}{2}f_1$ and $\frac{1}{2}\text{id}_{\{0,1\}} + \frac{1}{2}\text{NOT}$, as described earlier. Hence, the bit refreshment channel is not rigid. \blacklozenge

The non-rigidity of the bit refreshment channel has another peculiar consequence:


Remark 4.4.21. (Rigidity is Non-Composable – Temporal Localisability fails for Causal Dilations.)

Consider the two causal channels $-\{0,1\}-\boxed{\text{tr}}-\mathbf{1}-$ and $-\mathbf{1}-\boxed{r}-\{0,1\}-$ in CIT, both equipped with their primitive specifications. By Example 4.3.4 and Example 4.3.5, both of them have complete causal dilations (so they are rigid). However, their serial composition is the bit refreshment channel, which is not rigid; in particular, the composition of the complete causal dilations for $-\{0,1\}-\boxed{\text{tr}}-\mathbf{1}-$ and $-\mathbf{1}-\boxed{r}-\{0,1\}-$ is not a complete causal dilation of the bit refreshment, that is, the principle of temporal localisability fails for causal channels in CIT. \blackcross

Now, let us move on to the multipartite case. The analysis that led us in the unipartite case to the preliminary density theorem Theorem 4.4.14 generalises without difficulty (for simplicity, it is stated merely for the bipartite case):

Theorem 4.4.22. (*Preliminary Density Theorem for Bipartite Bell-Channels in CIT.*)

Let $-\begin{matrix} X_A \\ X_B \end{matrix}-\boxed{(T, \mathcal{C})}-\begin{matrix} Y_A \\ Y_B \end{matrix}-$ be a bipartite Bell-channel in CIT. Then the class of causal dilations of the form




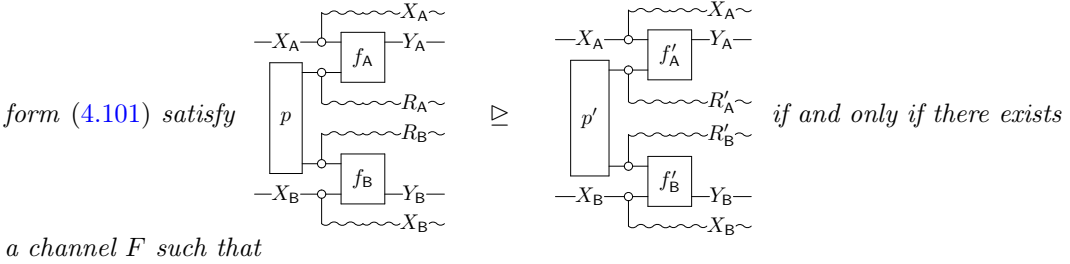
where p is a state on some system $R_A \times R_B$, where f_A and f_B are deterministic channels (functions), and where each component is given its primitive causal specification, is a dense class of dilations.

Clearly, we may in Theorem 4.4.22 moreover restrict to triples (f_A, f_B, p) for which the state p has locally full rank, meaning that the A- and B-marginals of p both have full rank. However, the rest of the analysis that led us in the unipartite case to the improved density theorem Theorem 4.4.15 is not immediately transferable. In particular, that p has locally full rank does not imply that p itself has full rank, indeed there can be strong correlations between the A- and B-parts of the state.

In fact, I will leave open the problem of finding an equivalent to Theorem 4.4.15, characterising the \succeq -maximal causal dilations and \succeq -equivalence among them. The following improvement of Theorem 4.4.10 (which appeared in disguise in the proof of Theorem 4.4.15 in the unipartite case) might be helpful in this regard:


Theorem 4.4.23. (Relations within the Dense Class for Bipartite Bell-Channels.)

Let $\begin{matrix} -X_A \\ -X_B \end{matrix} \boxed{(T, \mathcal{C})} \begin{matrix} Y_A \\ Y_B \end{matrix}$ be a bipartite Bell-channel in **CIT**. Then, two causal dilations of the



Proof. The ‘only if’-direction is clear from Theorem 4.4.10, by trashing the hidden ports corresponding to X_A and X_B (this is also what we did in the proof of Theorem 4.4.15). The ‘if’-

direction follows simply by noting that the identity



follows from (4.102). □

The fact that we have not obtained a multipartite counterpart to Theorem 4.4.15 in principle also leaves open the problem of finding a characterisation of rigidity in the multipartite case (i.e. an equivalent to Corollary 4.4.17). We do, however, have the following, which for simplicity is again only stated in the bipartite case:

Corollary 4.4.24. (*Sufficient Condition for Rigidity of Bipartite Bell-Channels in CIT.*)

Let $\begin{array}{c} -X_A \\ -X_B \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} -Y_A \\ -Y_B \end{array}$ be a bipartite Bell-channel in **CIT**. If T admits a unique convex decomposition into deterministic product channels, then (T, \mathcal{C}) is rigid.

Explicitly, (T, \mathcal{C}) is rigid if there exists a unique probability density $p : Y_A^{X_A} \times Y_B^{X_B} \rightarrow [0, 1]$ (here, $Y_i^{X_i}$ denotes the set of functions from X_i to Y_i) such that

$$T = \sum_{g_A, g_B} p(g_A, g_B) g_A \times g_B. \quad (4.103)$$

Proof. By an argument similar to that used in the unipartite case, we may restrict the dense class of dilations (4.101) to ensure that $f_A(\cdot, r_A) \times f_B(\cdot, r_B) \neq f_A(\cdot, r'_A) \times f_B(\cdot, r'_B)$ whenever the pairs (r_A, r_B) and (r'_A, r'_B) are distinct. However, the uniqueness of p in the representation (4.103) means that there is, up to relabelling, only one dilation of this kind, so it must be a complete dilation. □

Example 4.4.25. (Bipartite Rigidity without Marginal Rigidity.)

The bipartite Bell-channel $\begin{array}{c} -\{0, 1\} \\ -\{0, 1\} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \{0, 1\} \\ \{0, 1\} \end{array}$ with $T = \frac{1}{2} \text{id}_{\{0, 1\}} \otimes \text{NOT} + \frac{1}{2} \text{NOT} \otimes \text{id}_{\{0, 1\}}$ is rigid. Both of its marginals are $\frac{1}{2} \text{id}_{\{0, 1\}} + \frac{1}{2} \text{NOT} = \frac{1}{2} f_0 + \frac{1}{2} f_1$, i.e. they are the bit refreshment channel, which is not rigid. Hence, the rigidity of (T, \mathcal{C}) resides in a sense in the correlation between the local behaviours. ◆

Open Problem 4.4.26. (*Rigidity of Multipartite Bell-Channels.*)

Corollary 4.4.24 says that the uniqueness of convex decomposition into deterministic product channels implies rigidity. Is the converse true?

Open Problem 4.4.27. (*An Ultimate Density Theorem for Multipartite Bell-Channels.*)

For a bipartite (or generally multipartite) Bell-channel (T, \mathcal{C}) , is there an equivalent to Theorem 4.4.15, i.e. a dense class of \succeq -maximal causal dilations, and what is a criterion for \succeq -equivalence among such dilations?

Open Problem 4.4.28. (*Analysis of Other Causal Channels.*)

Is it possible to give an analysis of the causal dilations of a constructible causal channel in **CIT** which is not a Bell-channel?

4.5 Summary and Outlook

In this chapter, we have defined *causal channels* (Definition 4.1.3) and *causal dilations* (Definition 4.3.1), both of which have been heavily exemplified. The result that ultimately justifies our definition of a causal channel is the fact that *contractions* can be well-defined from knowledge of a channel and its causal specification alone (Theorem 4.2.6). An especially important class of causal channels are the *constructible* channels (Definition 4.1.15), which model physically realisable causal channels.

We have also seen how one may define *notions of contraction* in general (Definition 4.2.8), though the simplicity of this idea is slightly contaminated by the fact that we had to introduce *schemes* of causal channels (Definition 4.2.7), which physically represent ‘all sensible causal channels’ (e.g. the constructible ones), but mathematically are primarily needed so as to facilitate the *Coherence of Nested Contraction* (as for the constructible scheme in Theorem 4.2.13).

Given (a scheme and) a notion of contraction, it is possible to define the *causal-dilational ordering* (Definition 4.3.14), which is our final model for derivability among various dilations, hinted at already in the general introduction to the thesis. We have proved a handful of stability results about the causal-dilational ordering (Theorem 4.3.13, Theorem 4.3.24, Theorem 4.3.25) which all serve to consolidate the concepts of causal dilations and derivability.

Finally, we introduced the idea of *density theorems* and *rigidity* (Section 4.4), and have seen this exemplified most interestingly in the theory **CIT**, where in particular we saw that the bit refreshment channel is an example of non-rigidity since it allows two distinct convex decompositions (Example 4.4.20).

Several problems remain open for future investigation:

1. Open Problem 4.2.12: Given a universal theory, does there exist a notion of contraction in the scheme of all causal channels?
2. Can we find substitutes for the universal dilations of Chapter 2 so as to nail down more precisely what the causal-dilational ordering looks like, similar to what we did in Proposition 2.4.8?
3. In particular, how can we further the understanding of density theorems and rigidity in **CIT**, cf. Open Problem 4.4.26, Open Problem 4.4.27 and Open Problem 4.4.28?
4. Can we develop a metric version of the theory of causal channels and causal dilations, cf. the discussion in §4 of the prelude to this chapter?

Chapter 5

Rigidity and Quantum Self-Testing

§1. Introduction and Outline.

We concluded the preceding chapter by examining the causal-dilational ordering in **CIT** in a few special cases. The main lesson was that the dilational structure of Chapter 2 is fundamentally altered when causality is taken into account. In particular, the completeness principle no longer reigns globally, but rather becomes a volatile feature which depends on specifics of the causal channel in question. This phenomenon persists in the theory **QIT**, where it connects to the well-established field of quantum self-testing. That connection is the topic of this final chapter of the thesis.

(The reader is recommended at this point to review the material in the preliminary section of the thesis if it is unknown.)

Let us recall from the general introduction that *quantum self-testing* ([MY98, MY04, MYS12]) pertains to the following scenario:¹ Imagine that we interact with two separated computing devices, labelled by **A** and **B**. Device $i \in \{\mathbf{A}, \mathbf{B}\}$ accepts inputs from a finite set X_i and gives outputs from a finite set Y_i . By probing the devices many times, we may (under an i.i.d. assumption) establish the statistical input-output behaviour of the devices. These statistics are summarised by a collection of probability distributions $(P^x)_{x \in X}$ on the set of output pairs, $Y := Y_{\mathbf{A}} \times Y_{\mathbf{B}}$, indexed by the set of input pairs, $X := X_{\mathbf{A}} \times X_{\mathbf{B}}$. In general, the sets $X_{\mathbf{A}}, X_{\mathbf{B}}, Y_{\mathbf{A}}$ and $Y_{\mathbf{B}}$ are said to define a (*bipartite*) *Bell-scenario*, and a collection of probability distributions $(P^x)_{x \in X}$ is called a *behaviour*² for that Bell-scenario. Now assume that the devices work by sharing a bipartite quantum state $\rho \in \text{St}(\mathcal{H}_{\mathbf{A}} \otimes \mathcal{H}_{\mathbf{B}})$, and device i performing, on input $x_i \in X_i$, a projective quantum measurement to produce an outcome $y_i \in Y_i$, the measurement being described by the PVM $(\Pi_i^{x_i}(y_i))_{y_i \in Y_i}$. By the formalism of quantum theory, the input-output behaviour P is given by the Born rule, $P^{x_{\mathbf{A}}, x_{\mathbf{B}}}(y_{\mathbf{A}}, y_{\mathbf{B}}) = \text{tr}([\Pi_{\mathbf{A}}^{x_{\mathbf{A}}}(y_{\mathbf{A}}) \otimes \Pi_{\mathbf{B}}^{x_{\mathbf{B}}}(y_{\mathbf{B}})]\rho)$. We then ask the following:

Can the state and measurements be deduced from the input-output behaviour P ?

A configuration of state and measurements, $(\rho, \Pi_{\mathbf{A}}, \Pi_{\mathbf{B}})$, is called a (*tensor-product*³) *quantum strategy*. Clearly, there is always more than one quantum strategy which produces a

¹We restrict for simplicity to the bipartite case.

²Some authors use the term *correlation*. The term ‘behaviour’ is due to Cirelson ([Cir93]).

³There is in the infinite-dimensional case a more general notion of *quantum commuting strategies* ([DP16]). These are usually not considered in quantum self-testing, and neither shall we.

given input-output behaviour, since e.g. a local unitary rotation of state and measurements leaves the behaviour invariant. As such, taking the above question literally its answer is always in the negative. For some behaviours, however, it turns out that the strategy is ‘essentially unique’, in the sense that every strategy (ϱ, Π_A, Π_B) with behaviour P is ‘reducible’ to a fixed canonical strategy $(\tilde{\varrho}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, where the state $\tilde{\varrho} =: \tilde{\psi}$ is pure.⁴ This is the phenomenon of quantum self-testing.

To be precise, we say that (ϱ, Π_A, Π_B) is *reducible to* $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, if there exist Hilbert spaces $\mathcal{H}_A^{\text{res}}, \mathcal{H}_B^{\text{res}}$ (called *residual spaces*), isometries $W_i : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}}$, and a pure state ψ^{res} on $\mathcal{H}_A^{\text{res}} \otimes \mathcal{H}_B^{\text{res}} \otimes \mathcal{P}$ such that

$$(W\Pi^x(y) \otimes \mathbb{1}_{\mathcal{P}})|\psi\rangle = \tilde{\Pi}^x(y)|\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \quad \text{for all } x \in X, y \in Y, \quad (5.1)$$

where $W = W_A \otimes W_B$, $\Pi^x(y) = \Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B)$, $\tilde{\Pi}^x(y) = \tilde{\Pi}_A^{x_A}(y_A) \otimes \tilde{\Pi}_B^{x_B}(y_B)$ and where ψ is a purification of ϱ with purifying system \mathcal{P} .⁵ (This definition, which includes a purification of the state ϱ , is modelled on that of Refs. [ŠB19, RUV13].)

Intuitively, reducibility expresses that the strategy (ϱ, Π_A, Π_B) can be locally embedded into larger spaces and thereby realised as the strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, augmented by a *residual state* ψ^{res} which is left unmeasured. It is obvious that such a relation between strategies implies that their behaviours are the same – it is the converse which is interesting. One says that P *self-tests the strategy* $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ if every strategy (ϱ, Π_A, Π_B) with behaviour P is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$. Occasionally, one calls a behaviour P *rigid* if it self-tests some strategy, but we will refer to this as *rigidity a.t.t.l.* (*according to the literature*), so as to avoid confusion with rigidity in the sense of Definition 4.4.3 with which we ultimately want to link it.

Whereas the standard definition of self-testing reviewed above is mathematically unambiguous, its operational significance is unclear. The best way to expose this is by observing that is not at all obvious how one would formulate self-testing in a theory distinct from quantum theory – indeed, the reducibility condition is cast in a formalism specific to quantum theory, in terms of operators on Hilbert spaces. In particular, the projections $\Pi_i^{x_i}(y_i)$ are merely mathematical representations of certain measurement channels, guaranteed to exist by abstract theorems (as summarised in the preliminary section of the thesis); in fact, a general measurement channel is equivalent to a POVM, and whereas the reduction to PVMs is often justified in the literature on self-testing by reference to Naimark’s theorem, the meaning of this reduction remains somewhat obscure.⁶ Moreover, these unclarities are amplified when considering the fact that there could be ways for the two computing devices to locally establish an output on a given input which are more intricate than locally measuring a state;⁷ though it seems intuitively clear that we can standardise the form of more general ‘strategies’ to the form above, the meaning of the components deteriorates in that

⁴In fact, many treatments also take the state ϱ to be pure as well, as we did for simplicity in the general introduction. This, however, is unjustified as one should quantify universally over general strategies.

⁵It is easy to see that the choice of purification does not matter, since we may change the state ψ^{res} accordingly. The choices of vector representatives do not matter either, since we may absorb a potential phase into one of the isometries W_i .

⁶The reduction is often said to be ‘without loss of generality’ – whereas this is of course perfectly legitimate in cases where the objective is to prove a particular theorem using the self-testing phenomenon (as in e.g. Ref. [RUV13]), a statement of this kind is mathematically nonsensical in the absence of such a target theorem.

⁷For example, locally throwing a die and using the outcome to establish which of several shared states to use in a long sequence of various operations which ultimately make up an output, etc.

process.

In this chapter, we will see a fundamentally different way of looking at quantum self-testing. The idea is to consider the behaviour $P = (P^x)_{x \in X}$ not as a collection of probability distributions, but as a causal channel $\begin{array}{c} -\mathcal{X}_A- \\ \boxed{(P, \mathcal{C})} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \\ -\mathcal{Y}_B- \end{array}$ in **QIT**, with classical inputs on $\mathcal{X}_i := \mathbb{C}^{X_i}$ and classical outcomes on $\mathcal{Y}_i := \mathbb{C}^{Y_i}$, and where \mathcal{C} is the local causal specification that stems from the devices being separated. It is in fact easy to see that that the behaviours that can be realised by quantum strategies are precisely the Bell-channels in **QIT**.

The key is then to relate the causal dilations of this channel to the (tensor-product) quantum strategies as ordinarily perceived, and to relate the reducibility criterion (5.1) to the causal-dilational ordering. As such, self-testing will be recast operationally in terms of causally structured side-computations which the two computing devices may perform during our interaction with them.

*As in Section 4.4, we will assume throughout that the scheme **E** is that of constructible causal channels in **QIT**. In particular, all causal dilations will be constructible and derivability (\triangleright) will always refer to derivability using constructible channels.*

Density in Purifiable Theories. We start in Section 5.1 by making a few observations about the causal-dilational ordering in general purifiable theories (of which **QIT** is an example), in particular establishing a density theorem (Theorem 5.1.1) which slightly improves on Theorem 4.4.8 and Theorem 4.4.10. The density theorem says in the special case of **QIT** that *causal Stinespring dilations* form a dense class of dilations of a given Bell-channel. More precisely, in the case of a bipartite Bell-channel $\begin{array}{c} -\mathcal{X}_A- \\ \boxed{(P, \mathcal{C})} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \\ -\mathcal{Y}_B- \end{array}$, these are the causal dilations of the form

$$\begin{array}{c} -\mathcal{X}_A- \\ \boxed{\pi} \\ -\mathcal{X}_B- \end{array} \begin{array}{c} \boxed{\Sigma_A} \\ \text{---} \mathcal{E}_0 \text{---} \\ \boxed{\Sigma_B} \end{array} \begin{array}{c} -\mathcal{Y}_A- \\ \mathcal{E}_A \\ \mathcal{E}_B \\ -\mathcal{Y}_B- \end{array}, \quad (5.2)$$

with π a pure state and with Σ_A, Σ_B isometric quantum channels. It is (some of) those dilations which we will ultimately identify as representing the traditional quantum strategies.

In Section 5.1.A I will give an abstract recharacterisation of the channels (5.2), to the effect that an arbitrary isometric channel is of the form (5.2) precisely if it satisfies the correct non-signalling conditions (Theorem 5.1.2). This implies in particular that a multipartite channel is the behaviour of a (tensor-product) quantum strategy if and only if it admits a multipartite Stinespring dilation which is non-signalling among the different parties (Corollary 5.1.5). This result is quite surprising, since there are certainly channels which themselves are non-signalling without arising from quantum strategies, e.g. the PR box. As such, though the results presented in Section 5.1.A will not be put to use in the remainder of the chapter (in fact Section 5.1.A can be skipped without diminishing coherence in reading), they have been included because of their independent interest and connections to the work of Refs. [PR94, BGNP01, ESW02], which first contemplated the relationship between non-signalling and structural representations.

In Section 5.1.B we improve in purifiable theories on Theorem 4.4.10, giving a criterion (Lemma 5.1.8) for derivability among causal dilations in the dense class of dilations (5.2);

this will bring us closer to the language of quantum self-testing, in particular by emergence of the residual state ψ^{res} . The result also implies a partial ‘collapse’ of the causal-dilational ordering (Theorem 5.1.10), which in particular entails that the existence of a complete dilations with no acausal side-information is equivalent to all of the dilations (5.2) being \triangleright -equivalent. We will ultimately prove that this circumstance occurs for quantum self-testing, and this can be seen as justifying the intuition of some authors in the field who use the term ‘equivalence’ about the reducibility criterion (5.1), even though that criterion itself does not define an equivalence relation. (It also implies, however, that the condition (5.1) must be something slightly different from our notion of \triangleright -derivability, as detailed below.)

The Bridge to Quantum Self-Testing. In Section 5.2, we relate the framework proposed in this thesis to the standard framework of quantum self-testing. Specifically, the goal is to link rigidity and complete dilations of the causal behaviour channel (P, \mathcal{C}) to the standard notion of self-testing in terms of P . Basically, we face two challenges:

First, we must identify within the framework of causal dilations precisely what are the entities that correspond to quantum strategies for P in the usual sense. Then, we must rephrase in our language what the reducibility criterion (5.1) means in terms of these entities.

Strategies and Classically Bound Dilations. First, every ordinary quantum strategy (ϱ, Π_A, Π_B) will give rise to a causal dilation of the form (5.2); indeed, letting $\begin{array}{c} \text{---} \mathcal{C}^{X_i} \text{---} \\ \text{---} \mathcal{H}_i \text{---} \end{array} \boxed{\Lambda_i} \begin{array}{c} \text{---} \mathcal{C}^{Y_i} \text{---} \\ \text{---} \end{array}$ denote the ensemble of projective measurements corresponding to Π_i (cf. the definition in the preliminary section of the thesis), we obtain a dilation (5.2) by taking $\pi = \psi$ a purification of ϱ and $\Sigma_i = \hat{\Lambda}_i$ a Stinespring dilation of Λ_i . As it turns out, however, some dilations of the form (5.2) will not arise in this way from quantum strategies, and this is ultimately due to the fact that measurements can be dilated in peculiar ways, cf. Example 4.3.9. In the present context, those strange dilations reflect the fact (see Example 5.2.8) that it is mathematically possible to choose a state ϱ and channels Λ_A and Λ_B such that even though the channel

$$\begin{array}{c} \text{---} \mathcal{X}_A \text{---} \\ \text{---} \mathcal{H}_A \text{---} \\ \text{---} \mathcal{H}_B \text{---} \\ \text{---} \mathcal{X}_B \text{---} \end{array} \begin{array}{c} \boxed{\Lambda_A} \\ \boxed{\Lambda_B} \end{array} \begin{array}{c} \text{---} \mathcal{Y}_A \text{---} \\ \text{---} \mathcal{Y}_B \text{---} \end{array} \quad (5.3)$$

behaves as $\begin{array}{c} \text{---} \mathcal{X}_A \text{---} \\ \text{---} \mathcal{X}_B \text{---} \end{array} \boxed{(P, \mathcal{C})} \begin{array}{c} \text{---} \mathcal{Y}_A \text{---} \\ \text{---} \mathcal{Y}_B \text{---} \end{array}$, with classical inputs and outputs, the channels Λ_A and Λ_B are not themselves measurement channels, and thus the ‘strategy’ defined by the triple $(\varrho, \Lambda_A, \Lambda_B)$ has no counterpart in the ordinary framework of self-testing. We will resolve this basically by explicitly eliminating such ‘bad’ dilations, but providing the interpretation (Proposition 5.2.12) that it corresponds to restricting the class of imaginable dilations in a sensible way, namely to dilations which will be called *classically bound* (Definition 5.2.10).

Reducibility and Local Derivability. Secondly, as for the reducibility relation (5.1), there seems to be something fundamentally off with the direction of the relation, when compared to the philosophy that underlies the framework of causal dilations. Indeed, rigidity in terms of causal dilations means the existence of a \triangleright -largest dilation, whereas quantum self-testing would seem to assert the existence of a smallest (w.r.t. reducibility) strategy. The resolution to this paradox is offered by the collapse of the causal-dilational ordering as expressed by Corollary 5.1.12 mentioned above. More precisely, what we will show (Theorem 5.2.16) is that if $(\psi, \hat{\Lambda}_A, \hat{\Lambda}_B)$ defines a causal Stinespring dilation that corresponds to

the strategy (ϱ, Π_A, Π_B) and if $(\tilde{\psi}, \hat{\Lambda}_A, \hat{\Lambda}_B)$ defines one that corresponds to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, then the reducibility condition (5.1) holds if and only if there exist channels Γ_A and Γ_B such that

$$\begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \quad \uparrow \\
 \quad \psi \\
 \quad \downarrow \\
 \text{---}\mathcal{X}_B\text{---}
 \end{array}
 \begin{array}{c}
 \hat{\Lambda}_A \\
 \Gamma_A \\
 \text{tr} \\
 \hat{\Lambda}_B \\
 \Gamma_B
 \end{array}
 \begin{array}{c}
 \text{---}\mathcal{Y}_A\text{---} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 =
 \begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \quad \uparrow \\
 \quad \tilde{\psi} \\
 \quad \downarrow \\
 \text{---}\mathcal{X}_B\text{---}
 \end{array}
 \begin{array}{c}
 \hat{\Lambda}_A \\
 \hat{\Lambda}_B
 \end{array}
 \begin{array}{c}
 \text{---}\mathcal{Y}_A\text{---} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 \quad (5.4)$$

(The key to proving this is a technical result, Lemma 5.2.24.) We may call the above relation *local derivability*, since it asserts the \succeq -derivability of one dilation by another using a product of three channels in the environment. (Though this condition is somewhat theory-independent, it is not purely operational, since causal Stinespring dilations do not seem to admit an operational definition.) Using the above-mentioned collapse-corollary, however, we will show (Corollary 5.2.20) that this condition can be rephrased as rigidity of (P, \mathcal{C}) relative to classically bound dilations witnessed by a complete dilation with no causal side-information, along with the existence of a simple representative in the \succeq -equivalence class of causal Stinespring dilations, namely the one corresponding to the canonical strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$. Conjecturing that such a simple representative always exists (Conjecture 5.2.22) suggests a fully operational view on quantum self-testing.

Some Perks of a Reformulation. The recasting of quantum self-testing in the language of causal dilations not only points towards generalisations to other theories, but also loosens the original theory from linear operators and thus facilitates a new way of understanding some well-known consequences of quantum self-testing.

For example, it is easy to prove that self-testing implies the production of genuine randomness in a Bell-experiment (Proposition 5.3.2) and that extremality of a behaviour is a necessary condition for quantum self-testing (Proposition 5.3.4). The former fact connects to the work of Ref. [FFW11], and the latter (though seemingly common knowledge for many years) was first proven formally in Ref. [GKW+18].

It will also be clear from the reformulation why the canonical state and the canonical measurements can always be locally extracted from the state and measurements of arbitrary strategies. Extractibility of the state can be proven in a few lines using the original formulation, and is also easily proved in the reformulation using local derivability, though by a completely different argument (Proposition 5.3.6). Extractibility of the measurements does not seem to have been proved before in the formulation presented here (Proposition 5.3.7), but relates to some authors' alternative phrasing of self-testing (see e.g. Ref. [Kan17] and the discussion in Ref. [SB19]).

§2. Further Notes on the Existing Literature.

Quantum Self-Testing as Ordinarily Perceived. The concept of quantum self-testing is widely recognised as being introduced in Refs. [MY98] and [MY04] by D. Mayers and A. Yao, with Ref. [MY04] establishing much of the current terminology. However, Ref. [Eke91] by Artur Ekert contains already on an informal level some core ideas, in particular pointing towards the field of *device-independent cryptography*, and the mathematical results that some behaviours can be achieved by essentially unique quantum strategies were inde-

pendently reported already in Refs. [SW87] and [PR92], though with a foundational rather than cryptographic perspective.

The general mathematical definition of self-testing took its modern standard form in Ref. [MYS12], and it is the one employed in the vast majority of contemporary expositions on the subject. There are authors who, when proving rigidity results, seek alternative formulations, but to the best of my knowledge all such are either designed for specialised situations (e.g. [RUV13], Def. 5.5), or, though generally applicable, capture only certain aspects of the commonly used definition (e.g. [Kan17]). As such, a general operational definition of quantum self-testing has not been attempted before.

§3. Contributions.

The original contributions of this chapter are the following:

1. Proving a density theorem for Bell-channels in purifiable theories (Theorem 5.1.1) and providing an abstract characterisation of elements in the dense class, implying in particular a surprising new characterisation of tensor-product quantum behaviours in terms of non-signalling properties of their Stinespring dilations (Corollary 5.1.5).
2. Establishing a formal connection between the conventional definition of quantum self-testing ([MY98, MY04, MYS12, ŠB19]) and rigidity in the sense of Chapter 4 (Theorem 5.2.16, Corollary 5.2.20 and Conjecture 5.2.22).
3. Giving simple proofs for known implications of quantum self-testing, including the necessity of extremality of the behaviour (Proposition 5.3.2 and Proposition 5.3.4), and local extractibility of the canonical state (Proposition 5.3.6) and the canonical measurements (Proposition 5.3.7).

The contribution mentioned in 2. is in a sense the most important, as it was the original motivation for the thesis project itself. It points towards a purely operational interpretation of quantum self-testing, entailing in particular that

- the origin of *quantum strategies* (Definition 5.2.1) is explained, including the validity of restricting to projective measurements;
- the nature of the reducibility relation (Definition 5.2.3) is clarified, by exhibiting it as a strong version of derivability in the environment;
- the notions of reducibility and self-testing are put into a context of general theories;
- quantum self-testing is seen to imply rigidity in the sense of causal dilations, a fully operational notion, and a converse to this implication is conjectured.

§4. The Metric Aspect.

We will not touch at the concept of *robust* self-testing, an approximate version of the self-testing story in which the exact equality of behaviours and exact fulfilment of the reducibility condition are relaxed ([MYS12]). It seems quite obvious that an understanding of robustness in the framework of causal dilations hinges upon a metric version of Chapter 4, whose challenges have already been discussed.

5.1 Density Considerations in Purifiable Theories

Suppose that $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{(T, \mathcal{C})} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ is a bipartite Bell-channel in a purifiable theory Θ where every channel has a pure one-sided dilation (e.g. by virtue of Θ being universal, cf. Proposition 2.5.10). We have the following improvement of the general density theorem (Theorem 4.4.8), which as usual generalises without effort to the multipartite (or unipartite) case:

Theorem 5.1.1. (*Density Theorem for Bell-Channels in Purifiable Theories.*)
The causal dilations of the form

$$\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \boxed{\Sigma_A} \\ \boxed{\pi} \\ \boxed{\Sigma_B} \end{array} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{E}_A\text{---} \\ \text{---}\mathcal{E}_0\text{---} \\ \text{---}\mathcal{E}_B\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}, \quad (5.5)$$

where Σ_A, Σ_B and π are dilationally pure, constitute a dense class for the constructible

dilations of (T, \mathcal{C}) . Moreover, $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \boxed{\Sigma_A} \\ \boxed{\pi} \\ \boxed{\Sigma_B} \end{array} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{E}_A\text{---} \\ \text{---}\mathcal{E}_0\text{---} \\ \text{---}\mathcal{E}_B\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array} \supseteq \begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \boxed{\Sigma'_A} \\ \boxed{\pi'} \\ \boxed{\Sigma'_B} \end{array} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{E}'_A\text{---} \\ \text{---}\mathcal{E}'_0\text{---} \\ \text{---}\mathcal{E}'_B\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ if and only if there exist channels Γ_A, Γ_B and a dilationally pure channel Φ such that

$$\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \boxed{\Sigma_A} \\ \boxed{\pi} \\ \boxed{\Sigma_B} \end{array} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{E}_A\text{---} \\ \text{---}\mathcal{E}_0\text{---} \\ \text{---}\mathcal{E}_B\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array} \begin{array}{c} \boxed{\Gamma_A} \\ \boxed{\Phi} \\ \boxed{\Gamma_B} \end{array} \begin{array}{c} \text{---}\mathcal{E}'_A\text{---} \\ \text{---}\mathcal{E}'_0\text{---} \\ \text{---}\mathcal{E}'_B\text{---} \end{array} = \begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_A\text{---} \end{array} \begin{array}{c} \boxed{\Sigma'_A} \\ \boxed{\pi'} \\ \boxed{\Sigma'_B} \end{array} \begin{array}{c} \text{---}\mathcal{Y}_B\text{---} \\ \text{---}\mathcal{E}'_A\text{---} \\ \text{---}\mathcal{E}'_0\text{---} \\ \text{---}\mathcal{E}'_B\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array} \quad (5.6)$$

Proof. The density statements follow from Theorem 4.4.8, simply by additionally forming for each component a pure dilation, thus moving further up in the \supseteq -ordering. As for the characterisation of the derivability relation \supseteq , that follows directly from the statements in Theorem 4.4.10, observing that Φ can be taken dilationally pure by choosing a pure dilation and including a trash in one of the channels Γ_i if necessary. \square

As mentioned in the introduction, a (*tensor-product*) *quantum strategy* will be encoded in the isometric channel (5.5) in the theory **QIT**. This isometric channel will reveal all relevant information about the strategy with regards to self-testing. We will call it a *causal Stinespring dilation* of (T, \mathcal{C}) . In a general purifiable theory, the channels (5.5) are pure if the theory is localisable and we may thus more generally call them *pure causal dilations*.

As was the case in **CIT** (cf. Remark 4.4.16), the channels of the dense class are complete dilations of T if we disregard causality.⁸ Hence, the dilations (5.5) are all equivalent by

⁸This is no coincidence; in a complete and localisable theory, Theorem 4.4.8 can always be such improved, by picking complete dilations of each component.

means of a channel acting on the inaccessible interface; the intricacy arises because for two channels in the dense class to be related in the causal-dilational order, they must be related by a channel with a particular structure as prescribed by Eq. (5.6).

5.1.A An Abstract Characterisation of Pure Causal Dilations

In Chapter 2, we saw that many theories have the DiVincenzo Property, which allows one to conclude that a channel $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{T} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ has the structure $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{T_1} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ merely provided that it satisfies the necessary non-signalling condition – in the case of **QIT**, this statement was raised and illuminated in Ref. [BGNP01], and then fully proved in Ref. [ESW02].

On the other hand, it has been known since the work of Ref. [PR94] that more intricate structures generally can not be concluded based on fulfilment of non-signalling conditions; for example, not every channel $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{T} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ satisfying the same non-signalling conditions as the channel

$$\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \boxed{s} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \boxed{T_A} \\ \boxed{T_B} \end{array} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array} \quad (5.7)$$

(i.e. compatible with the local causal specification \mathcal{C} given by $\mathcal{C}(y_i) = \{x_i\}$) is of that form. In the language of Chapter 4, not every causal channel with local specification is constructible.

What we will prove in the present subsection is that, under mild dilational assumptions, every pure channel which satisfies the necessary non-signalling conditions is of the form (5.7). In fact, we have the following:

Theorem 5.1.2. (*Structure from Non-Signalling.*)

Suppose that Θ is localisable, and that there is a state on every system in Θ . Then, a dilationally pure channel

$$\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \boxed{T} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \text{---}\mathbb{Y}_A\text{---} \\ \text{---}\mathbb{Z}\text{---} \\ \text{---}\mathbb{Y}_B\text{---} \end{array} \quad (5.8)$$

is of the form

$$\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \boxed{s} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \begin{array}{c} \boxed{T_A} \\ \boxed{T_B} \end{array} \begin{array}{c} \text{---}\mathbb{Y}_A\text{---} \\ \text{---}\mathbb{Z}\text{---} \\ \text{---}\mathbb{Y}_B\text{---} \end{array} \quad (5.9)$$

if and only if it is compatible with the causal specification \mathcal{C} given by $\mathcal{C}(\text{ports}(\mathbb{Z})) = \emptyset$, $\mathcal{C}(\text{ports}(\mathbb{Y}_A)) = \{x_A\}$ and $\mathcal{C}(\text{ports}(\mathbb{Y}_B)) = \{x_B\}$. If Θ is moreover universal and purifiable, then the components T_A , T_B and s may be taken pure in that case.

Remark 5.1.3. (More Parties.)

The statement and proof generalise to the arbitrary multipartite setting by induction. In fact, the proof can be seen as the induction step executed in going from the case of unipartite to bipartite Bell-scenarios. \spadesuit

Remark 5.1.4. (Approximate Generalisation.)

If D is a dilational metric on Θ (cf. Definition 3.3.1), it is easy to see from the proof that the first statement in Theorem 5.1.2 admits an approximate generalisation, such that if T is ε -close to being suitably non-signalling, then T is 3ε -close to a channel of the form (5.9). \spadesuit

Proof. The ‘only if’-direction is clear. As for the ‘if’-direction, first observe that, since Θ is spatially localisable, Θ has the DiVincenzo Property (Proposition 2.3.20). Since T is compatible with \mathcal{C} , the non-signalling property implied by the condition $\mathcal{C}(\text{ports}(\mathbb{Z}) \cup \{y_B\}) = \{x_B\}$ therefore yields that T is of the form

$$\begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \begin{array}{|c|} \hline S_B \\ \hline \end{array} \\
 \begin{array}{|c|} \hline T_A \\ \hline \end{array} \\
 \text{---}\mathbb{Y}_A\text{---} \\
 \text{---}\mathbb{Z}\text{---} \\
 \text{---}\mathbb{Y}_B\text{---}
 \end{array}
 \tag{5.10}$$

for some channels S_B and T_A . If we can show that S_B is of the form

$$\begin{array}{c}
 \text{---} \\
 \begin{array}{|c|} \hline s \\ \hline \end{array} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \begin{array}{|c|} \hline T_B \\ \hline \end{array} \\
 \text{---}\mathbb{Y}_B\text{---} \\
 \text{---}\mathbb{Z}\text{---}
 \end{array}
 , \text{ then}$$

we are done proving the first statement in the theorem. To this end, however, observe that because of the non-signalling property implied by the condition $\mathcal{C}(\text{ports}(\mathbb{Z}) \cup \{y_A\}) = \{x_A\}$, the channel T is also of the form

$$\begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \begin{array}{|c|} \hline S'_A \\ \hline \end{array} \\
 \begin{array}{|c|} \hline T'_B \\ \hline \end{array} \\
 \text{---}\mathbb{Y}_A\text{---} \\
 \text{---}\mathbb{Z}\text{---} \\
 \text{---}\mathbb{Y}_B\text{---}
 \end{array}
 \tag{5.11}$$

for some channels S'_A and T'_B . In general, the equality of the two forms (5.10) and (5.11) of T is all that can be concluded on the basis of compatibility with the specification \mathcal{C} . If we could somehow reverse T_A and move it to the other side of this equality, effectively isolating S_B , we would have the desired form of S_B , but this is not possible in general. *However*, since T is dilationally pure, we can squeeze out this opportunity: If we pick (by Proposition 2.3.19) a reversible dilation R_A of T_A , then dilational purity of T implies that

$$\begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \begin{array}{|c|} \hline S_B \\ \hline \end{array} \\
 \begin{array}{|c|} \hline R_A \\ \hline \end{array} \\
 \text{---}\mathbb{Y}_A\text{---} \\
 \text{---}\mathbb{Z}\text{---} \\
 \text{---}\mathbb{Y}_B\text{---}
 \end{array}
 =
 \begin{array}{c}
 \begin{array}{|c|} \hline t \\ \hline \end{array} \\
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \begin{array}{|c|} \hline T \\ \hline \end{array} \\
 \text{---}\mathbb{Y}_A\text{---} \\
 \text{---}\mathbb{Z}\text{---} \\
 \text{---}\mathbb{Y}_B\text{---}
 \end{array}
 =
 \begin{array}{c}
 \begin{array}{|c|} \hline t \\ \hline \end{array} \\
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \begin{array}{|c|} \hline S'_A \\ \hline \end{array} \\
 \begin{array}{|c|} \hline T'_B \\ \hline \end{array} \\
 \text{---}\mathbb{Y}_A\text{---} \\
 \text{---}\mathbb{Z}\text{---} \\
 \text{---}\mathbb{Y}_B\text{---}
 \end{array}
 \tag{5.12}$$

for some state t , and applying an inverse⁹ R_A^- to R_A on both sides yields the identity

$$\begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \boxed{S_B} \\
 \text{---}\mathcal{Z}\text{---} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 =
 \begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \text{---}\mathcal{X}_B\text{---} \\
 \boxed{S'_A} \\
 \text{---}\mathcal{Z}\text{---} \\
 \boxed{T'_B} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 \cdot
 \begin{array}{c}
 \boxed{t} \\
 \text{---}\mathcal{X}_A\text{---} \\
 \boxed{R_A^-} \\
 \text{---}\mathcal{Z}\text{---} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 . \quad (5.13)$$

Now, by finally inserting an arbitrary state as input to \mathcal{X}_A and trashing it, we obtain

$$\begin{array}{c}
 \text{---}\mathcal{X}_A\text{---} \\
 \boxed{S_B} \\
 \text{---}\mathcal{Z}\text{---} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 =
 \begin{array}{c}
 \boxed{s} \\
 \text{---}\mathcal{Z}\text{---} \\
 \boxed{T'_B} \\
 \text{---}\mathcal{Y}_B\text{---}
 \end{array}
 \quad (5.14)$$

for some state s . Letting $T_B = T'_B$, we conclude that T is of the form Eq. (5.9), as desired.

To prove the second statement in the theorem, first observe that we may without loss of generality assume that S_B in Eq. (5.10) is dilationally pure. Indeed, by redefining T_A to include a trash, we may replace S_B with a universal dilation of S_B , and by Proposition 2.5.10 this dilation is dilationally pure. By the same argument, the state s in (5.14) may be taken

to be pure and a universal dilation of $\begin{array}{c} \boxed{s} \\ \text{---}\mathcal{Z}\text{---} \\ \boxed{\text{tr}} \end{array}$ without loss of generality. Then, however,

Eq. (5.14) implies by purity of S_B and universality of the dilation s that $T_B (= T'_B)$ is dilationally pure. Inserting (5.14) into Eq. (5.10), and once again replacing s by a universal

dilation $\begin{array}{c} \boxed{s'} \\ \text{---}\mathcal{Z}\text{---} \\ \boxed{\text{tr}} \end{array}$ of $\begin{array}{c} \boxed{s} \\ \text{---}\mathcal{Z}\text{---} \\ \boxed{\text{tr}} \end{array}$ then similarly implies (since $\begin{array}{c} \boxed{s'} \\ \text{---}\mathcal{Z}\text{---} \\ \boxed{T_B} \end{array}$ is a universal dilation, T_B being pure and reversible) that the modified T_A is dilationally pure. □

Theorem 5.1.2 immediately implies a recharacterisation in terms of non-signalling of the channels in the dense class of Theorem 5.1.1. In a purifiable theory such as **QIT**, however, it also directly yields a surprising recharacterisation of the entire class of Bell-channels in terms of non-signalling in their pure dilations:

Given an arbitrary bipartite channel $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \\ \boxed{T} \\ \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$, let us say that T is *purely non-signalling* if T has a dilationally pure bipartite dilation $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \\ \boxed{\Sigma} \\ \text{---}\mathcal{E}_A\text{---} \\ \text{---}\mathcal{E}_B\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ which is compatible with the specification \mathcal{C} given by $\mathcal{C}(\text{ports}(\mathbb{E}_i) \cup \{y_i\}) = \{x_i\}$ for $i = A, B$. This requirement does not imply that every pure bipartite dilation will be non-signalling, but merely says there


⁹We might pick R_A to be dilationally pure, and for this reason reversible; in **QIT**, this corresponds to picking R_A isometric, i.e. of ‘going to the church of the larger Hilbert space’. As such, the following joke may help to remember the idea behind the proof of Theorem 5.1.2: *Why is the church of the larger Hilbert space better than the church of Catholicism? Because, while in the Catholic church your sins can be forgiven, in the church of the larger Hilbert space your sins can be undone.*

exist ones which are.¹⁰ In the specific theory **QIT**, the condition is a matter of whether there exists a bipartite Stinespring dilation which is non-signalling, i.e. whether T can be implemented in a non-signalling way ‘in the Church of the larger Hilbert space’.

We have the following:

Corollary 5.1.5. (*Purely Non-Signalling means Bell.*)

Let Θ be a universal, localisable and purifiable theory. Then, an arbitrary bipartite channel



is purely non-signalling if and only if it is a Bell-channel, i.e. of the form

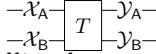
(5.15)

Remark 5.1.6. Again, the generalisation to the multipartite case is straightforward. \spadesuit

Proof. First recall that every channel in Θ has a one-sided pure dilation, by the last statement in Proposition 2.5.10.

The ‘if’-direction is not difficult (the reader is encouraged to think about the special case of **QIT**, where pure one-sided dilations mean Stinespring dilations): If T is a Bell-channel, we may take it to be of the form (5.15) with s pure. By dilating T_A and T_B to pure channels, we thus obtain a pure (by localisability) bipartite dilation of T which evidently has the required non-signalling properties, so T is purely non-signalling.

The ‘only if’-direction is the surprise: Given a purely non-signalling channel T , any pure bipartite dilation which witnesses this must by Theorem 5.1.2 (with $\mathbb{Z} = \mathbb{I}$) factor as a Bell-channel; but then, trashing the hidden interfaces gives the Bell-structure (5.15) for T itself. \square

Corollary 5.1.7. In **QIT**, there exist non-signalling channels  which are not purely non-signalling. For example, the PR box is non-signalling but not purely non-signalling.

Corollary 5.1.5 gives a novel characterisation of the tensor-product quantum behaviours in **QIT**, i.e. those that arise as Bell-channels. It is unexpected because, a priori, there is no reason to think that ‘non-signalling in the Church’ should favour tensor-product behaviours over more general behaviours. By previous observations, Corollary 5.1.5 applies not only to the theory **QIT** but also the theory **QIT**[∞]. In that theory, a slightly more general model for the implementable behaviours is that arising from so-called *commuting-operator strategies* ([DP16]); however, Corollary 5.1.5 means that the behaviours from this class which are not already tensor-product behaviours do not admit non-signalling Stinespring dilations. (The approximate generalisation of Corollary 5.1.5 even implies, in light of the recently established fact that there exist commuting-operator behaviours which are not even close to being tensor-product behaviours ([JNV+20]), that some commuting-operator behaviours are not even close to having non-signalling Stinespring dilations.)

It is also worth observing (still in **QIT**[∞]) that if the systems \mathcal{X}_i and \mathcal{Y}_i are finite-dimensional, then a slightly more detailed book-keeping in the proof of Corollary 5.1.5 will

¹⁰In general, there will be ones which are not non-signalling, since given any one that is non-signalling we can swap \mathbb{E}_A with \mathbb{E}_B and generically obtain one that is not.

reveal that the bipartite system $\mathcal{H}_A \otimes \mathcal{H}_B$ corresponding to the state s can be taken finite-dimensional if and only if the hidden systems in the non-signalling Stinespring dilation can be taken finite-dimensional. Since there are Bell-channels in \mathbf{QIT}^∞ which cannot be realised using a state on a finite-dimensional system, this means that some bipartite channels $\begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{T} \begin{array}{c} -\mathcal{Y}_A- \\ -\mathcal{Y}_B- \end{array}$ in \mathbf{QIT}^∞ have a non-signalling bipartite Stinespring dilation with infinite-dimensional hidden systems, but no one with finite-dimensional hidden systems.

5.1.B A Partial Collapse of the Causal-Dilational Ordering

It is of interest to better understand the relation (5.6) of derivability between the dilations of Theorem 5.1.1. As it turns out, purifiability of a theory entails a recharacterisation of this relation:

Lemma 5.1.8. (*Recharacterisation of \triangleright among Pure Causal Dilations.*)

Suppose that Θ is purifiable and localisable, and let $\begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{(T, \mathcal{E})} \begin{array}{c} -\mathcal{Y}_A- \\ -\mathcal{Y}_B- \end{array}$ be a bipartite Bell-channel in Θ . Then, two pure causal dilations satisfy

if and only if there exist dilationally pure channels $\hat{\Gamma}_A, \hat{\Gamma}_B, \Phi$ and a dilationally pure state π^{res} , such that

$$\begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{\Sigma_A} \begin{array}{c} -\mathcal{Y}_A- \\ -\mathcal{Y}_B- \end{array} \begin{array}{c} \mathcal{E}_A \\ \mathcal{E}_B \end{array} \begin{array}{c} \hat{\Gamma}_A \\ \hat{\Gamma}_B \end{array} \begin{array}{c} \mathcal{E}'_A \\ \mathcal{E}'_B \end{array} \\ \begin{array}{c} \pi \\ \pi \end{array} \begin{array}{c} \mathcal{E}_0 \\ \mathcal{E}_0 \end{array} \begin{array}{c} \Phi \\ \Phi \end{array} \begin{array}{c} \mathcal{E}'_0 \\ \mathcal{E}'_0 \end{array} \\ \begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{\Sigma_B} \begin{array}{c} -\mathcal{Y}_A- \\ -\mathcal{Y}_B- \end{array} \begin{array}{c} \mathcal{E}_A \\ \mathcal{E}_B \end{array} \begin{array}{c} \hat{\Gamma}_A \\ \hat{\Gamma}_B \end{array} \begin{array}{c} \mathcal{E}'_A \\ \mathcal{E}'_B \end{array} \\ \begin{array}{c} \pi \\ \pi \end{array} \begin{array}{c} \mathcal{E}_0 \\ \mathcal{E}_0 \end{array} \begin{array}{c} \Phi \\ \Phi \end{array} \begin{array}{c} \mathcal{E}'_0 \\ \mathcal{E}'_0 \end{array} \\ \begin{array}{c} -\mathcal{X}_A- \\ -\mathcal{X}_B- \end{array} \boxed{\Sigma'_A} \begin{array}{c} -\mathcal{Y}_A- \\ -\mathcal{Y}_B- \end{array} \begin{array}{c} \mathcal{E}'_A \\ \mathcal{E}'_B \end{array} \begin{array}{c} \text{id} \\ \text{id} \end{array} \begin{array}{c} \mathcal{E}'_A \\ \mathcal{E}'_B \end{array} \\ \begin{array}{c} \pi' \\ \pi' \end{array} \begin{array}{c} \mathcal{E}'_0 \\ \mathcal{E}'_0 \end{array} \begin{array}{c} \pi^{\text{res}} \\ \pi^{\text{res}} \end{array} \begin{array}{c} \text{id} \\ \text{id} \end{array} \begin{array}{c} \mathcal{E}'_A \\ \mathcal{E}'_B \end{array} \end{array} \quad (5.16)$$

where each channel is given its primitive specification (the identities thus stalling the outputs), and where some of the wires of the hidden interfaces are drawn dotted and unlabelled for the sake of clearer perception.

The channel Φ may be taken identical to that of Eq. (5.6).

Remark 5.1.9. (Interpretation.)


The main difference between the above condition and the original condition (5.6) is that all involved channels above are dilationally pure. In \mathbf{QIT} , one would say that (5.16) is the ‘purified’ version of (5.6), or that it corresponds to viewing the condition (5.6) ‘in the Church of the larger Hilbert space’. The surprise of Lemma 5.1.8 is that even in this purified view the condition is rather simple, with no shenanigans playing out on the right hand side – only a harmless *residual* state π^{res} is adjoined. \spadesuit

Proof. It is clear that (5.16) implies (5.6), simply by trashing all systems corresponding to dotted wires. Conversely, if (5.6) holds as an equation between channels, then by forming dilationally pure dilations $\hat{\Gamma}_A$ of Γ_A and $\hat{\Gamma}_B$ of Γ_B , dilational purity of the right hand side of (5.6) implies the identity (5.16) for some state π^{res} . That π^{res} must in fact be dilationally pure follows from dilational purity of the left hand side of (5.16).


(We use in this proof that compositions of pure channels are pure; this hinges on localisability.) \square

Theorem 5.1.10. (*Partial Collapse.*)

A pure causal dilation of (T, \mathcal{C}) is derivable from a given pure causal dilation




if and only if it is \triangleright -equivalent to a dilation of the form




for some dilationally pure channel Φ , where each channel is given its primitive specification (the identities thus stalling the \mathcal{G}_i -outputs). In fact, the channel Φ can be taken to be the one that derives it according to condition (5.6).

Remark 5.1.11. The significance of this theorem is that within the class of pure causal dilations, the pre-order \triangleright implodes (up to equivalence) to simple modifications which redistribute the acausal side-information of \mathcal{E}_0 to acausal side-information in \mathcal{E}'_0 along with side-information in \mathcal{G}_A and \mathcal{G}_B , by means of the pure channel Φ . \spadesuit


Proof. It is clear that the dilation (5.17) is derivable from the dilation



any \triangleright -equivalent dilation must be as well. Conversely, we find for any derivable dilation

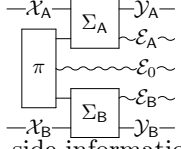


by Lemma 5.1.8 a pure state π^{res} and pure channels $\hat{\Gamma}_A, \hat{\Gamma}_B$ and Φ such that (5.16) holds. As such, it is also derivable from the dilation (5.17). But it is in fact \triangleright -equivalent to it: By Theorem 2.5.11, the pure channels $\hat{\Gamma}_A$ and $\hat{\Gamma}_B$ are reversible, so by applying left-inverses $\tilde{\Gamma}_A$ and $\tilde{\Gamma}_B$, with their primitive specifications, we obtain the identity

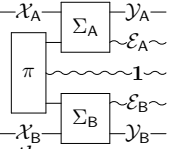


showing that the converse derivability holds as well, as desired. (Plainly speaking, we have managed to transfer to the other side of the equality symbol the circuitry in the environment that derives one dilation from the other.) \square

Recall from Example 4.3.8 the notion of a dilation having *no acausal side-information*.

If a pure causal dilation  has $\mathcal{E}_0 = \mathbf{1}$, then it is clearly \succeq -equivalent to a dilation with no acausal side-information, so we might as well use the term about that dilation itself. In this case (i.e. when $\mathcal{E}_0 = \mathbf{1}$), the dilation (5.17) is \succeq -equivalent to it for any Φ , and this gives us the following two corollaries to Theorem 5.1.10, which will be important in the context of quantum self-testing:

Corollary 5.1.12. (*Derivability Collapses to Equivalence.*)

If  is a pure causal dilation of (T, \mathcal{C}) with no acausal side-information, then the pure causal dilations which are derivable from it are precisely those which are \succeq -equivalent to it.

Proof. Obvious from Theorem 5.1.10. □

Corollary 5.1.13. *The following are equivalent:*

1. (T, \mathcal{C}) has a complete dilation with no acausal side-information.
2. All pure causal dilations of (T, \mathcal{C}) are \succeq -equivalent.

Moreover, in this case any pure causal dilation is complete.

Proof. If 2. holds, then the dense class collapses to a single level; hence, any pure causal dilation is a complete dilation of (T, \mathcal{C}) . Clearly, we find among the pure causal dilations one with $\mathcal{E}_0 = \mathbf{1}$, simply by picking anyone and applying some Φ which merges \mathcal{E}_0 with \mathcal{E}_A .

If 1. holds, then any given complete dilation with no acausal side-information can be derived from a pure complete dilation, by density of pure dilations. We may moreover assume without loss of generality that this pure dilation has $\mathcal{E}_0 = \mathbf{1}$. In other words, (T, \mathcal{C}) has a pure causal dilation which is complete and has no acausal side-information. In particular, it must be possible to derive any pure causal dilation from it, so Corollary 5.1.12 implies that 2. holds. □

What we now aim to get at, is that when $(T, \mathcal{C}) = (P, \mathcal{C})$ encodes the behaviour in a Bell-scenario, quantum self-testing is essentially equivalent to the two conditions of Corollary 5.1.13. There are, however, two subtleties related to this statement.

First, it turns out that we have to consider rigidity relative to a sub-class of possible dilations, because some dilations of (P, \mathcal{C}) will not be reflected in the space of ordinary quantum strategies; this is related to the peculiarities of dilating measurements, cf. Example 4.3.9.

Secondly, in the usual formulation of self-testing a specific strategy \tilde{S} is singled out as canonical, and the relations that other strategies S bear to it is not a symmetric one, and thus cannot possibly be \succeq -equivalence; what is actually going on is that \tilde{S} is smallest in the \succeq -equivalence class, w.r.t. a pre-order (namely, reducibility) which refines the pre-order \succeq .

5.2 Rigidity in QIT – Quantum Self-Testing

Consider a bipartite Bell-channel $\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \end{array} \begin{array}{c} \boxed{(P, \mathcal{C})} \\ \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ in **QIT**. Suppose that the systems \mathcal{X}_i and \mathcal{Y}_i are embeddings of classical systems, $\mathcal{X}_i = \mathbb{C}^{X_i}$ and $\mathcal{Y}_i = \mathbb{C}^{Y_i}$, for finite sets X_i and Y_i . Suppose moreover that the channel P is classical on its interfaces, i.e. that

$$\begin{array}{c} \Delta_{X_A} \\ \Delta_{X_B} \end{array} \begin{array}{c} \boxed{(P, \mathcal{C})} \\ \Delta_{Y_A} \\ \Delta_{Y_B} \end{array} = \begin{array}{c} \boxed{(P, \mathcal{C})} \end{array}, \quad (5.19)$$

where Δ_Z is the decoherence channel on \mathbb{C}^Z given by $\Delta_Z(A) = \sum_{z \in Z} |z\rangle\langle z| A |z\rangle\langle z|$. As discussed earlier, P can then be thought of as a channel in **CIT**, determined by the states

$$P^{(x_A, x_B)} := P(|x_A\rangle\langle x_A| \otimes |x_B\rangle\langle x_B|), \quad x_A \in X_A, x_B \in X_B, \quad (5.20)$$

which are classical states on $\mathbb{C}^{Y_A} \otimes \mathbb{C}^{Y_B} \cong \mathbb{C}^{Y_A \times Y_B}$, that is, probability distributions on $Y_A \times Y_B$. In the traditional language of self-testing, we are given a *behaviour* or *correlation* $(P^{(x_A, x_B)})_{x_A \in X_A, x_B \in X_B}$ for the bipartite Bell-scenario with input sets X_A, X_B and output sets Y_A, Y_B . Let us denote by $X := X_A \times X_B$ and $Y := Y_A \times Y_B$ the total input and output sets, and let us generically write x and y for the tuples (x_A, x_B) and (y_A, y_B) , respectively.

By Theorem 5.1.1, the isometric channels

$$\begin{array}{c} \mathcal{X}_A \\ \mathcal{H}_A \\ \mathcal{H}_B \\ \mathcal{X}_B \end{array} \begin{array}{c} \boxed{\Sigma_A} \\ \boxed{\Sigma_B} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{E}_A \\ \mathcal{E}_B \\ \mathcal{Y}_B \end{array} \quad (5.21)$$

form a dense class for constructible causal dilations of (P, \mathcal{C}) . The goal is now to build our way from the channels (5.21) to the traditional (tensor-product) quantum strategies which realise the input-output behaviour $(P^x)_{x \in X}$, and to connect rigidity of the causal channel (P, \mathcal{C}) to the traditional definition of quantum self-testing.

5.2.A The Standard Definition of Quantum Self-Testing

Let us begin by recalling the traditional notions from the literature (recall in particular that measurements are in this traditional conception taken to be projective):

Definition 5.2.1. (Quantum Strategies.)

A (*finite-dimensional tensor-product*) *quantum strategy* is a triple (ϱ, Π_A, Π_B) , where ϱ is a state on some bipartite finite-dimensional system $\mathcal{H}_A \otimes \mathcal{H}_B$, and where, for $i = A, B$, $\Pi_i = (\Pi_i^{x_i})_{x_i \in X_i}$ is a collection of PVMs on \mathcal{H}_i , that is, families $\Pi_i^{x_i} = (\Pi_i^{x_i}(y_i))_{y_i \in Y_i}$ of orthogonal projections on \mathcal{H}_i with $\sum_{y_i \in Y_i} \Pi_i^{x_i}(y_i) = \mathbf{1}_{\mathcal{H}_i}$. ■

Given a quantum strategy (ϱ, Π_A, Π_B) , we will generically denote the system $\mathcal{H}_A \otimes \mathcal{H}_B$ by \mathcal{H} , and for $x = (x_A, x_B)$ and $y = (y_A, y_B)$ the tensor-product projection $\Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B)$ by $\Pi^x(y)$.

Definition 5.2.2. (Behaviour of a Quantum Strategy.)

The behaviour of the quantum strategy (ϱ, Π_A, Π_B) is the behaviour $P = (P^x)_{x \in X}$ given by

$$P^x(y) = \text{tr}(\Pi^x(y)\varrho) = \text{tr}([\Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B)]\varrho) \quad (5.22)$$

for $x = (x_A, x_B) \in X$ and $y = (y_A, y_B) \in Y$. ■

Observe that the collection $\Pi_i = (\Pi_i^{x_i})_{x_i \in X_i}$ of projective measurements can be packed (cf. the observations in the preliminary section of the thesis) in the ensemble $\begin{array}{c} \text{---} \mathbb{C}^{X_i} \text{---} \\ \text{---} \mathcal{H}_i \text{---} \end{array} \boxed{\Lambda_i} \begin{array}{c} \text{---} \mathbb{C}^{Y_i} \text{---} \\ \text{---} \mathcal{H}_i \text{---} \end{array}$, and as such the behaviour P is nothing but the channel

$$\begin{array}{c} \text{---} \mathbb{C}^{X_A} \text{---} \\ \text{---} \mathcal{H}_A \text{---} \end{array} \boxed{\Lambda_A} \begin{array}{c} \text{---} \mathbb{C}^{Y_A} \text{---} \\ \text{---} \mathcal{H}_B \text{---} \end{array} \\ \begin{array}{c} \boxed{\varrho} \\ \text{---} \mathcal{H}_B \text{---} \end{array} \boxed{\Lambda_B} \begin{array}{c} \text{---} \mathbb{C}^{Y_B} \text{---} \\ \text{---} \mathcal{H}_B \text{---} \end{array} \quad (5.23)$$

A behaviour P is often called a *quantum behaviour* (or, *quantum correlation*) if it is the behaviour of some quantum strategy, and in our terminology the quantum behaviours are thus simply the Bell-channels in **QIT**.

Definition 5.2.3. (Reducibility of Strategies.)

Let (ϱ, Π_A, Π_B) and $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ be two quantum strategies, with $\tilde{\psi}$ a pure state. We say that (ϱ, Π_A, Π_B) is *reducible to* $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ if there exists a purification $\psi \in \text{St}(\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{P})$ of ϱ , systems $\mathcal{H}_A^{\text{res}}$ and $\mathcal{H}_B^{\text{res}}$, a pure state ψ^{res} on $\mathcal{H}_A^{\text{res}} \otimes \mathcal{H}_B^{\text{res}} \otimes \mathcal{P}$ and isometries $W_i : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \tilde{\mathcal{H}}_i^{\text{res}}$ such that

$$[W_A \Pi_A^{x_A}(y_A) \otimes W_B \Pi_B^{x_B}(y_B) \otimes \mathbf{1}_{\mathcal{P}}] |\psi\rangle = [\tilde{\Pi}_A^{x_A}(y_A) \otimes \tilde{\Pi}_B^{x_B}(y_B)] |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \quad \text{for all } x \in X, y \in Y, \quad (5.24)$$

or, with $W = W_A \otimes W_B$, more compactly

$$[W \Pi^x(y) \otimes \mathbf{1}_{\mathcal{P}}] |\psi\rangle = \tilde{\Pi}^x(y) |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \quad \text{for all } x \in X, y \in Y, \quad (5.25)$$

where, by usual abuse of notation, we write $|\cdot\rangle$ to denote vector representatives of pure states.¹¹ ■

Remark 5.2.4. (Terminology and Scope.)

The term *reducible* is not standard in the literature, indeed the condition is rarely separated from the self-testing definition and explicitly named. The domain of the reducibility relation is slightly odd, since ϱ can be arbitrary but $\tilde{\psi}$ is assumed pure, but the relation restricts to a pre-order on the class of pure-state strategies. Alternatively, it could be made a pre-order on the class of all strategies by rephrasing it to quantify also over purifications of potentially mixed state $\tilde{\varrho}$ of the strategy $(\tilde{\varrho}, \tilde{\Pi}_A, \tilde{\Pi}_B)$. This notion is apparently never considered in the literature on self-testing. ✂

¹¹The condition is not dependent on the choice of vector representatives, since a phase may be absorbed into one of the isometries W_i .

Definition 5.2.5. (Quantum Self-Testing and Rigidity According to the Literature.)
 Suppose that $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ is a quantum strategy with behaviour P , whose state $\tilde{\psi}$ is pure. We say that P *self-tests the strategy* $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ if every strategy (ϱ, Π_A, Π_B) with behaviour P is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$.

A behaviour P is called *rigid according to the literature* (for short, *rigid a.t.t.l.*) if P self-tests some quantum strategy. ■

Remark 5.2.6. (To be Pure or Not to be Pure.)

In Definition 5.2.5 we quantify over strategies whose state ϱ may be non-pure. In much work, this quantification is (seemingly unconsciously) restricted to strategies for which the state ϱ is pure, thus obtaining an a priori weaker condition. Later results will imply (cf. Remark 5.2.18) that if the isometries W_i in condition (5.24) can be chosen independently of ϱ (only depending on the PVMs $\Pi_i^{x_i}$), and if the behaviour P is extremal in the convex set of quantum behaviours, then these two definitions are actually equivalent. However, I do not know whether this is generally the case, and the question apparently has not been considered in the literature. The strong version of quantifying over all strategies agrees with the definition that appears in the recent review article [ŠB19], and is also used in Ref. [RUV13]. ✂

Example 5.2.7. (Rigidity a.t.t.l. of the CHSH-Behaviour.)

Recall the CHSH-behaviour from Example 4.1.17. This behaviour self-tests the strategy which we explicitly used to define it in Example 4.1.17. Historically, the CHSH-behaviour was the first behaviour to be proven rigid a.t.t.l. ([SW87, PR92, Cir93]), though the terminology of self-testing was only coined later ([MY98, MY04]), in the context of a slightly different behaviour. The CHSH-behaviour remains the quintessential example of quantum self-testing, and there are by now several different proofs for its rigidity ([MYS12, MS13, ŠB19]). ♦

5.2.B Classically Bound Dilations

Consider again the channel (5.21) from the dense class of causal Stinespring dilations of the Bell-channel (P, \mathcal{C}) , i.e. the channel

$$\begin{array}{c}
 \text{---} \mathcal{X}_A \text{---} \boxed{\Sigma_A} \text{---} \mathcal{Y}_A \text{---} \\
 \boxed{\pi} \text{---} \mathcal{H}_A \text{---} \boxed{\Sigma_A} \text{---} \mathcal{E}_A \text{---} \\
 \text{---} \mathcal{H}_B \text{---} \boxed{\Sigma_B} \text{---} \mathcal{E}_B \text{---} \\
 \text{---} \mathcal{X}_B \text{---} \boxed{\Sigma_B} \text{---} \mathcal{Y}_B \text{---}
 \end{array}
 \quad . \quad (5.26)$$

When analysing density in **CIT**, we were successful by introducing the marginals of the individual components,

$$\begin{array}{c}
 \text{---} \mathcal{X}_i \text{---} \boxed{\Lambda_i} \text{---} \mathcal{Y}_i \text{---} \\
 \text{---} \mathcal{H}_i \text{---} \boxed{\Lambda_i} \text{---}
 \end{array}
 :=
 \begin{array}{c}
 \text{---} \mathcal{X}_i \text{---} \boxed{\Sigma_i} \text{---} \mathcal{Y}_i \text{---} \\
 \text{---} \mathcal{H}_i \text{---} \boxed{\Sigma_i} \text{---} \mathcal{E}_i \text{---} \boxed{\text{tr}}
 \end{array}
 \quad \text{and} \quad
 \boxed{\varrho} \text{---} \mathcal{H}_A \text{---} :=
 \begin{array}{c}
 \boxed{\pi} \text{---} \mathcal{H}_A \text{---} \\
 \boxed{\pi} \text{---} \mathcal{E}_0 \text{---} \boxed{\text{tr}} \\
 \boxed{\pi} \text{---} \mathcal{H}_B \text{---}
 \end{array}
 , \quad (5.27)$$

so that

$$\begin{array}{c}
\text{---}\mathcal{X}_A\text{---} \\
\text{---}\mathcal{X}_B\text{---}
\end{array}
\boxed{P}
\begin{array}{c}
\text{---}\mathcal{Y}_A\text{---} \\
\text{---}\mathcal{Y}_B\text{---}
\end{array}
=
\begin{array}{c}
\text{---}\mathcal{X}_A\text{---} \\
\text{---}\mathcal{X}_B\text{---}
\end{array}
\begin{array}{c}
\boxed{\varrho} \\
\text{---}\mathcal{H}_A\text{---} \\
\text{---}\mathcal{H}_B\text{---}
\end{array}
\begin{array}{c}
\boxed{\Lambda_A} \\
\boxed{\Lambda_B}
\end{array}
\begin{array}{c}
\text{---}\mathcal{Y}_A\text{---} \\
\text{---}\mathcal{Y}_B\text{---}
\end{array}
. \quad (5.28)$$

If we can understand the triples $(\varrho, \Lambda_A, \Lambda_B)$ which satisfy Eq. (5.28), then we understand the possible triples $(\pi, \Sigma_A, \Sigma_B)$ defining the dilations (5.26), for these arise by Stinespring dilating the components of the former triples. And now, for the first time in our story, quantum self-testing finally meets causal dilations: Indeed, as we observed above, any triple $(\varrho, \Lambda_A, \Lambda_B)$ that corresponds to an ordinary quantum strategy (i.e. for which Λ_A and Λ_B are ensembles of projective measurements) is a triple that satisfies Eq. (5.28); hence, any ordinary quantum strategy defines a causal Stinespring dilation (5.26). The challenge is now that there might be other triples $(\varrho, \Lambda_A, \Lambda_B)$ satisfying Eq. (5.28) than those arising from such strategies.

First of all, there might of course be triples for which Λ_A and Λ_B correspond to ensembles of measurements which are not projective. This is not problematic, since, as we will show in due time, any causal Stinespring dilation corresponding to such a POVM-strategy is derivable from one corresponding to a PVM-strategy, i.e. the dense class can be further thinned – for this, we use Naimark’s theorem.¹²

A much more peculiar problem, however, is that even though the channel P has classical inputs and outputs, the channels Λ_A and Λ_B actually need not be measurement ensembles:


Example 5.2.8. (Strange Realisations of the CHSH-Behaviour.)

Consider the strategy for the CHSH-behaviour which we defined in Example 4.1.17. Or, more generally, consider for the Bell-scenario with $X_A = X_B = \{0, 1\}$ and $Y_A = Y_B = \{+1, -1\}$ any strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ for which $\tilde{\psi}$ is a pure state on $\mathbb{C}^2 \otimes \mathbb{C}^2$, and all projections $\tilde{\Pi}_i^{x_i}(\pm 1)$ have rank one (i.e. project onto 1-dimensional subspaces). The channel

$$\begin{array}{c}
\text{---}\mathbb{C}^{X_A}\text{---} \\
\text{---}\mathbb{C}^{X_B}\text{---}
\end{array}
\begin{array}{c}
\boxed{\tilde{\psi}} \\
\text{---}\mathbb{C}^2\text{---} \\
\text{---}\mathbb{C}^2\text{---}
\end{array}
\begin{array}{c}
\boxed{\tilde{\Lambda}_A} \\
\boxed{\tilde{\Lambda}_B}
\end{array}
\begin{array}{c}
\text{---}\mathbb{C}^{Y_A}\text{---} \\
\text{---}\mathbb{C}^{Y_B}\text{---}
\end{array}
. \quad (5.29)$$

is the behaviour of this strategy, and it has classical inputs and outputs. Now, however, observe the following: Every projective rank-one measurement on \mathbb{C}^2 can be realised as a unitary conjugation followed by a measurement in the computational basis $(|0\rangle, |1\rangle)$; in other words, there exists an ensemble of unitary conjugations $([\mathcal{Z}_i^{x_i}])_{x_i \in X_i}$ on \mathbb{C}^2 , such that, with $[\tilde{\mathcal{Z}}]_i$ denoting the channel that encodes this ensemble, we have $\begin{array}{c} \text{---}\mathbb{C}^{X_i}\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{\tilde{\Lambda}_i} \begin{array}{c} \text{---}\mathbb{C}^{Y_i}\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} = \begin{array}{c} \text{---}\mathbb{C}^{X_i}\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{[\tilde{\mathcal{Z}}]_i} \begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{\Delta} \begin{array}{c} \text{---}\mathbb{C}^{Y_i}\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array}$, where Δ is the measurement in the computational basis of \mathbb{C}^2 (under some identification of \mathbb{C}^{Y_i} with \mathbb{C}^2). As we have seen in Example 4.3.9, the channel Δ can be written as a convex combination of unitary conjugations, that is, $\begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{\Delta} \begin{array}{c} \text{---}\mathbb{C}^{Y_i}\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} = \begin{array}{c} \text{---}\mathbb{C}^2\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array} \boxed{[\mathcal{V}]_i} \begin{array}{c} \text{---}\mathbb{C}^{Y_i}\text{---} \\ \text{---}\mathbb{C}^2\text{---} \end{array}$ for some classical state τ_i on \mathbb{C}^{m_i} (e.g. with $m_i = 2$), and some ensemble of unitaries $[\mathcal{V}]_i$. In total, we can therefore rewrite the channel (5.29) as

¹²This is the correct way of stating that the measurements in a ‘quantum strategy’ can without loss of generality be assumed to be projective.



Now, if we regroup components so as to make a new triple $(\varrho, \Lambda_A, \Lambda_B)$ for which $\varrho = \tilde{\psi} \otimes \tau_A \otimes \tau_B$ and Λ_i is the channel given by the composition of $[\tilde{\mathcal{Z}}]_i$ and $[\mathcal{V}]_i$, then the channel Λ_i will generally not have classical outputs. In fact, on input $|x_i\rangle\langle x_i|$ from \mathbb{C}^{X_i} and $|k\rangle\langle k|$ from \mathbb{C}^{m_i} , the channel acts as a unitary conjugation from \mathbb{C}^2 to $\mathbb{C}^{Y_i} \cong \mathbb{C}^2$. The resulting triple $(\varrho, \Lambda_A, \Lambda_B)$ has no chance of being interpreted as a strategy in the traditional sense – it corresponds to a ‘strategy’ in which the outcomes are not produced by measurements, but by unitarily conjugating a state ϱ whose conjugates happen to be classical states on $\mathbb{C}^{Y_A} \otimes \mathbb{C}^{Y_B}$. \blacklozenge

The issue illustrated by Example 5.2.8 is not merely a formal nuisance, but has operational ramifications as well. Indeed, the representation (5.30) gives a convex decomposition of e.g. the CHSH-behaviour (interpreting the classical distribution $\tau_A \otimes \tau_B$ as the weights), and it is a convex decomposition into channels which individually map every classical input $|x_A\rangle\langle x_A| \otimes |x_B\rangle\langle x_B|$ to a pure state; since the CHSH-behaviour itself yields a non-pure state on any classical input, these pure states cannot all be identical, so if we dilate $\tau_A \otimes \tau_B$ by forming a copy to obtain an acausal dilation of the CHSH-behaviour (as in Example 4.3.8), then this acausal dilation *does not factor*. In other words, it is possible to have side-information correlated with the outputs, before the inputs have been given, a circumstance which drastically contradicts the usual conception of self-testing ([SB19]). The twist is, of course, that the outputs provided by the strange strategy of Example 5.2.8 are not classical, and thus cannot be interpreted in the usual framework of self-testing. The unwelcome acausal dilation arising from this is in turn related to the strange dilations of measurements which we saw in Example 4.3.9.

In this subsection, we will eliminate this problem, more or less by forcing it away. (We must also enforce classicality of the inputs, but though we ultimately pick the same solution, this problem is of a different character – indeed, it is the honest users of the open interfaces, and not the potentially malicious agents implementing the channel, who provide the inputs.)

The solution we choose, though not overly elegant, is essentially to dismiss the dilations that arise from the strange strategies of Example 5.2.8. That this should be sensible relies on the fact that it is the honest users of the channel who control the open interfaces; as such they may well choose to only provide classical inputs, and to measure the outputs received from the channels. Under this course of action, the class of dilations which we deemed reasonable would shrink:

Definition 5.2.9. (Classical Dilations.)

Let $\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \end{array} \begin{array}{|c} \hline (P, \mathcal{C}) \\ \hline \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ be a Bell-channel with classical inputs and outputs. A causal dilation $\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \\ \mathbb{D} \end{array} \begin{array}{|c} \hline (L, \mathcal{E}) \\ \hline \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \\ \mathbb{E} \end{array}$ of (P, \mathcal{C}) is called *classical* if

Recall from Definition 4.4.3 the concept of *rigidity relative to \mathbf{D}* , namely the notion that a class of causal dilations \mathbf{D} has a \succeq -largest element.

Now, it is possible to prove that rigidity a.t.t.l. of P (Definition 5.2.5) implies rigidity of (P, \mathcal{C}) relative to the class of classical dilations in our sense. However, we shall eventually prove a different result, and there are two reasons for that.

First of all, the above implication is not optimal: Self-testing in the sense of Definition 5.2.5 implies rigidity relative to a slightly larger class of dilations than the classical ones, and up to some small caveats the converse is true as well. On the other hand, I do not know that rigidity relative to classical dilations should conversely imply self-testing in the usual sense.

Secondly, though the problem exposed by Example 5.2.8 is ultimately a child of the strange dilations of measurements (Example 4.3.9), there seems to be something odd about the assumption that we can ignore any dilations of the measurements that are performed by the honest users (cf. the condition (5.31)). After all, this apparently goes against the very idea of asking about the possible dilations that the accessible channel could have come from. More concretely: The devices which allegedly execute our measurements could themselves be manufactured by malicious agents.

Fortunately, these two issues are solved by the same move, namely, by not assuming that absolutely no information leaks to the environment from those measurements, but only restraining the allowed dilations of the measurements so as to exclude the strange acausal dilations. What we must require is that only primitive dilations¹³ of a measurement count as reasonable (cf. Example 4.3.3):

Definition 5.2.10. (Classically Bound Dilations.)

Let $\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} (P, \mathcal{C}) \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ be a Bell-channel with classical inputs and outputs. A causal dilation

$\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} (L, \mathcal{E}) \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ of (P, \mathcal{C}) is called *classically bound* if there exists a causal dilation $(\check{L}, \check{\mathcal{E}})$ of (P, \mathcal{C}) , such that

for some primitive dilations $\bar{\Delta}_{X_i}$ of Δ_{X_i} and $\bar{\Delta}_{Y_i}$ of Δ_{Y_i} .

¹³That is, dilations with no acausal side-information.

Remark 5.2.11. The reason for allowing the dilation $(\check{L}, \check{\mathcal{E}})$ to be distinct from (L, \mathcal{E}) is that if we force $(\check{L}, \check{\mathcal{E}}) = (L, \mathcal{E})$, it does not seem obvious that any dilation derivable from a classically bound dilation is itself classically bound, and this property is desirable for the interpretation of classically bound dilations as ‘the allowed ones’. \spadesuit

Every classical dilation is classically bound, as witnessed by choosing $(\check{L}, \check{\mathcal{E}}) = (L, \mathcal{E})$, $\overline{\Delta}_{X_i} = \Delta_{X_i}$ and $\overline{\Delta}_{Y_i} = \Delta_{Y_i}$. Classical boundedness is however slightly less restrictive than classicality. In fact we have the following result, which expresses that we have achieved exactly the desired:

Proposition 5.2.12. *(Concrete Recharacterisation of Classical Boundedness.)*

The classically bound dilations of (P, \mathcal{E}) are precisely those which are derivable from a dilation of the form

$$\begin{array}{c}
 \text{---} \mathcal{X}_A \text{---} \boxed{\hat{\Lambda}_A} \text{---} \mathcal{Y}_A \text{---} \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \psi \text{---} \mathcal{E}_0 \text{---} \\
 \text{---} \mathcal{X}_B \text{---} \boxed{\hat{\Lambda}_B} \text{---} \mathcal{Y}_B \text{---} \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \hat{\mathcal{E}}_A \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \hat{\mathcal{E}}_B
 \end{array} , \quad (5.33)$$

where ψ is a purification of a state ϱ , and where $\hat{\Lambda}_A$ and $\hat{\Lambda}_B$ are Stinespring dilations of measurement ensembles Λ_A and Λ_B such that

$$\begin{array}{c}
 \text{---} \mathcal{X}_A \text{---} \boxed{P} \text{---} \mathcal{Y}_A \text{---} \\
 \text{---} \mathcal{X}_B \text{---} \boxed{P} \text{---} \mathcal{Y}_B \text{---}
 \end{array} = \begin{array}{c}
 \text{---} \mathcal{X}_A \text{---} \boxed{\Lambda_A} \text{---} \mathcal{Y}_A \text{---} \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \varrho \text{---} \\
 \text{---} \mathcal{X}_B \text{---} \boxed{\Lambda_B} \text{---} \mathcal{Y}_B \text{---}
 \end{array} .$$

Proof. Let us start by observing that the dilations of the form

$$\begin{array}{c}
 \boxed{\hat{\Delta}_{X_A}} \text{---} \mathcal{X}_A \text{---} \boxed{\Sigma_A} \text{---} \mathcal{Y}_A \text{---} \boxed{\hat{\Delta}_{Y_A}} \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \pi \text{---} \mathcal{E}_0 \text{---} \\
 \text{---} \mathcal{X}_B \text{---} \boxed{\Sigma_B} \text{---} \mathcal{Y}_B \text{---} \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \hat{\mathcal{E}}_A \\
 \left. \begin{array}{l} \mathcal{H}_A \\ \mathcal{H}_B \end{array} \right\} \hat{\mathcal{E}}_B
 \end{array} , \quad (5.34)$$

where π , Σ_A and Σ_B are isometric, and where $\hat{\Delta}_{X_i}$ are Stinespring dilations of Δ_{X_i} and $\hat{\Delta}_{Y_i}$ of Δ_{Y_i} , constitute a dense class within the class of classically bound dilations. First, any classically bound dilation (L, \mathcal{E}) can be derived from such a dilation: Pick a dilation $(\check{L}, \check{\mathcal{E}})$ such that

$$\begin{array}{c}
 \boxed{\hat{\Delta}_{X_A}} \text{---} \\
 \boxed{\hat{\Delta}_{X_B}} \text{---} \\
 \text{---} \mathbb{D} \text{---}
 \end{array}
 \begin{array}{c}
 \boxed{\hat{\Delta}_{Y_A}} \text{---} \\
 \boxed{\hat{\Delta}_{Y_B}} \text{---} \\
 \text{---} \mathbb{E} \text{---}
 \end{array}
 \begin{array}{c}
 \text{---} \\
 \boxed{(\check{L}, \check{\mathcal{E}})} \\
 \text{---}
 \end{array}
 \begin{array}{c}
 \text{---} \\
 \boxed{(L, \mathcal{E})} \\
 \text{---}
 \end{array}
 \begin{array}{c}
 \text{---} \mathcal{X}_A \text{---} \mathcal{Y}_A \text{---} \\
 \text{---} \mathcal{X}_B \text{---} \mathcal{Y}_B \text{---} \\
 \text{---} \mathbb{D} \text{---} \mathbb{E} \text{---}
 \end{array} \geq (5.35)$$

Theorem 5.2.13. (*Density of PVM-Dilations.*)
The dilations of (P, \mathcal{C}) of the form

$$\begin{array}{c}
\text{--- } \mathcal{X}_A \text{---} \boxed{\hat{\Lambda}_A} \text{--- } \mathcal{Y}_A \text{---} \\
\boxed{\psi} \text{--- } \mathcal{H}_A \text{---} \boxed{\hat{\Lambda}_A} \text{---} \mathcal{E}_A \text{---} \\
\text{--- } \mathcal{E}_0 \text{---} \\
\boxed{\psi} \text{--- } \mathcal{H}_B \text{---} \boxed{\hat{\Lambda}_B} \text{---} \mathcal{E}_B \text{---} \\
\text{--- } \mathcal{X}_B \text{---} \boxed{\hat{\Lambda}_B} \text{---} \mathcal{Y}_B \text{---}
\end{array}
, \tag{5.39}$$

where ψ is a purification of a state ρ , and where $\hat{\Lambda}_A$ and $\hat{\Lambda}_B$ are Stinespring dilations of projective measurement ensembles Λ_A and Λ_B such that $\begin{array}{c} \text{--- } \mathcal{X}_A \text{---} \\ \text{--- } \mathcal{X}_B \text{---} \end{array} \boxed{P} \begin{array}{c} \text{--- } \mathcal{Y}_A \text{---} \\ \text{--- } \mathcal{Y}_B \text{---} \end{array} = \begin{array}{c} \text{--- } \mathcal{X}_A \text{---} \\ \text{--- } \mathcal{X}_B \text{---} \end{array} \begin{array}{c} \boxed{\rho} \\ \boxed{\rho} \end{array} \begin{array}{c} \boxed{\Lambda_A} \\ \boxed{\Lambda_B} \end{array} \begin{array}{c} \text{--- } \mathcal{Y}_A \text{---} \\ \text{--- } \mathcal{Y}_B \text{---} \end{array}$, is dense in the class of constructible classically bound dilations of (P, \mathcal{C}) .

Remark 5.2.14. (Re-Recharacterisation of Classically Bound Dilations.)

Obviously, a slightly stronger statement is also true, namely that a dilation is classically bound if and only if it can be derived from a dilation of the form (5.39). \blacktimes

Remark 5.2.15. (Relation to the Usual Reduction.)

Theorem 5.2.13 is a precise mathematical statement. In contrast, the standard comment in the literature in motivating the definition of quantum strategies (Definition 5.2.1) – namely, that measurements can ‘without loss of generality’ be assumed to be projective – is in most contexts where it is used not a precise mathematical statement. Indeed, a ‘without loss of generality’-claim is formally the claim that the statement $\forall s \in S : Q(s)$ follows from the statement $\forall s \in S_0 : Q(s)$, where $S_0 \subseteq S$, and in the absence of such a predicate Q it is meaningless. (To some authors, the predicate $Q(s)$ seems to be implicitly ‘the strategy s is reducible to the canonical strategy’, but this of course is meaningless too, since the relation of being reducible to the canonical strategy has not even been *defined* for strategies which do not have projective measurements.) \blacktimes

Proof. We have already seen in Proposition 5.2.12 that the class is dense when Λ_A, Λ_B range over measurement ensembles, so we need only prove that any such dilation can be derived from one in which the measurements are projective.

By the well-known theorem of Naimark (see the preliminary section), any measurement on \mathcal{H} with outcomes in Y , $\begin{array}{c} \text{--- } \mathcal{H} \text{---} \\ \text{--- } \mathcal{H} \text{---} \end{array} \boxed{M} \begin{array}{c} \text{--- } \mathcal{Y} \text{---} \\ \text{--- } \mathcal{Y} \text{---} \end{array}$, can be written as $\begin{array}{c} \text{--- } \mathcal{H} \text{---} \\ \text{--- } \mathcal{H} \text{---} \end{array} \begin{array}{c} \boxed{\phi^{\text{Nai}}} \\ \boxed{\phi^{\text{Nai}}} \end{array} \begin{array}{c} \boxed{M^{\text{Nai}}} \\ \boxed{M^{\text{Nai}}} \end{array} \begin{array}{c} \text{--- } \mathcal{Y} \text{---} \\ \text{--- } \mathcal{Y} \text{---} \end{array}$,

where ϕ^{Nai} is a pure state on some system \mathcal{K}^{Nai} , and where M^{Nai} is a projective measurement on $\mathcal{H} \otimes \mathcal{K}^{\text{Nai}}$. By an inductive argument we can easily extend this to ensembles: If

$$\begin{array}{c} \text{--- } \mathcal{C}^{X_i} \text{---} \\ \text{--- } \mathcal{H}_i \text{---} \end{array} \boxed{\Lambda_i} \begin{array}{c} \text{--- } \mathcal{Y}_i \text{---} \\ \text{--- } \mathcal{Y}_i \text{---} \end{array}
\quad \text{is an ensemble of measurements, it can be written as}
\quad \begin{array}{c} \text{--- } \mathcal{C}^{X_i} \text{---} \\ \text{--- } \mathcal{H}_i \text{---} \end{array} \begin{array}{c} \boxed{\phi_i^{\text{Nai}}} \\ \boxed{\phi_i^{\text{Nai}}} \end{array} \begin{array}{c} \boxed{\Lambda_i^{\text{Nai}}} \\ \boxed{\Lambda_i^{\text{Nai}}} \end{array} \begin{array}{c} \text{--- } \mathcal{Y}_i \text{---} \\ \text{--- } \mathcal{Y}_i \text{---} \end{array}
,$$

with ϕ_i^{Nai} a pure state and Λ_i^{Nai} an ensemble of projective measurements. Then, however, if

$$\hat{\Lambda}_i^{\text{Nai}} \text{ is a Stinespring dilation of } \Lambda_i^{\text{Nai}}, \text{ the channel } \begin{array}{c} \text{--- } \mathcal{C}^{X_i} \text{---} \\ \text{--- } \mathcal{H}_i \text{---} \end{array} \begin{array}{c} \boxed{\hat{\Lambda}_i^{\text{Nai}}} \\ \boxed{\hat{\Lambda}_i^{\text{Nai}}} \end{array} \begin{array}{c} \text{--- } \mathcal{Y}_i \text{---} \\ \text{--- } \mathcal{Y}_i \text{---} \end{array}$$

dilation of Λ_i , so the desired follows by redefining the state ρ as $\rho \otimes \phi_A^{\text{Nai}} \otimes \phi_B^{\text{Nai}}$. \square

Theorem 5.2.13 concludes our strive for locating quantum strategies within the framework of causal dilations: They correspond to the channels (5.39), forming a dense class of classically bound dilations. Whereas it would evidently have been cleaner if that class were dense in the class of all (constructible) dilations, the reader should think of this result as explaining the origin of the traditional quantum strategies. This understanding will be solidified and improved when in the next subsection we see how the relation between strategies used in the conventional definition of self-testing can be re-expressed in terms of operations on the inaccessible interface of (5.39).

5.2.C The Bridge to Quantum Self-Testing

Suppose that (ϱ, Π_A, Π_B) is a quantum strategy, and that Λ_A and Λ_B are the associated measurement ensembles. What does the channel (5.39), i.e. the channel

$$(5.40)$$

look like? That of course depends on the concrete choice of purification and Stinespring dilations, but by the completeness properties of Chapter 2, all choices are equivalent by means of channels acting locally on the inaccessible systems. Therefore, we may fix them at our convenience without compromising density of the dilation class. In equations, the channel Λ_i which encodes the measurement ensemble Π_i is given by

$$\Lambda_i(A \otimes B) = \sum_{y_i \in Y_i, x_i \in X_i} \text{tr}[\Pi^{x_i}(y_i)A] |y_i\rangle\langle y_i| \langle x_i| B |x_i\rangle \quad \text{for } A \in \text{End}(\mathcal{H}_i), B \in \text{End}(\mathcal{X}_i). \quad (5.41)$$

A particularly nice choice of Stinespring dilation of this channel is the isometric channel $\hat{\Lambda}_i$ corresponding to the isometry $S_i : \mathcal{H}_i \otimes \mathcal{X}_i \rightarrow \mathcal{H}_i \otimes \mathcal{X}_i \otimes \mathcal{Y}_i$ given by¹⁴

$$S_i = \sum_{x_i \in X_i, y_i \in Y_i} \Pi_i^{x_i}(y_i) \otimes |x_i\rangle \otimes |y_i\rangle \otimes \langle x_i|. \quad (5.42)$$

With this choice, the isometric channel (5.40) corresponds to the isometry $S : \mathcal{X} \rightarrow \mathcal{H} \otimes \mathcal{X} \otimes \mathcal{Y}$ given by

$$S = \sum_{x \in X, y \in Y} [\Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B) \otimes \mathbf{1}_{\mathcal{E}_0}] |\psi\rangle \otimes |x\rangle \otimes |y\rangle \otimes \langle x|, \quad (5.43)$$

¹⁴Observe that if the measurements $\Pi_i^{x_i}$ had not been projective but rather general POVMs, $E_i^{x_i} = (E_i^{x_i}(y_i))_{y_i \in Y_i}$, then we would have had to include additionally a copy of \mathcal{Y}_i in the inaccessible system, and the Stinespring isometry S_i would instead have been $\sum_{x_i \in X_i, y_i \in Y_i} \sqrt{E_i^{x_i}(y_i)} \otimes |x_i\rangle \otimes |y_i\rangle \otimes \langle x_i|$.

where $\mathcal{X} = \mathcal{X}_A \otimes \mathcal{X}_B$, $\mathcal{Y} = \mathcal{Y}_A \otimes \mathcal{Y}_B$ (and as usual $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$), and where $|x\rangle = |x_A, x_B\rangle = |x_A\rangle \otimes |x_B\rangle$ and $|y\rangle = |y_A, y_B\rangle = |y_A\rangle \otimes |y_B\rangle$. (Observe that $\hat{\mathcal{E}}_i = \mathcal{H}_i \otimes \mathcal{X}_i$.)

Qualitatively, this is the pivotal moment in our storyline: We now see the quantities from the reducibility condition (5.24) emerge in the guise of a causal Stinespring dilation of the behaviour channel. The fundamental link between the reducibility criterion and these casual dilations is contained in the following result:

Theorem 5.2.16. (*Reducibility among Quantum Strategies in Terms of PVM-Dilations.*)

Let $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ be a quantum strategy for which the state $\tilde{\psi}$ is pure and has locally full rank. Then, a quantum strategy (ϱ, Π_A, Π_B) is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ in the sense of Definition 5.2.3 if and only if there exist channels Γ_A, Γ_B such that

$$(5.44)$$

where $\hat{\Lambda}_i$ and $\tilde{\Lambda}_i$ are Stinespring dilations of the corresponding measurement ensembles Λ_i and $\tilde{\Lambda}_i$, respectively, and where ψ is a purification of ϱ .

Remark 5.2.17. (On the Full-Rank Assumption.)

The technical full-rank assumption is, to the best of my knowledge, satisfied in all known examples of quantum self-testing. Also, it is only needed for the ‘if’-direction of the statement. \spadesuit

Remark 5.2.18. (Reducibility of Pure-State Strategies versus General Strategies.)

From Theorem 5.2.16, we may argue that in order to establish self-testing it is often enough to quantify over pure-state strategies (cf. Remark 5.2.6):

Suppose that every pure-state strategy (ψ, Π_A, Π_B) with behaviour P is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, and that this is witnessed by isometries W_i which do not depend on ψ . (This is often the case, cf. the so-called *swap-method* ([SB19]), which produces isometries W_i depending only on Π_i .) By the proof of Theorem 5.2.16, the channels Γ_i are simply marginals of the isometric conjugations by W_i , so Γ_i can then be chosen independently of ψ . Suppose moreover that the behaviour P is extremal (this is true, for example, if the reducibility of pure-state strategies can be established merely based on the value of the behaviour in a so-called *non-local game* ([SB19])). Now, if (ϱ, Π_A, Π_B) is a general strategy with behaviour P , then, writing $\varrho = \sum_{k=1}^n p_k \psi_k$ by the spectral theorem, with ψ_1, \dots, ψ_n pure and $p_1, \dots, p_n > 0$, we conclude by extremality that all of the pure-state strategies (ψ_k, Π_A, Π_B) have behaviour P . We thus find (by the state-independence assumption) channels Γ_A, Γ_B

such that

for all $k = 1, \dots, n$, and hence by forming

the convex combination on each side we have

, which

is exactly the condition (5.44), so (ϱ, Π_A, Π_B) is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$.

In summary, self-testing relative to pure-state strategies along with extremality of the behaviour implies, under the state-independence assumption, the full self-testing condition. (The converse holds as well, since the full self-testing condition always implies extremality of the behaviour, as detailed in Section 5.3.A.) \spadesuit

Proof of Theorem 5.2.16. By Lemma 5.1.8, the condition (5.44) is equivalent to the existence of isometric channels $\hat{\Gamma}_A, \hat{\Gamma}_B$ and a pure state ψ^{res} on $\mathcal{H}_A^{\text{res}} \otimes \mathcal{H}_B^{\text{res}} \otimes \mathcal{E}_0$, such that

$$= \quad (5.45)$$

with our convenient choice of Stinespring dilations $\hat{\Lambda}_i$ and $\hat{\Lambda}_i$ from above. Eq. (5.45) is an identity between two isometric channels, and can thus be rephrased as an identity between the two isometries which represent them (possibly absorbing a phase). Let $V_i : \mathcal{H}_i \otimes \mathcal{X}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}} \otimes \mathcal{X}_i$ be an isometry representing the channel $\hat{\Gamma}_i$, and let $|\psi\rangle, |\tilde{\psi}\rangle$ and $|\psi^{\text{res}}\rangle$ represent the states. Then, Eq. (5.45) is equivalent to the equation

$$\begin{aligned} & \sum_{x \in X, y \in Y} [V_A^{x_A} \Pi_A^{x_A}(y_A) \otimes V_B^{x_B} \Pi_B^{x_B}(y_B) \otimes \mathbb{1}_{\mathcal{E}_0}] |\psi\rangle \otimes |y\rangle \otimes \langle x| \\ &= \sum_{x \in X, y \in Y} [\tilde{\Pi}_A^{x_A}(y_A) \otimes \tilde{\Pi}_B^{x_B}(y_B)] |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \otimes |x\rangle \otimes |y\rangle \otimes \langle x|, \end{aligned} \quad (5.46)$$

where $V_i^{x_i} : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}} \otimes \mathcal{X}_i$ denotes the isometry $V_i(\mathbb{1}_{\mathcal{H}_i} \otimes |x_i\rangle)$. Now, since the system $(|y\rangle \otimes \langle x|)_{x \in X, y \in Y}$ is orthonormal, this identity can be rephrased component-wise as

$$\forall x \in X, y \in Y : [V^x \Pi^x(y) \otimes \mathbb{1}_{\mathcal{E}_0}] |\psi\rangle = \tilde{\Pi}^x(y) |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \otimes |x\rangle, \quad (5.47)$$

with the abbreviations $V^x = V_A^{x_A} \otimes V_B^{x_B}$, and, as usual, $\Pi^x(y) = \Pi_A^{x_A}(y_A) \otimes \Pi_B^{x_B}(y_B)$, $\tilde{\Pi}^x(y) = \tilde{\Pi}_A^{x_A}(y_A) \otimes \tilde{\Pi}_B^{x_B}(y_B)$. In comparison, the reducibility condition (Eq. (5.25)) reads

$$\forall x \in X, y \in Y : [W \Pi^x(y) \otimes \mathbb{1}_{\mathcal{E}_0}] |\psi\rangle = \tilde{\Pi}^x(y) |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle, \quad (5.48)$$

with $W = W_A \otimes W_B$, for some isometries $W_i : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}}$ and some pure state ψ^{res} . If Eq. (5.48) holds, then Eq. (5.47) follows by taking $V_i = W_i \otimes \mathbb{1}_{\mathcal{X}_i}$, so that $V_i^{x_i} = W_i \otimes |x_i\rangle$. Hence, it is clear that (5.44) follows from the reducibility condition. That conversely Eq. (5.48) follows from Eq. (5.47) is not obvious, and it is the content of Lemma 5.2.24 below (using the assumption that $\tilde{\psi}$ has locally full rank). Hence, the reducibility condition follows from (5.44), finishing the proof. \square

The main technicalities of the above proof reside in Lemma 5.2.24 below. We will postpone this lemma to the end of the subsection, however, in order to first discuss the consequences of Theorem 5.2.16 and phrase it more transparently in our language.

When P self-tests $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, Theorem 5.2.16 apparently says that the causal dilation corresponding to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ can be derived from the causal dilation corresponding to any strategy (ϱ, Π_A, Π_B) . This, of course, is the *wrong way around* compared to our conception of rigidity. The point is, however, that as in the proof of Theorem 5.1.10, the purified relation (5.45) allows us to *move the channels $\hat{\Gamma}_i$ to the other side*, thus realising also that the dilation corresponding to (ϱ, Π_A, Π_B) is derivable from the dilation corresponding to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$. As such, self-testing really implies that all dilations corresponding to strategies are \succeq -equivalent (as in Corollary 5.1.12).

Consequently, rigidity a.t.t.l. implies in our language rigidity relative to classically bound dilations, but in the rather strong sense that all pure causal dilations are \succeq -equivalent. In this light, there is nothing special about the ‘canonical’ strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$; any representative of the equivalence class will be a complete dilation of (P, \mathcal{C}) . However, Eq. (5.44) expresses that the dilation corresponding to the strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ is particularly simple, because it can be derived from the others using *local* operations in the environment.

More precisely, suppose we define the following pre-order on the class of pure causal dilations:

Definition 5.2.19. (Local Derivability.)

Let (P, \mathcal{C}) be a Bell-channel. For two pure causal dilations we write

$$\begin{array}{ccc}
 \begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \\ \begin{array}{|c|} \hline \pi \\ \hline \end{array} \\ \text{---}\mathcal{X}_B\text{---} \end{array} & \begin{array}{c} \begin{array}{|c|} \hline \Sigma_A \\ \hline \end{array} \\ \begin{array}{|c|} \hline \Sigma_B \\ \hline \end{array} \end{array} & \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \\ \text{---}\mathcal{E}_A\text{---} \\ \text{---}\mathcal{E}_B\text{---} \\ \text{---}\mathcal{E}_0\text{---} \end{array} \\
 & \succeq_{loc.} & \\
 \begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \\ \begin{array}{|c|} \hline \pi' \\ \hline \end{array} \\ \text{---}\mathcal{X}_B\text{---} \end{array} & \begin{array}{c} \begin{array}{|c|} \hline \Sigma'_A \\ \hline \end{array} \\ \begin{array}{|c|} \hline \Sigma'_B \\ \hline \end{array} \end{array} & \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \\ \text{---}\mathcal{E}'_A\text{---} \\ \text{---}\mathcal{E}'_B\text{---} \\ \text{---}\mathcal{E}'_0\text{---} \end{array}
 \end{array} \tag{5.49}$$

and say that the latter is *locally derivable* from the former if there exist channels Γ_A, Γ_B and Γ_0 such that

$$\begin{array}{ccc}
 \begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \\ \begin{array}{|c|} \hline \pi \\ \hline \end{array} \\ \text{---}\mathcal{X}_B\text{---} \end{array} & \begin{array}{c} \begin{array}{|c|} \hline \Sigma_A \\ \hline \end{array} \\ \begin{array}{|c|} \hline \Sigma_B \\ \hline \end{array} \end{array} & \begin{array}{c} \text{---}\mathcal{Y}_B\text{---} \\ \text{---}\mathcal{E}_A\text{---} \\ \text{---}\mathcal{E}_0\text{---} \\ \text{---}\mathcal{E}_B\text{---} \\ \text{---}\mathcal{E}'_A\text{---} \\ \text{---}\mathcal{E}'_0\text{---} \\ \text{---}\mathcal{E}'_B\text{---} \end{array} \\
 & = & \\
 \begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \\ \begin{array}{|c|} \hline \pi' \\ \hline \end{array} \\ \text{---}\mathcal{X}_B\text{---} \end{array} & \begin{array}{c} \begin{array}{|c|} \hline \Sigma'_A \\ \hline \end{array} \\ \begin{array}{|c|} \hline \Sigma'_B \\ \hline \end{array} \end{array} & \begin{array}{c} \text{---}\mathcal{Y}_B\text{---} \\ \text{---}\mathcal{E}'_A\text{---} \\ \text{---}\mathcal{E}'_0\text{---} \\ \text{---}\mathcal{E}'_B\text{---} \end{array}
 \end{array} \tag{5.50}$$

■

Then, the pre-order $\succeq_{loc.}$ refines the pre-order \succeq , and we have the following which will constitute our best relation between self-testing (that is, rigidity a.t.t.l.) and rigidity in our sense:

Corollary 5.2.20. (*Rigidity versus Rigidity a.t.t.l.*)

Let $\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \end{array} \boxed{(P, \mathcal{E})} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$ be bipartite Bell-channel, and let $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ be a quantum strategy with behaviour P for which the state $\tilde{\psi}$ is pure and has locally full rank. Then, the following are equivalent:

1. P self-tests $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$.
2. All the causal Stinespring dilations of (P, \mathcal{E}) corresponding to quantum strategies are \succeq -equivalent, and $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ is a $\succeq_{loc.}$ -smallest element of this equivalence class.
3. (P, \mathcal{E}) is rigid relative to classically bound dilations and has a complete dilation with no acausal side-information. Moreover, the \succeq -equivalence class of causal Stinespring dilations corresponding to quantum strategies has a $\succeq_{loc.}$ -smallest element, namely the dilation corresponding to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$.

Proof. The equivalence of 1. and 2. follows from Theorem 5.2.16 and the considerations that follow. The equivalence of 2. and 3. follows by Corollary 5.1.12 (or rather, an adaption of this corollary to the class of classically bound dilations). \square

Now, of the statements appearing in Corollary 5.2.20 only the statement ‘ (P, \mathcal{E}) is rigid relative to classically bound dilations and has a complete dilation with no acausal side-information.’ is of a purely operational nature. Indeed, as far as I can see, the causal Stinespring dilations themselves a priori have no operational interpretation.

Consider therefore the following question:

Open Problem 5.2.21. (*Existence of Simple Representatives.*)

Does every \succeq -equivalence class of [classically bound] causal Stinespring dilations contain a $\succeq_{loc.}$ -smallest element?

If it has an affirmative answer, then, by Corollary 5.2.20, quantum self-testing can be given a purely operational formulation. Indeed, the following will be true (assuming that also the full-rank condition on the state is always satisfied):

Conjecture 5.2.22. (*Equivalence of Quantum Self-Testing with Acausal Rigidity.*)

Given a bipartite Bell-channel $\begin{array}{c} \mathcal{X}_A \\ \mathcal{X}_B \end{array} \boxed{(P, \mathcal{E})} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}$, the following are equivalent:

1. P is rigid a.t.t.l., i.e. self-tests (in the sense of Definition 5.2.5) some strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$.
2. (P, \mathcal{E}) is rigid relative to classically bound dilations, and has a complete dilation with no acausal side-information.

On the other hand, there seems to me to be no good reasons why it should not be possible to have rigidity by means of a complete dilation with non-trivial acausal side-information:

Open Problem 5.2.23. (*Rigidity Beyond the Usual Assumptions of Quantum Self-Testing.*)

Does there exist a Bell-channel (P, \mathcal{E}) which is rigid relative to classically bound dilations, but for which no complete dilation has no acausal side-information?

We end the subsection by establishing the missing technical ingredient for the proof of Theorem 5.2.16:

Lemma 5.2.24. (Technical Result for Theorem 5.2.16).

Let (ϱ, Π_A, Π_B) and $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ be quantum strategies, for which $\tilde{\psi}$ is a pure state with locally full rank. Suppose that there exists a purification ψ of ϱ with purifying space \mathcal{P} , a pure state ψ^{res} on $\mathcal{H}_A^{\text{res}} \otimes \mathcal{H}_B^{\text{res}} \otimes \mathcal{P}$, and isometries $V_i^{x_i} : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}} \otimes \mathcal{X}_i$, such that, with $V^x = V_A^{x_A} \otimes V_B^{x_B}$,

$$[V^x \Pi^x(y) \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle = \tilde{\Pi}^x(y) |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \otimes |x\rangle \quad \text{for all } x \in X, y \in Y. \quad (5.51)$$

Then, there exist x_i -independent isometries $W_i : \mathcal{H}_i \rightarrow \tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}}$ such that, with $W = W_A \otimes W_B$,

$$[W \Pi^x(y) \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle = \tilde{\Pi}^x(y) |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \quad \text{for all } x \in X, y \in Y. \quad (5.52)$$

In other words, (ϱ, Π_A, Π_B) is reducible to $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$.

Remark 5.2.25. (On an Approximate Generalisation.)

It is clear that the condition (5.51) implies that the behaviours of the two strategies (ϱ, Π_A, Π_B) and $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ are identical. Nevertheless, it has never been considered in the literature that this condition, in which the local isometries are allowed to depend on the local inputs (x_i) , should actually be the correct notion of reducibility among quantum strategies. The result of this lemma gives a possible explanation for this. It is worth noting, however, that the relationship between the two conditions (5.51) and (5.52) is probably more subtle in the approximate case (of so-called *robust* self-testing), since we use at a point in the proof a full-rank argument without regards to the fact that some Schmidt coefficients may be small. \spadesuit

Proof. First, we argue that there exists an x_A -independent isometry $W_A : \mathcal{H}_A \rightarrow \tilde{\mathcal{H}}_A \otimes \mathcal{H}_A^{\text{res}}$ such that the operator identity

$$W_A = [\mathbb{1}_{\tilde{\mathcal{H}}_A \otimes \mathcal{H}_A^{\text{res}}} \otimes \langle x_A |] V_A^{x_A} \quad (5.53)$$

holds on the subspace $\text{supp}(\varrho_A)$ of \mathcal{H}_A , for all $x_A \in X_A$. (Here, ϱ_A is the A-marginal of ϱ .)

By summing over $y \in Y$ in condition (5.51), we obtain the identity

$$[V^x \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle = |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \otimes |x\rangle \quad (5.54)$$

for all $x \in X$. Now, let

$$|\psi\rangle = \sum_{j=1}^r \sqrt{p(j)} |\psi_A(j)\rangle \otimes |\psi_{-A}(j)\rangle \quad (5.55)$$

be a Schmidt decomposition of $|\psi\rangle$ relative to the factorisation $\mathcal{H}_A \otimes \mathcal{H}_{-A}$ where $\mathcal{H}_{-A} = \mathcal{H}_B \otimes \mathcal{P}$, with $p(1), \dots, p(r) > 0$. We then have by Eq. (5.54), for $x = (x_A, x_B) \in X$,

$$\begin{aligned}
|\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \otimes |x_{\text{B}}\rangle &= [\mathbb{1}_{\tilde{\mathcal{H}} \otimes \mathcal{H}^{\text{res}} \otimes \mathcal{P} \otimes \mathcal{X}_{\text{B}}} \otimes \langle x_{\text{A}}|] [V^x \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle \\
&= \sum_{j=1}^r \sqrt{p(j)} [\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}| V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle \otimes [V_{\text{B}}^{x_{\text{B}}} \otimes \mathbb{1}_{\mathcal{P}}] |\psi_{-\text{A}}(j)\rangle].
\end{aligned} \tag{5.56}$$

The left hand side of Eq. (5.56) is independent of x_{A} . The isometry $V_{\text{B}}^{x_{\text{B}}}$ is also independent of x_{A} , whence the orthonormal system $([V_{\text{B}}^{x_{\text{B}}} \otimes \mathbb{1}_{\mathcal{P}}] |\psi_{-\text{A}}\rangle(j))_{j=1, \dots, r}$ is independent of x_{A} . Therefore,

$$[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}| V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle] \tag{5.57}$$

must be independent of x_{A} , for all $j = 1, \dots, r$. (In this step we would loose something in an approximate argument, when dividing by $\sqrt{p(j)}$.) Observing that

$$\text{span}\{|\psi_{\text{A}}(j)\rangle \mid j = 1, \dots, r\} = \text{supp}(\varrho_{\text{A}}), \tag{5.58}$$

we conclude that the operator $[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}| V_{\text{A}}^{x_{\text{A}}}]$ restricted to $\text{supp}(\varrho_{\text{A}})$ is independent of x_{A} . Moreover, it acts isometrically:

By the Eq. (5.56),

$$1 = \|\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \otimes |x_{\text{B}}\rangle\|^2 = \sum_{j=1}^r p(j) \|[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}| V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle]\|^2, \tag{5.59}$$

and by extremality of 1 in the convex set $[0, 1]$ this forces $\|[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}| V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle]\| = 1$ for all $j = 1, \dots, r$ (here we would lose again in the approximate version). It follows that

$$\|[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes |x_{\text{A}}\rangle\langle x_{\text{A}}|] V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle\| = 1, \tag{5.60}$$

and since $\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes |x_{\text{A}}\rangle\langle x_{\text{A}}|$ is a projection and $V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle$ is a unit vector, we must therefore have

$$[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes |x_{\text{A}}\rangle\langle x_{\text{A}}|] V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle = V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle, \tag{5.61}$$

so that

$$(V_{\text{A}}^{x_{\text{A}}})^* [\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes |x_{\text{A}}\rangle\langle x_{\text{A}}|] V_{\text{A}}^{x_{\text{A}}} |\psi_{\text{A}}(j)\rangle = |\psi_{\text{A}}(j)\rangle \tag{5.62}$$

for all $j = 1, \dots, r$, which is precisely the statement that $[\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}|] V_{\text{A}}^{x_{\text{A}}}$ acts isometrically on $\text{supp}(\varrho_{\text{A}})$.

Altogether, we conclude as desired the existence of an isometry $W_{\text{A}} : \mathcal{H}_{\text{A}} \rightarrow \tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}$ such that $W_{\text{A}} = [\mathbb{1}_{\tilde{\mathcal{H}}_{\text{A}} \otimes \mathcal{H}_{\text{A}}^{\text{res}}} \otimes \langle x_{\text{A}}|] V_{\text{A}}^{x_{\text{A}}}$ on $\text{supp}(\varrho_{\text{A}})$, for any $x_{\text{A}} \in X_{\text{A}}$.

By the exact same argument, we find an x_B -independent isometry $W_B : \mathcal{H}_B \rightarrow \tilde{\mathcal{H}}_B \otimes \mathcal{H}_B^{\text{res}}$ such that $W_B = [\mathbb{1}_{\tilde{\mathcal{H}}_B \otimes \mathcal{H}_B^{\text{res}}} \otimes \langle x_B |] V_B^{x_B}$ on $\text{supp}(\varrho_B)$ for any $x_B \in X_B$.

Next, we argue that the operator identity $W_i = [\mathbb{1}_{\tilde{\mathcal{H}}_i \otimes \mathcal{H}_i^{\text{res}}} \otimes \langle x_i |] V_i^{x_i}$ holds not only on $\text{supp}(\varrho_i) = \text{span}\{|\psi_i(j)\rangle \mid j = 1, \dots, r\}$, but in fact on the a priori larger space

$$\begin{aligned} & \text{span}\{\Pi_i^{x_i}(y_i) |\psi_i(j)\rangle \mid j = 1, \dots, r, x_i \in X_i, y_i \in Y_i\} \\ & (= \text{span}\{\Pi_i^{x_i}(y_i) |\phi_i\rangle \mid |\phi_i\rangle \in \text{supp}(\varrho_i), x_i \in X_i, y_i \in Y_i\}). \end{aligned} \quad (5.63)$$

This will imply, by (5.51), that the local isometry $W := W_A \otimes W_B$ witnesses condition (5.52), and thus altogether finish the proof. To demonstrate the assertion, we show that – under the assumption that the marginal $\tilde{\varrho}_i$ of $\tilde{\psi}$ has full rank – the subspace (5.63) in fact coincides with $\text{supp}(\varrho_i)$.¹⁵

By rewriting Eq. (5.54) in terms of density matrices (i.e. $[V^x \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle\langle\psi| [V^x \otimes \mathbb{1}_{\mathcal{P}}]^* = |\tilde{\psi}\rangle\langle\tilde{\psi}| \otimes |\psi^{\text{res}}\rangle\langle\psi^{\text{res}}| \otimes |x\rangle\langle x|$) and marginalising to site $i \in \{A, B\}$, we obtain

$$V_i^{x_i} \varrho_i (V_i^{x_i})^* = \tilde{\varrho}_i \otimes \varrho_i^{\text{res}} \otimes |x_i\rangle\langle x_i|, \quad (5.64)$$

which implies (since $(V_i^{x_i})^*$ is surjective) the following equality of subspaces:

$$V_i^{x_i} [\text{supp}(\varrho_i)] = \text{supp}(\tilde{\varrho}_i) \otimes \text{supp}(\varrho_i^{\text{res}}) \otimes \text{span}\{|x_i\rangle\}. \quad (5.65)$$

By a similar rewriting, the condition (5.51) implies that

$$V_i^{x_i} \Pi_i^{x_i}(y_i) \varrho_i V_i^{x_i*} = \tilde{\Pi}_i^{x_i}(y_i) \tilde{\varrho}_i \otimes \varrho_i^{\text{res}} \otimes |x_i\rangle\langle x_i|, \quad (5.66)$$

where this time, before marginalising to i , the ket side was obtained from (5.51) by summing only over the elements of Y_{-i} , whereas the bra side was obtained by summing over all elements of Y . This analogously implies the following equality of subspaces:¹⁶

$$V_i^{x_i} \Pi_i^{x_i}(y_i) [\text{supp}(\varrho_i)] = \tilde{\Pi}_i^{x_i}(y_i) [\text{supp}(\tilde{\varrho}_i)] \otimes \text{supp}(\varrho_i^{\text{res}}) \otimes \text{span}\{|x_i\rangle\}. \quad (5.67)$$

Now, by the full-rank assumption, $\tilde{\Pi}_i^{x_i}(y_i) [\text{supp}(\tilde{\varrho}_i)] \subseteq \text{supp}(\tilde{\varrho}_i)$, and since $V_i^{x_i}$ has a left-inverse (indeed, it is an isometry), we thus conclude from (5.65) and (5.67) that $\Pi_i^{x_i}(y_i) [\text{supp}(\varrho_i)] \subseteq \text{supp}(\varrho_i)$. This proves the desired assertion, and ultimately finishes the proof of the lemma. \square

5.3 Some Implications of a Reformulation

In this final section, I will highlight how some of the (very few) general results which are known about quantum self-testing can be easily and transparently proved using the equivalent formulation provided by Theorem 5.2.16.

¹⁵It may be observed from the argument that this conclusion relies only on the weaker circumstance that $\text{supp}(\tilde{\varrho}_i)$ is invariant under the operators $(\tilde{\Pi}_i^{x_i}(y_i))_{y_i \in Y_i, x_i \in X_i}$, and thus we could alternatively have stated the lemma under this weaker assumption.

¹⁶Eq. (5.65) actually follows from Eq. (5.67) by summing over $y_i \in Y_i$, but a direct argument for (5.65) was given first for the sake of transparency.

The first such result is that a behaviour P which self-tests some strategy is necessarily extremal in the convex set of all (tensor-product) quantum behaviours. The second and third result assert that if P self-tests the strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, then for any strategy (ϱ, Π_A, Π_B) with behaviour P , the state $\tilde{\psi}$ can be locally extracted from the state ϱ , and the measurement ensemble $\tilde{\Lambda}_i$ can be extracted from the measurement ensemble Λ_i on the local support of the state ϱ (in a suitable sense). This third result clarifies the connection to other notions of self-testing of measurements, cf. Def. 3 in Ref. [SB19], as e.g. employed in Ref. [Kan17].

5.3.A Rigidity, Security and Extremality

It has for some time been part of the folklore in quantum self-testing that a necessary condition for P to self-test a strategy is extremality of P in the convex set of all behaviours realisable by quantum strategies. A formal general proof of this, however, appears to have been first given in Ref. [GKW⁺18]. As observed earlier in Ref. [FFW11], extremality in the set of quantum behaviours is equivalent to a natural notion of *security*, and thus a valid way to show the implication is by demonstrating that self-testing implies security in this sense. This implication can be seen as a precise version of the remarkable insight of Ekert's in Ref. [Eke91], which was mentioned already in the general introduction to the thesis.

Below, I will give a definition of the security notion in the language of dilations, prove that it implies extremality, and then prove that self-testing implies security. This yields altogether a different proof for the necessity of extremality, and also directly the proof for security, with the interpretation that the randomness generated in a Bell-experiment is genuinely random, in the sense of being unknowable before the inputs to the channel are provided.

Recall from Example 4.3.8 the notion of *acausal side-information*.

Definition 5.3.1. (Security.)

Let \mathbf{D} be a class of causal dilations of a causal channel $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{(P, \mathcal{C})} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$. We say that (P, \mathcal{C}) is *secure (against acausal side-information) relative to \mathbf{D}* if any acausal dilation $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{(L, \mathcal{E})} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ in \mathbf{D} is trivial, i.e. factors as $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{(P, \mathcal{C})} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array} \begin{array}{c} \text{---}\mathcal{D}\text{---} \\ \text{---}\mathcal{E}\text{---} \end{array} \boxed{(S, \mathcal{D})} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ for some causal channel (S, \mathcal{D}) . ■

The reader should have in mind the case where \mathbf{D} is class of classically bound dilations; security then means that any reasonable side-information about the channel, which is present before the inputs have been supplied, must be completely independent of the channel ('reasonable' referring to classical boundedness). In particular, any randomness generated by the channel (P, \mathcal{C}) is *fresh*, in the sense that it must be uncorrelated with pre-existing randomness.

We have the following:

Proposition 5.3.2. (Self-Testing implies Security.)

If $\begin{array}{c} \text{---}\mathcal{X}_A\text{---} \\ \text{---}\mathcal{X}_B\text{---} \end{array} \boxed{(P, \mathcal{C})} \begin{array}{c} \text{---}\mathcal{Y}_A\text{---} \\ \text{---}\mathcal{Y}_B\text{---} \end{array}$ is a Bell-channel in **QIT** for which P self-tests some quantum strategy, then (P, \mathcal{C}) is secure relative to classically bound dilations.

Proof. By Corollary 5.2.20, (P, \mathcal{C}) has a complete dilation for classically bound dilations, with no acausal side-information. The statement should be intuitively clear from this circumstance, but let us prove it explicitly:

Let $\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \left[\begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \\ \mathbb{E}_0 \end{array} \right] (K, \mathcal{F})$ be a dilation of (P, \mathcal{C}) which is complete for classically bound dilations and which has no acausal side-information. By Lemma 4.3.21, any classically bound dilation is thus of the form

$$\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \left[\begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \\ \mathbb{E}_0 \end{array} \right] (K, \mathcal{F}) \quad \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \\ \mathbb{E}_0 \end{array} \left[\begin{array}{c} \mathbb{D} \\ \mathbb{E} \end{array} \right] (G, \mathcal{B}) \quad (5.68)$$

for some causal channel (G, \mathcal{B}) . If the dilation (5.68) is acausal, which is to say that $\mathcal{F}(\text{ports}(\mathbb{E}_0) \cap \mathcal{B}(\text{ports}(\mathbb{E}))) = \emptyset$, then, since $\mathcal{F}(e_0) \neq \emptyset$ for all $e_0 \in \text{ports}(\mathbb{E}_0)$, we must have $\text{ports}(\mathbb{E}_0) \cap \mathcal{B}(\text{ports}(\mathbb{E})) = \emptyset$, i.e. $\mathcal{B}(\text{ports}(\mathbb{E})) \subseteq \text{ports}(\mathbb{D})$. This however implies that (G, \mathcal{B})

factors as $\begin{array}{c} \sim \mathbb{E}_0 \sim \\ \sim \mathbb{D} \sim \end{array} \left[\begin{array}{c} \text{tr} \\ (S, \mathcal{D}) \end{array} \right] \begin{array}{c} \mathbb{E} \\ \mathbb{E} \end{array}$, and the desired follows. \square

Remark 5.3.3. (On the Significance of Security.)

As demonstrated by the proof, most of this statement has nothing to do with self-testing per se; it is true in any theory that the existence of a complete dilation with no acausal side-information implies security. What makes the statement *interesting* in the context of quantum self-testing, is that a quantum behaviour can give random outputs and still have a complete dilation with no acausal side-information.

This is not true, for example, in **CIT**: We have seen that some Bell-channels in **CIT** have complete dilations (cf. Example 4.4.25 and the examples following Corollary 4.4.17), but any such causal channel which has a complete dilation with no acausal side-information must be a deterministic channel. Indeed, by a consideration as the one in Example 4.3.8 (or as in Proposition 5.3.4 below), such a channel is extremal. \spadesuit

Proposition 5.3.4. (Security and Extremality.)

Let $\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \left[\begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \right] (P, \mathcal{C})$ be a Bell-channel in **QIT**. If (P, \mathcal{C}) is secure relative to classically bound dilations, then P is extremal in the convex set of quantum behaviours from \mathcal{X} to \mathcal{Y} .

Proof. First, observe that, by density of one-sided dilations (due to the constructibility assumption), (P, \mathcal{C}) is secure if and only if one-sided acausal dilations are trivial. What we will do is to establish a correspondence between the classically bound one-sided acausal dilations of (P, \mathcal{C}) and the convex decompositions of P into quantum behaviours. Conceptually, this correspondence is analogous to that in the case of **CIT** as expressed by Theorem 4.4.14 or Theorem 4.4.15.

More precisely, given a convex decomposition $P = \sum_{k=1}^n s_k P_k$ with P_k a quantum behaviour and $s_1, \dots, s_n > 0$, we may write $\begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \left[\begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \right] P = \begin{array}{c} -\mathcal{X}_A \\ -\mathcal{X}_B \end{array} \left[\begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \right] \tilde{P}$, where σ is the

state $\sum_{k=1}^n s_k |k\rangle\langle k|$ on \mathbb{C}^n and \tilde{P} the channel that reads the value k in the \mathbb{C}^n -register and yields P_k accordingly. Giving \tilde{P} the obvious causal specification $\tilde{\mathcal{C}}$, this identity moreover holds between causal channels as well. Also, $(\tilde{P}, \tilde{\mathcal{C}})$ is constructible since each P_k is constructible. Letting $\left[\begin{array}{c} \mathbb{C}^n \\ \mathbb{C}^n \end{array} \right] \tilde{\sigma}$ denote the classical copy of σ , the channel

$$(5.69)$$

is evidently a classically bound acausal dilation¹⁷ of (P, \mathcal{C}) . It moreover encodes the initial convex decomposition of P , as it is given by $\sum_{k=1}^n s_k P_k \otimes |k\rangle\langle k|$. Now, if (P, \mathcal{C}) is secure relative to classically bound dilations, then the acausal dilation (5.69) must factor, and this implies that $P_1 = P_2 = \dots = P_n$. Hence, P is extremal. \square

Proposition 5.3.2 and Proposition 5.3.4 together imply that if P self-tests some strategy (i.e. is rigid a.t.t.l.), then P is necessarily extremal.

Open Problem 5.3.5. (*Extremality versus Rigidity*)

Does every extremal quantum behaviour self-test some quantum strategy?

5.3.B Extraction of the Canonical State

It is well-known that the usual definition of quantum self-testing implies that the state $\tilde{\psi}$ of a strategy $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$ which is self-tested by P can be locally extracted from the state ϱ of any strategy (ϱ, Π_A, Π_B) with behaviour P . Indeed, if in the reducibility condition

$$[W_A \Pi_A^{x_A}(y_A) \otimes W_B \Pi_B^{x_B}(y_B) \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle = [\tilde{\Pi}_A^{x_A}(y_A) \otimes \tilde{\Pi}_B^{x_B}(y_B)] |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle \quad (5.70)$$

we sum over $y_A \in Y_A$ and $y_B \in Y_B$ (for any fixed x_A, x_B), then we obtain

$$[W_A \otimes W_B \otimes \mathbb{1}_{\mathcal{P}}] |\psi\rangle = |\tilde{\psi}\rangle \otimes |\psi^{\text{res}}\rangle, \quad (5.71)$$

or, more compactly,

$$[W \otimes \mathbb{1}_{\mathcal{P}}] \psi [W^* \otimes \mathbb{1}_{\mathcal{P}}] = \tilde{\psi} \otimes \psi^{\text{res}}, \quad (5.72)$$

with $W = W_A \otimes W_B$. In pictures, this last condition reads

$$(5.73)$$

with Ξ_i the isometric channel corresponding to conjugation by W_i . This condition is a purified version of the condition that there exist some (not necessarily isometric) channels M_A and M_B such that

¹⁷Recall the tacit assumption that we only consider constructible dilations – therefore, it is important for the argument that $(\tilde{P}, \tilde{\mathcal{C}})$ is constructible.

$$\begin{array}{c} \rho \\ \hline \mathcal{H}_A \quad \mathcal{H}_B \\ \hline \end{array} \begin{array}{c} \boxed{M_A} \\ \boxed{M_B} \end{array} \begin{array}{c} \tilde{\mathcal{H}}_A \\ \tilde{\mathcal{H}}_B \end{array} = \begin{array}{c} \tilde{\psi} \\ \hline \tilde{\mathcal{H}}_A \\ \tilde{\mathcal{H}}_B \end{array}, \quad (5.74)$$

asserting local extractability of $\tilde{\psi}$ from ρ . Though this argument is so simple that it hardly needs improvement, it is instructive to see how local extractability of the canonical state $\tilde{\psi}$ is manifested in the self-testing formulation offered by Theorem 5.2.16. There will be no summing over y_A, y_B in equations involving operators; another argument altogether (which would also be valid in most other theories) proves extractability:

Proposition 5.3.6. (*Self-Testing implies Local Extractability of the Canonical State.*)
Suppose that P self-tests $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$. If (ρ, Π_A, Π_B) is any strategy with behaviour P , then there exist channels M_A and M_B such that (5.74) holds.

Proof. By Theorem 5.2.16, there are channels Γ_i such that

$$\begin{array}{c} \mathcal{X}_A \\ \hline \rho \\ \hline \mathcal{X}_B \end{array} \begin{array}{c} \boxed{\hat{\Lambda}_A} \\ \boxed{\hat{\Lambda}_B} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array} \begin{array}{c} \boxed{\Gamma_A} \\ \boxed{\Gamma_B} \end{array} = \begin{array}{c} \mathcal{X}_A \\ \hline \tilde{\psi} \\ \hline \mathcal{X}_B \end{array} \begin{array}{c} \boxed{\hat{\Lambda}_A} \\ \boxed{\hat{\Lambda}_B} \end{array} \begin{array}{c} \mathcal{Y}_A \\ \mathcal{Y}_B \end{array}. \quad (5.75)$$

By applying inverses to the reversible channels $\hat{\Lambda}_i$ on both sides, we obtain

$$\begin{array}{c} \mathcal{X}_A \\ \hline \rho \\ \hline \mathcal{X}_B \end{array} \begin{array}{c} \boxed{M'_A} \\ \boxed{M'_B} \end{array} \begin{array}{c} \mathcal{X}_A \\ \tilde{\mathcal{H}}_A \\ \tilde{\mathcal{H}}_B \\ \mathcal{Y}_B \end{array} = \begin{array}{c} \mathcal{X}_A \\ \hline \tilde{\psi} \\ \hline \mathcal{X}_B \end{array}, \quad (5.76)$$

for some channels M'_i . The desired then follows by inserting arbitrary states into the systems \mathcal{X}_i and trashing them. \square

5.3.C Extraction of Canonical Measurements

Finally, let us prove that self-testing implies local extractability of the canonical measurement ensembles. This connects to alternative formulations of self-testing (e.g. as in Ref. [Kan17] in the case of the ‘tilted CHSH-behaviour’), but as far as I know the general statement below has not been presented before. The proof provided here is pictorial, in the same style as the one above for local extractability of the state, and in principle it generalises beyond **QIT** to other purifiable theories.

Proposition 5.3.7. (*Self-Testing implies Local Extractability of Canonical Measurements.*)
Suppose that P self-tests $(\tilde{\psi}, \tilde{\Pi}_A, \tilde{\Pi}_B)$, and assume that $\tilde{\psi}$ has locally full rank. If (ρ, Π_A, Π_B) is any strategy with behaviour P , then there exist channels N_A and N_B such that

$$\begin{array}{c} \mathcal{X}_i \\ \hline \tilde{\mathcal{H}}_i \end{array} \begin{array}{c} \boxed{N_i} \\ \boxed{\Lambda_i} \end{array} \mathcal{Y}_i = \begin{array}{c} \mathcal{X}_i \\ \hline \tilde{\mathcal{H}}_i \end{array} \begin{array}{c} \boxed{\tilde{\Lambda}_i} \end{array} \mathcal{Y}_i \quad (5.77)$$

for $i = A, B$, where Λ_i and $\tilde{\Lambda}_i$ are the measurement ensembles corresponding to $(\Pi_i^{x_i})_{x_i \in X_i}$ and $(\tilde{\Pi}_i^{x_i})_{x_i \in X_i}$, respectively.

Remark 5.3.8. (Generalisation to the Approximate Case.)

As observed in the proof, approximate versions of this result (e.g. in terms of the diamond-distance) will generally lose factors, but the loss depends only on the Schmidt coefficients of the fixed and known state $\tilde{\psi}$. \blacktimes

Proof. It is enough to show the existence of N_A . We know that (5.44) holds for some channels Γ_A, Γ_B . We can actually prove the desired from the weaker circumstance that occurs after purifying and moving channels to the other side, namely from the assumption that

$$(5.78)$$

for some state ψ^{res} and some channels $\tilde{\Gamma}_A$ and $\tilde{\Gamma}_B$.

We may assume without loss of generality that $\tilde{\Gamma}_B$ is reversible. (If not, we simply replace it by a pure reversible dilation, absorb the pure state arising on the left hand side into ψ , noting that the modified dilated measurement channel $\hat{\Lambda}_B \parallel \text{id}$ is a Stinespring dilation of $\Lambda_B \otimes \text{tr}$, which is extractibility-equivalent to Λ_B .)

Under this assumption, first observe that by applying a left-inverse $\hat{\Lambda}_A^-$ to $\hat{\Lambda}_A$ on both sides, and then inserting an arbitrary state on \mathcal{X}_A and trashing \mathcal{X}_A , yields

$$(5.79)$$

for some channel N'_A . By inserting this fragment back into Eq. (5.78), we find

$$(5.80)$$

Now, cancelling $\tilde{\Gamma}_B$ (which we assumed reversible) and afterwards $\hat{\Lambda}_B$, and then inserting an arbitrary state on \mathcal{X}_B and trashing \mathcal{X}_B , yields

$$(5.81)$$

Trashing furthermore the two lower systems of ψ^{res} finally yields

$$(5.82)$$

for some channels N_A and M_A . Now, since $\tilde{\psi}$ has locally full rank¹⁸ it is a universal dilation of its marginal, so we must have

$$(5.83)$$

which even more strongly than desired asserts that any dilation of Λ_A can be extracted by means of N_A from some dilation of $\tilde{\Lambda}_A$. \square

5.4 Summary and Outlook

In this chapter, we have seen a precise connection between quantum self-testing as it is traditionally envisioned and the framework of causal dilations presented in the thesis.

Specifically, we established that quantum strategies form a dense class of *classically bound* dilations (Definition 5.2.10) of the behaviour channel, and we then formulated the usual reducibility condition between strategies in terms of these dilations (Theorem 5.2.16). By virtue of a partial collapse in the causal-dilational ordering due to purifiability of **QIT**, we were able to relate this condition to the existence of a complete causal dilation with no acausal side-information, along with the existence of a simple representative in the \succeq -equivalence class of strategies (Corollary 5.2.20). It has moreover been conjectured that such a representative always exists and that quantum self-testing can thus be recast in purely operational terms (Conjecture 5.2.22).

We have also seen that the causal Stinespring dilations of a quantum behaviour can be surprisingly recharacterised by non-signalling conditions (Theorem 5.1.2) and seen how this implies a recharacterisation of the set of quantum behaviours (Corollary 5.1.5).

Many problems are left open for future work:

¹⁸This is the first argument which would not be robust to errors.

1. Can we give a proof of the rigidity of e.g. the CHSH-behaviour without reference to Corollary 5.2.20, but instead based in the language of causal dilations rather than linear operators?
2. Do the Open Problem 5.2.21 and the Conjecture 5.2.22 have affirmative answers?
3. Does the Open Problem 5.2.23 have an affirmative answer?
4. What happens to the causal-dilational ordering, and rigidity by extension, if we replace the constructible scheme $\mathbf{E} = \text{Cons}(\mathbf{QIT})$ with another? Is it feasible, and is it relevant?
5. Is Corollary 5.1.5 useful?

It is very interesting to observe, as mentioned in Remark 5.2.25, that the relationship between quantum strategies that seems to arise from the framework of causal dilations allows local isometries which depend on the local input. Though it happens in the exact case that this dependence can be eliminated (that is the content of Lemma 5.2.24), the situation in the approximate case (of *robust* self-testing [MYS12]) may very well be different. As such, ‘robustness’ results could in the framework of this thesis turn out to be of an entirely different quantitative character than in the usual framework. This is probably one of the most interesting problems for future research.

Conclusion

The results of this thesis were motivated by the desire to recast quantum self-testing in a new, operational language. We have seen that such a recasting is viable (Chapter 5) in a framework of causally structured dilations (Chapter 4) which applies to physical theories of a very general kind (Chapter 1). In the course of contemplating dilations systematically, we have also seen that a rather rich theory is possible based on a handful of principles pertaining only to the structure of dilations (Chapter 2), and considered how a metric version of this theory can be created (Chapter 3).

Each chapter has already been concluded with summaries and detailed overviews of open problems for the reader to consult (Section 1.4, Section 2.6, Section 3.6, Section 4.5, Section 5.4).

Let me reiterate, however, what could be the three most interesting open directions for future work.

First, there is the problem of understanding in more detail the causal-dilational ordering. Systematic considerations (in the style of Chapter 2) may be possible, but it might of course also be the case that general techniques simply lack, and that its structure depends heavily on both the theory and the causal channel in question.

Secondly, there is the problem of extending the metric considerations of Chapter 3 to the causal setting of Chapter 4. Possible challenges in this regard have already been mentioned in the prelude to Chapter 4. Such an extension is not only interesting on theoretical grounds, but also because it would seem the proper way of recasting *robust* self-testing in the language of causal dilations, cf. the results of Chapter 5. As mentioned in Section 5.4, this might ultimately yield robustness results of a different kind.

Finally, there is the curious question of whether it might be possible to derive a self-testing result not by reference to Corollary 5.2.20, but rather by giving a proof within the language of causal dilations. If such a proof were to be found, which made no reference to linear operators, it would likely yield a new and enlightening perspective on quantum self-testing.

One of the most general – but also most vague – points that surface from the work in this thesis is that causal dilations seem to formally cover the intuitive notion of ‘implementations’ of a physical process. We saw this already in the general introduction with the bit refreshment channel (and it was formalised in Section 4.4.C), and it is also the message of the link between quantum strategies and causal dilations (Theorem 5.2.13). In a sense, this correspondence between implementations and causal dilations is the moral of the thesis. It

is possible that there is a precise and eloquent way of expressing this connection – or that it is even, when properly viewed, tautological – but at this point I do not know the answer to that question.

References

- [AC04] Samson Abramsky and Bob Coecke, *A categorical semantics of quantum protocols*, Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004, IEEE, 2004, pp. 415–425.
- [AH16] Hanne Andersen and Brian Hepburn, *Scientific method*, The Stanford Encyclopedia of Philosophy (Edward N. Zalta, ed.), Metaphysics Research Lab, Stanford University, summer 2016 ed., 2016.
- [Att14] Stephane Attal, *Lecture 6: Quantum channels*, Lecture notes available online at <http://math.univ-lyon1.fr/~attal/chapters.html>, 2014.
- [Awo10] Steve Awodey, *Category theory*, Oxford university press, 2010.
- [Bae06] John C. Baez, *Quantum quandaries: A category-theoretic perspective*, The structural foundations of quantum gravity (2006), 240–265.
- [Bar07] Jonathan Barrett, *Information processing in generalized probabilistic theories*, Physical Review A **75** (2007), no. 3, 032304.
- [BB84] Charles H. Bennett and Gilles Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Proceedings of the International Conference on Computers, Systems and Signal Processing, vol. 1, 1984, pp. 175–179.
- [BCF⁺96] Howard Barnum, Carlton M. Caves, Christopher A. Fuchs, Richard Jozsa, and Benjamin Schumacher, *Noncommuting mixed states cannot be broadcast*, Physical Review Letters **76** (1996), no. 15, 2818–2821.
- [BCP⁺14] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner, *Bell nonlocality*, Reviews of Modern Physics **86** (2014), no. 2, 419–478.
- [BCS12] Harry Buhrman, Matthias Christandl, and Christian Schaffner, *Complete insecurity of quantum protocols for classical two-party computation*, Physical review letters **109** (2012), no. 16, 160501.
- [Bel64] John S. Bell, *On the Einstein Podolsky Rosen paradox*, Physics Physique Fizika **1** (1964), no. 3, 195–200.
- [BF14] John C. Baez and Tobias Fritz, *A bayesian characterization of relative entropy*, arXiv preprint arXiv:1402.3067 (2014).
- [BGNP01] David Beckman, Daniel Gottesman, Michael Nielsen, and John Preskill, *Causal and localizable quantum operations*, Physical Review A **64** (2001), no. 5, 052309.

- [BMR92] Samuel L. Braunstein, Ady Mann, and Michael Revzen, *Maximal violation of Bell inequalities for mixed states*, Physical Review Letters **68** (1992), no. 22, 3259–3261.
- [Boh52] David Bohm, *A suggested interpretation of the quantum theory in terms of "hidden" variables. I*, Physical review **85** (1952), no. 2, 166–179.
- [BP07] Samuel L Braunstein and Arun K Pati, *Quantum information cannot be completely hidden in correlations: implications for the black-hole information paradox*, Physical review letters **98** (2007), no. 8, 080502.
- [Bru14] Časlav Brukner, *Quantum causality*, Nature Physics **10** (2014), no. 4, 259–263.
- [BS10] John C. Baez and Mike Stay, *Physics, topology, logic and computation: a Rosetta Stone*, New structures for physics, Springer, 2010, pp. 95–172.
- [BT24] Stefan Banach and Alfred Tarski, *Sur la décomposition des ensembles de points en parties respectivement congruentes*, Fund. math **6** (1924), no. 1, 244–277.
- [Bur69] Donald Bures, *An extension of Kakutani's theorem on infinite product measures to the tensor product of semifinite W^* -algebras*, Transactions of the American Mathematical Society **135** (1969), 199–212.
- [BW11] Howard Barnum and Alexander Wilce, *Information processing in convex operational theories*, Electronic Notes in Theoretical Computer Science **270** (2011), no. 1, 3–15.
- [BW16] ———, *Post-classical probability theory*, Quantum Theory: Informational Foundations and Foils, Springer, 2016, pp. 367–420.
- [Can84] Georg Cantor, *Über unendliche, lineare Punktmannichfaltigkeiten*, Mathematische Annalen **23** (1884), no. 4, 453–488.
- [CDKW16] Bob Coecke, Ross Duncan, Aleks Kissinger, and Quanlong Wang, *Generalised compositional theories and diagrammatic reasoning*, Quantum Theory: Informational Foundations and Foils, Springer, 2016, pp. 309–366.
- [CDP09] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti, *Theoretical framework for quantum networks*, Physical Review A **80** (2009), no. 2, 022339.
- [CDP10] ———, *Probabilistic theories with purification*, Physical Review A **81** (2010), no. 6, 062348.
- [CDP11] ———, *Informational derivation of quantum theory*, Physical Review A **84** (2011), no. 1, 012311.
- [CDPV13] Giulio Chiribella, Giacomo Mauro D'Ariano, Paolo Perinotti, and Benoit Valiron, *Quantum computations without definite causal structure*, Physical Review A **88** (2013), no. 2, 022318.
- [CFS16] Bob Coecke, Tobias Fritz, and Robert W. Spekkens, *A mathematical theory of resources*, Information and Computation **250** (2016), 59–86.
- [CH17] Oscar Cunningham and Chris Heunen, *Purity through factorisation*, arXiv preprint arXiv:1705.07652 (2017).

- [Chi14a] Giulio Chiribella, *Dilation of states and processes in operational-probabilistic theories*, arXiv preprint arXiv:1412.8539 (2014).
- [Chi14b] ———, *Distinguishability and copiability of programs in general process theories*, Int J Software Informatics **1** (2014), no. 2, 209–223.
- [Cho75] Man-Duen Choi, *Completely positive linear maps on complex matrices*, Linear algebra and its applications **10** (1975), no. 3, 285–290.
- [CHSH69] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt, *Proposed experiment to test local hidden-variable theories*, Physical review letters **23** (1969), no. 15, 880–884.
- [Chu] The Church of the larger Hilbert space on *quantiki.org*, <https://www.quantiki.org/wiki/church-larger-hilbert-space>.
- [Cir80] Boris S. Cirel’son, *Quantum generalizations of Bell’s inequality*, Letters in Mathematical Physics **4** (1980), no. 2, 93–100.
- [Cir93] ———, *Some results and problems on quantum Bell-type inequalities*, Hadronic Journal Supplement **8** (1993), no. 4, 329–345.
- [CL13] Bob Coecke and Raymond Lal, *Causal categories: Relativistically interacting processes*, Foundations of Physics **43** (2013), no. 4, 458–501.
- [Coe11] Bob Coecke, *A universe of processes and some of its guises*, Deep beauty: Understanding the quantum world through mathematical innovation (2011), 129–186.
- [Coe14] ———, *Terminality implies non-signalling*, arXiv preprint arXiv:1405.3681 (2014).
- [CS16] Giulio Chiribella and Robert W. Spekkens, *Quantum Theory: Informational Foundations and Foils*, Springer, 2016.
- [CS18] Andrea Coladangelo and Jalex Stark, *Unconditional separation of finite and infinite-dimensional quantum correlations*, arXiv preprint arXiv:1804.05116 (2018).
- [CS20] ———, *An inherently infinite-dimensional quantum correlation*, Nature communications **11** (2020), no. 1, 1–6.
- [CT09] Matthias Christandl and Ben Toner, *Finite de Finetti theorem for conditional probability distributions describing physical theories*, Journal of mathematical physics **50** (2009), no. 4, 042104.
- [Deu85] David Deutsch, *Quantum theory, the Church–Turing principle and the universal quantum computer*, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences **400** (1985), no. 1818, 97–117.
- [Dir81] Paul Adrien Maurice Dirac, *The principles of quantum mechanics*, no. 27, Oxford university press, 1981.
- [DKSW07] Giacomo Mauro D’Ariano, Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner, *Reexamination of quantum bit commitment: The possible and the impossible*, Physical Review A **76** (2007), no. 3, 032328.

- [dP67] John de Pillis, *Linear transformations which preserve Hermitian and positive semidefinite operators*, Pacific Journal of Mathematics **23** (1967), no. 1, 129–137.
- [DP16] Kenneth J. Dykema and Vern Paulsen, *Synchronous correlation matrices and Connes’ embedding conjecture*, Journal of Mathematical Physics **57** (2016), no. 1, 015214.
- [DS05] Igor Devetak and Peter W. Shor, *The capacity of a quantum channel for simultaneous transmission of classical and quantum information*, Communications in Mathematical Physics **256** (2005), no. 2, 287–303.
- [Ein05] Albert Einstein, *Zur Elektrodynamik bewegter Körper*, Annalen der Physik **4** (1905), 891–921.
- [Ein16] A Einstein, *Die Grundlage der allgemeinen Relativitätstheorie*, Annalen der Physik **354** (1916), no. 7, 769–822.
- [Eke91] Artur K. Ekert, *Quantum cryptography based on Bell’s theorem*, Physical Review Letters **67** (1991), no. 6, 661–663.
- [EML45] Samuel Eilenberg and Saunders Mac Lane, *General theory of natural equivalences*, Transactions of the American Mathematical Society **58** (1945), no. 2, 231–294.
- [End01] Herbert B. Enderton, *A mathematical introduction to logic*, Elsevier, 2001.
- [Enr] Enriched Categories on nLab, <https://ncatlab.org/nlab/show/enriched+category>.
- [EPR35] Albert Einstein, Boris Podolsky, and Nathan Rosen, *Can quantum-mechanical description of physical reality be considered complete?*, Physical review **47** (1935), no. 10, 777–780.
- [ESW02] Tilo Eggeling, Dirk Schlingemann, and Reinhard F Werner, *Semicausal operations are semilocalizable*, EPL (Europhysics Letters) **57** (2002), no. 6, 782–788.
- [Fey82] Richard P. Feynman, *Simulating physics with computers*, Int. J. Theor. Phys **21** (1982), no. 6/7, 467–488.
- [Fey87] ———, *Negative probability*, Quantum implications: Essays in honour of David Bohm (1987), 235–248.
- [FFW11] Torsten Franz, Fabian Furrer, and Reinhard F. Werner, *Extremal quantum correlations and cryptographic security*, Physical review letters **106** (2011), no. 25, 250502.
- [Fol99] Gerald B. Folland, *Real analysis: Modern techniques and their applications*, 2 ed., vol. 40, John Wiley & Sons, 1999.
- [Fri20] Tobias Fritz, *A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics*, Advances in Mathematics **370** (2020), 107239.

- [GKW⁺18] Koon Tong Goh, Jędrzej Kaniewski, Elie Wolfe, Tamás Vértesi, Xingyao Wu, Yu Cai, Yeong-Cherng Liang, and Valerio Scarani, *Geometry of the set of quantum correlations*, Physical Review A **97** (2018), no. 2, 022104.
- [Göd29] Kurt Gödel, *Über die Vollständigkeit des Logikkalküls*, Doctoral Thesis, 1929.
- [Göd31] ———, *Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I*, Monatshefte für Mathematik und Physik **38** (1931), no. 1, 173–198.
- [Gra] *Wikipedia-page on various graph products*, https://en.wikipedia.org/wiki/Graph_product.
- [Gri05] David J. Griffiths, *Introduction to Quantum Mechanics*, 2 ed., Pearson, 2005.
- [Han09] Ernst Hansen, *Measure Theory*, 4 ed., University of Copenhagen, Department of Mathematical Sciences, 2009.
- [Har01] Lucien Hardy, *Quantum theory from five reasonable axioms*, arXiv preprint quant-ph/0101012 (2001).
- [Har07] Lucien Hardy, *Towards quantum gravity: a framework for probabilistic theories with non-fixed causal structure*, Journal of Physics A: Mathematical and Theoretical **40** (2007), no. 12, 3081–3099.
- [Har10] ———, *A formalism-local framework for general probabilistic theories including quantum theory*, arXiv preprint arXiv:1005.5164 (2010).
- [Has09] Masahito Hasegawa, *On traced monoidal closed categories*, Mathematical Structures in Computer Science **19** (2009), no. 2, 217–244.
- [HPS08] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman, *An introduction to mathematical cryptography*, vol. 1, Springer, 2008.
- [Jam72] Andrzej Jamiołkowski, *Linear transformations which preserve trace and positive semidefiniteness of operators*, Reports on Mathematical Physics **3** (1972), no. 4, 275–278.
- [JLF13] Min Jiang, Shunlong Luo, and Shuangshuang Fu, *Channel-state duality*, Physical Review A **87** (2013), no. 2, 022310.
- [JNV⁺20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen, *MIP^{*}=RE*, arXiv preprint arXiv:2001.04383 (2020).
- [JS91] André Joyal and Ross Street, *The Geometry of Tensor Calculus, I*, Advances in mathematics **88** (1991), no. 1, 55–112.
- [JSV96] André Joyal, Ross Street, and Dominic Verity, *Traced Monoidal Categories*, Mathematical Proceedings of the Cambridge Philosophical Society, Cambridge University Press, 1996, pp. 447–468.
- [Kan17] Jędrzej Kaniewski, *Self-testing of binary observables based on commutation*, Physical Review A **95** (2017), no. 6, 062323.

- [KSW08a] Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner, *A continuity theorem for Stinespring’s dilation*, Journal of Functional Analysis **255** (2008), no. 8, 1889–1904.
- [KSW08b] ———, *The information-disturbance tradeoff and the continuity of Stinespring’s representation*, IEEE transactions on information theory **54** (2008), no. 4, 1708–1717.
- [KU17] Aleks Kissinger and Sander Uijlen, *A categorical semantics for causal structure*, 2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), IEEE, 2017, pp. 1–12.
- [Kuh12] Thomas S. Kuhn, *The Structure of Scientific Revolutions*, University of Chicago press, 2012.
- [Kun80] Kenneth Kunen, *Set theory – an introduction to independence proofs*, Elsevier, 1980.
- [Lan61] Rolf Landauer, *Irreversibility and heat generation in the computing process*, IBM journal of research and development **5** (1961), no. 3, 183–191.
- [Law62] Francis William Lawvere, *The category of probabilistic mappings*, Lecture Notes available online at <https://ncatlab.org/nlab/files/lawvereprobability1962.pdf>, 1962.
- [Law73] Francis William Lawvere, *Metric spaces, generalized logic, and closed categories*, Rendiconti del seminario matematico e fisico di Milano **43** (1973), no. 1, 135–166.
- [LC97] Hoi-Kwong Lo and Hoi Fung Chau, *Is quantum bit commitment really possible?*, Physical Review Letters **78** (1997), no. 17, 3410–3413.
- [Lo97] Hoi-Kwong Lo, *Insecurity of quantum secure computations*, Physical Review A **56** (1997), no. 2, 1154–1162.
- [Mac17] Peter Machamer, *Galileo Galilei*, The Stanford Encyclopedia of Philosophy (Edward N. Zalta, ed.), Metaphysics Research Lab, Stanford University, summer 2017 ed., 2017.
- [May97] Dominic Mayers, *Unconditionally secure quantum bit commitment is impossible*, Physical review letters **78** (1997), no. 17, 3414–3417.
- [Mic] *Quote by Albert A. Michelson*, https://en.wikiquote.org/wiki/Albert_A._Michelson.
- [ML13] Saunders Mac Lane, *Categories for the Working Mathematician*, vol. 5, Springer Science & Business Media, 2013.
- [MPSS18] Kavan Modi, Arun Kumar Pati, Aditi Sen(De), and Ujjwal Sen, *Masking quantum information is impossible*, Physical review letters **120** (2018), no. 23, 230501.
- [MS13] Carl A. Miller and Yaoyun Shi, *Optimal robust self-testing by binary nonlocal XOR games*, 8th Conference on the Theory of Quantum Computation, Communication and Cryptography, 2013, pp. 254–262.

- [MY98] Dominic Mayers and Andrew Yao, *Quantum cryptography with imperfect apparatus*, Proceedings 39th Annual Symposium on Foundations of Computer Science (Cat. No. 98CB36280), IEEE, 1998, pp. 503–509.
- [MY04] ———, *Self testing quantum apparatus*, Quantum Information & Computation **4** (2004), no. 4, 273–286.
- [MYS12] Matthew McKague, Tzyh Haur Yang, and Valerio Scarani, *Robust self-testing of the singlet*, Journal of Physics A: Mathematical and Theoretical **45** (2012), no. 45, 455304.
- [Nai40] Mark Naimark, *Spectral functions of a symmetric operator*, Izvestiya Rossiiskoi Akademii Nauk. Seriya Matematicheskaya **4** (1940), no. 3, 277–318.
- [NC02] Michael Nielsen and Isaac Chuang, *Quantum Computation and Quantum Information*, American Association of Physics Teachers, 2002.
- [OCB12] Ognjan Oreshkov, Fabio Costa, and Časlav Brukner, *Quantum correlations with no causal order*, Nature communications **3** (2012), no. 1, 1–8.
- [PB00] Arun Kumar Pati and Samuel L. Braunstein, *Impossibility of deleting an unknown quantum state*, Nature **404** (2000), no. 6774, 164–165.
- [Pea90] Giuseppe Peano, *Sur une courbe, qui remplit toute une aire plane*, Mathematische Annalen **36** (1890), no. 1, 157–160.
- [Pen71] Roger Penrose, *Applications of negative dimensional tensors*, Combinatorial mathematics and its applications **1** (1971), 221–244.
- [Per17] Paolo Perinotti, *Causal structures and the classification of higher order quantum computations*, Time in physics, Springer, 2017, pp. 103–127.
- [PHHH06] Marco Piani, Michal Horodecki, Pawel Horodecki, and Ryszard Horodecki, *Properties of quantum nonsignaling boxes*, Physical Review A **74** (2006), no. 1, 012305.
- [Pin18] Steven Pinker, *Enlightenment now: The case for reason, science, humanism, and progress*, Penguin, 2018.
- [PR92] Sandu Popescu and Daniel Rohrlich, *Which states violate Bell’s inequality maximally?*, Physics Letters A **169** (1992), no. 6, 411–414.
- [PR94] Sandu Popescu and Daniel Rohrlich, *Quantum nonlocality as an axiom*, Foundations of Physics **24** (1994), no. 3, 379–385.
- [R⁺18] Hans Rosling et al., *Factfulness: Ten reasons we’re wrong about the world – and why things are better than you think*, Flatiron Books, 2018.
- [Rus] *Wikipedia-page on Russell’s paradox*, https://en.wikipedia.org/wiki/Russell%27s_paradox#cite_note-2.
- [RUV13] Ben W. Reichardt, Falk Unger, and Umesh Vazirani, *Classical command of quantum systems*, Nature **496** (2013), no. 7446, 456–460.
- [ŠB19] Ivan Šupić and Joseph Bowles, *Self-testing of quantum systems: a review*, arXiv preprint arXiv:1904.10042 (2019).

- [Sch] *Wikipedia-page on Cantor-Schröder-Bernstein theorem*, https://en.wikipedia.org/wiki/Schroeder-Bernstein_theorem.
- [Sch17] René L. Schilling, *Measures, Integrals and Martingales*, Cambridge University Press, 2017.
- [Sel04] Peter Selinger, *Towards a semantics for higher-order quantum computation*, Proceedings of the 2nd International Workshop on Quantum Programming Languages, TUCS General Publication, vol. 33, Citeseer, 2004, pp. 127–143.
- [Sel10] ———, *A survey of graphical languages for monoidal categories*, New structures for physics, Springer, 2010, pp. 289–355.
- [Sem] Semicartesian Monoidal Categories on *nLab*, <https://ncatlab.org/nlab/show/semicartesian+monoidal+category>.
- [Sha48] Claude E. Shannon, *A mathematical theory of communication*, The Bell System Technical Journal **27** (1948), no. 3, 379–423.
- [Sti55] William Forrest Stinespring, *Positive functions on C^* -algebras*, Proceedings of the American Mathematical Society **6** (1955), no. 2, 211–216.
- [SW87] Stephen J. Summers and Reinhard F. Werner, *Maximal violation of Bell's inequalities is generic in quantum field theory*, Communications in Mathematical Physics **110** (1987), no. 2, 247–259.
- [TCR10] Marco Tomamichel, Roger Colbeck, and Renato Renner, *Duality between smooth min- and max-entropies*, IEEE Transactions on information theory **56** (2010), no. 9, 4674–4681.
- [Thi] Thin Categories on *nLab*, <https://ncatlab.org/nlab/show/thin+category>.
- [Tom12] Marco Tomamichel, *A framework for non-asymptotic quantum information theory*, Ph.D. thesis, ETH Zürich, 2012.
- [vD13] Wim van Dam, *Implausible consequences of superstrong nonlocality*, Natural Computing **12** (2013), no. 1, 9–12.
- [Wat] John Watrous, *CS 766/QIC 820 Theory of Quantum Information (Fall 2011)*, Lecture notes available online at <https://cs.uwaterloo.ca/~watrous/TQI-notes/TQI-notes.pdf>.
- [Wei72] Karl Weierstrass, *Über continuirliche Functionen eines reellen Arguments, die für keinen Werth des letzteren einen bestimmten Differentialquotienten besitzen*, Math. Werke (1872), 71–74.
- [Wie83] Stephen Wiesner, *Conjugate coding*, ACM Sigact News **15** (1983), no. 1, 78–88.
- [Wol19] Michael Wolf, *Mathematical introduction to quantum information processing (growing lecture notes, SS2019)*, Lecture notes available online at https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MA5057_2019S/QIPlecture.pdf, 2019.
- [WZ82] William K. Wootters and Wojciech H. Zurek, *A single quantum cannot be cloned*, Nature **299** (1982), no. 5886, 802–803.