Commitment Algorithms

Katriel Cohn-Gordon

Merton College

University of Oxford

Supervised by Dr. Chris Heunen

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Master of Science
in
Mathematics and the Foundations of Computer Science

Trinity 2012

Abstract

In this dissertation we examine *commitment protocols* in the graphical language for the foundations of quantum mechanics.

The previous sentence sums up the background material in Part I. Chapter 1 begins with an introduction to classical algorithms and classical cryptography, extending these to the quantum case, which – at least ordinarily – is formulated in terms of Hilbert spaces. Chapter 2 presents the categorical formulation of quantum mechanics: the graphical language for concurrent processes and its semantics in terms of †-compact categories. In particular, it examines the interpretation in the "toy" category **Rel** of sets and relations.

In Part II we move to more original work. We begin in Chapter 3 with a review of classical bit commentment, and present an abstract definition of a commitment protocol in a general †-compact category. We interpret it first in the usual category **FHilb** of Hilbert spaces and continuous linear maps, and reproduce the impossibility proof of Mayers (and, independently Lo and Chau). In doing so, we note the additional hypotheses used beyond those of a †-compact category, and suggest how to formulate some of them categorically. We prove in particular that one of them does not hold in **Rel**, so that the proof fails – indeed, we show that bit commitment according to our definition is possible there. In Chapter 4 we discuss the tension this causes with Clifton et al.'s characterisation theorem, and suggest how to remedy it.

We conclude with some potential extensions of this work.

Acknowledgements

I would first like to thank my supervisor Chris Heunen. I began pestering him nearly six months ago, and since then he has always been ready to answer questions, direct me towards useful reading, and keep me on track. Without him there would be no dissertation.

Thanks to all of the Quantum Group for providing such a pleasant environment in which to work; in particular, thanks to my fellow first-year students for always being ready to discuss work, and to my friends from MFoCS, Merton and elsewhere for always being ready not to.

Thanks to the organisers of the outstanding Oxford MFoCS course; I remember my excitement last year upon receiving an offer, and it has more than lived up to my expectations.

Finally, thanks to my parents, without whom I wouldn't be here today.

Contents

Chapter 0

Introduction

The modern study of quantum theory is not restricted to the von Neumann formalism: much work has been done in exploring generalisations of quantum mechanics and the physical features which uniquely pin it down. In the field of *categorical quantum mechanics*, a large proportion of the usual theory of quantum mechanics has been reformulated in terms of †-compact categories. In that language, we shift the focus of our analyses from states to processes acting on quantum systems, and study algorithms in terms of the interacting processes from which they are built.

Reformulation of quantum mechanics, however, is not limited to category theorists. Clifton, Bub, and Halvorson [2003] consider a physical theory to be specified by its algebra of observables, and investigate in particular the class of theories whose observables form a C^*-algebra. They locate quantum mechanics in this space of theories first through a set of "quantum" axioms, and then present an equivalent axiomatisation in terms of its information-theoretic properties.

One of these axioms states the non-existence of a particular type of cryptographic algorithm, known as bit commitment. It is easy to show that such an algorithm does not exist using only classical computers; within the last twenty years it has also been proved impossible even if quantum computers are available to us. The reason is roughly that although transmitting quantum systems can avoid the obstacles to commitment in classical mechanics, in permitting the use of quantum phenomena we open the doors to entanglement. By carefully entangling the quantum systems which they send, a dishonest party to a commitment protocol can unfairly manipulate the results. The existence of entangled states was in fact shown by Clifton et al. to be *equivalent* to the impossibility of bit commitment, using an argument from the world of C^*-algebras.

The CBH theorem demonstrates the relevance of bit commitment schemata to theories of quantum mechanics. It is natural, therefore, to ask about such schemata in

categorical quantum mechanics, and to investigate whether the known impossibility results also hold there. In this dissertation, we present an axiomatisation of bit commitment in a †-compact category and explore its interpretations. We'll see that although the usual impossibility proof is certainly valid in **FHilb**, it uses properties which are not captured by the †-compact structure, and hence does not lift to a proof in a general †-compact category.

This fact raises the question of whether bit commitment is indeed impossible in a †-compact category. Since we already know it is not possible in **FHilb**, we turn to our other example: the category **Rel** of sets and relations between them. To interpret bit commitment there we first need to give the semantics of our standard tools: states, classical structures, unitary morphisms, and so on. Once we do this, we demonstrate that in fact there do exist secure bit commitment protocols in **Rel**, and thus that they are not proved impossible by categorical quantum mechanics.

0.0.1 Plan of attack

This dissertation is split into two parts. We first review some of the background material that we'll use later, both from ordinary quantum mechanics and from the category-theoretic formulation. This is not meant to comprise a full introduction to the subject, and, in general, we expect readers to have come across the basic concepts of the field before. Nevertheless, we have tried to remain as self-contained as possible.

A reader acquainted with the basic concepts of quantum mechanics can safely skip Chapter 1, and with those of categorical quantum mechanics Chapter 2 – although readers may wish to glance through §2.5 for the semantics of the graphical calculus in **Rel**.

0.0.2 Novel Work

The main novel section of this thesis is the axiomatisation of bit commitment, and its interpretation in the standard models. Although the author has read many informal definitions, to the best of his knowledge none of them use a rigorous framework for the definition of a protocol.

Some new work was also needed in the interpretation of categorical quantum mechanics in CPM(**Rel**): Heunen and Boixo [2011] worked out classical structures but completely positive unitaries were new to the author, and many of the simple results about **Rel** needed some working out.

Part I

Background

Chapter 1

Classical and Quantum Algorithms

This chapter – an introduction to quantum mechanics as used in computer science – is based on Nielsen and Chuang [2000], the lecture notes [Coecke] and the Oxford University course *Quantum Computer Science*, among others. It describes the subset of quantum mechanics used in quantum information theory, in terms of states of Hilbert spaces and linear transformations between them.

1.1 Quantum Mechanics

As the components in modern computers shrink, quantum mechanics begins to take effect. If we accept this as a feature of computers then new abilities open up: just as a classical computer operates on bits using bit operations, a quantum computer operates on *qubits* using unitary operations.

Qubits are quantum systems that can be used to encode information. Just as with ordinary, classical data, a qubit may be placed in a state $|0\rangle$ or $|1\rangle$ to indicate the value of a bit. Unlike classical data, however, there are many pairs of states which can be so used, and some of these pairs are mutually exclusive: receiving a qubit does not convey useful information unless the choice of encoding scheme is also known.

Moreover, any linear combination of states is also a state, called their *superposition*. We think of a qubit in a superposition state as being in some sense in all of the states at the same time. This is distinct from being in some state with an associated probability, although that too is possible. A probability distribution over states (or, equivalently, a state with some probabilistic degree of uncertainty) is known as a *mixed state*.

Mixed states are particularly useful for cryptographic protocols, because there are distinct probability distributions over states that lead to the same mixed state. These

can be used to provide a certain amount of information to one party without revealing too much.

1.1.1 The Von Neumann formalism

The usual axiomatisation of quantum mechanics is known as the *von Neumann formalism*. Using it, we take the states of a quantum system to lie in a Hilbert space, and state-transforming operations to be unitary linear maps of Hilbert spaces.

1.1.1 Definition. A *Hilbert space* is a complex vector space equipped with a positive-definite conjugate-symmetric sesquilinear form $\langle - \,|\, - \rangle$ inducing a complete norm.

We adopt the bra-ket notation, writing states of a Hilbert space within angle brackets. At the very least, we would like the maps between Hilbert spaces to be *linear*; that is, to preserve the vector space structure. A function $X : \mathcal{H}_1 \to \mathcal{H}_2$ is called a *linear operator* if $X\big(a\phi + b\psi\big) = aX\phi + bX\psi$ for all $a, b \in \mathbb{C}$ and $\phi, \psi \in \mathcal{H}_1$.

Every Hilbert space $\mathcal{H}$ has a *dual* $\mathcal{H}^* \coloneqq \{\text{linear operators } \mathcal{H} \to \mathbb{C}\}$. It is not hard to verify that the dual of a Hilbert space is a Hilbert space, and that Hilbert spaces are canonically isomorphic to their double duals. In fact, this isomorphism arises just from the vector space structure: from the inner product we can do even better.

1.1.2 Proposition. *The operator $x \mapsto \langle - \,|\, x\rangle$ is an isometric anti-isomorphism between a Hilbert space $\mathcal{H}$ and its dual $\mathcal{H}^*$.*

1.1.3 Corollary. Every linear operator $X : \mathcal{H}_1 \to \mathcal{H}_2$ between Hilbert spaces has a unique *adjoint* linear operator $X^\dagger : \mathcal{H}_2 \to \mathcal{H}_1$ such that

$$\langle X^\dagger \psi \,|\, \phi\rangle = \langle \psi \,|\, X\phi\rangle.$$

We normally write $\langle \psi | X | \phi \rangle$ for the common value.

Proofs. Standard functional analysis. The Proposition is [Bollobás, 1990, §9, Corollary 10] and the Corollary is treated in Chapter 11, ibid. □

Now, we may regard states $\psi \in \mathcal{H}$ as linear operators $\mathbb{C} \to \mathcal{H} :: 1 \mapsto \psi$, so that states of $\mathcal{H}^*$, being linear maps $\mathcal{H} \to \mathbb{C}$, are adjoints of states of $\mathcal{H}$. We then see that the composite $\phi \circ \psi$ is exactly the inner product $\langle \phi \,|\, \psi\rangle$, justifying the Dirac notation $|\psi\rangle$ for elements of $\mathcal{H}$ and $\langle \psi|$ for elements of $\mathcal{H}^*$.

Some special linear operators will be of particular interest to us.

1.1.4 Definition. A linear operator X is

- *unitary* if its adjoint is its inverse; equivalently, if it is surjective and preserves the inner product
- *self-adjoint* if it is equal to its adjoint
- *idempotent* if $X^2 = X$
- a *projector* if it is self-adjoint and idempotent

1.1.5 Proposition. *Every self-adjoint operator X admits a unique* spectral decomposition

$$X = \sum_i a_i P_i,$$

where $a_i \in \mathbb{R}$, P_i are projectors and $P_i P_j = 0$ for all $i \neq j$.

Proof. Again, standard functional analysis, simplified by the fact that all operators are compact in finite dimensions. See [Bollobás, 1990, §14]. □

Von Neumann's approach to quantum mechanics is then summarised by

1.1.6 Postulate. *The states of a quantum system are given by rays in some Hilbert space, and transformations of the system are given by unitary operations on that Hilbert space.*

The physical interpretation of a measurement on some quantum system is not entirely clear. Nevertheless, it is easy to give a purely mathematical explanation.

1.1.7 Definition. A *measurement* on some quantum system is given by some self-adjoint operator H. If the spectral decomposition of H is $\sum_i a_i P_i$ then the *outcomes* of the measurement are the a_i. When the measurement is performed, one of the a_i is returned to the observer as an indication of the outcome, and the corresponding projector P_i is applied to the system. Each P_i occurs on measuring the (normalised) state ψ with probability $\langle \psi | P_i | \psi \rangle$.

Since projectors are idempotent, repeating a measurement will always give the same value.

1.1.2 Compound Systems

We define the *direct sum* of Hilbert space $\mathcal{H}_1$ and $\mathcal{H}_2$ as

$$\mathcal{H}_1 \oplus \mathcal{H}_2 \coloneqq \Big\{(\psi_1, \psi_2) : \psi_1 \in \mathcal{H}_1, \psi_2 \in \mathcal{H}_2\Big\}.$$

It is easy to show that $\mathcal{H}_1 \oplus \mathcal{H}_2$ is a Hilbert space of dimension $n + m$, since bases of $\mathcal{H}_1$ and $\mathcal{H}_2$ induce a canonical choice of basis on their direct sum.

The direct sum of Hilbert spaces represents the system of pairs of states – but, according to quantum mechanics, this is not the same as the *compound* system representing the states of both systems at the same time. Instead, the states of a compound system lie in the *tensor product space*

$$\mathcal{H}_1 \otimes \mathcal{H}_2 \coloneqq \left\{ \begin{pmatrix} c_{11} & \dots & c_{1m} \\ \vdots & \ddots & \vdots \\ c_{n1} & \dots & c_{nm} \end{pmatrix} \;\middle|\; c_{ij} \in \mathbb{C}, n = \dim \mathcal{H}_1, m = \dim \mathcal{H}_2 \right\}.$$

Intuitively, whereas the basis of the direct sum is given by the disjoint union of the bases of its components, the tensor product basis is given by the *cartesian product* of its component bases.

Note that we have given a basis-dependent construction of the tensor product space. This is not necessary: we can define the tensor product as the free Hilbert space on the direct sum, modulo bilinearity of $\otimes$. Alternatively, it can be characterised (up to isomorphism) as the universal space of bilinear maps; that is, the space through which any bilinear map from the direct sum factors.

The passage from the direct sum to the tensor product has far-reaching consequences. For example, it is easy to show that there are state of the tensor product space which are not products of states of the components – for example, if we write $\{e_1, e_2\}$ for a basis of $\mathbb{C}^2$, the *Bell state* $e_1 \otimes e_1 + e_2 \otimes e_2 \in \mathbb{C}^2 \otimes \mathbb{C}^2$ is not $e_i \otimes e_j$ for any i, j. We call such states *entangled*; their existence can be summed up by the slogan "the whole is more than the sum of its parts", since they mean that the joint state of a compound system cannot necessarily be given in terms of the marginal states of each of its component systems.

1.1.3 Density matrices

Sometimes we have quantum systems about whose state we only have partial information. Suppose we know that such a system is in state $|\psi_1\rangle$ with probability p_1,

$|\psi_2\rangle$ with probability p_2, and so on (which is *not* the same as certainly being in the linear superposition state $\sum p_i|\psi_i\rangle$). We say that this system has *density matrix*

$$\rho \coloneqq \sum_i p_i |\psi_i\rangle\langle\psi_i|.$$

By linear extension from the pure-state case, the expectation of a measurement is given by $\mathbb{E}A = \mathrm{tr}(\rho A) = \sum_i p_i \langle\psi_i|A|\psi_i\rangle$. The basic postulates of quantum mechanics can then be rephrased to work not with state vectors $|\psi\rangle$ but "one level up", with density matrices. These we call *mixed states*. Note that any "true" state $|\psi\rangle$ can be regarded as the mixed state $|\psi\rangle\langle\psi|$ – ψ with probability 1 – and hence states as we defined them embed into the space of mixed states via the map $\mathsf{Pure} : |\psi\rangle \mapsto |\psi\rangle\langle\psi|$. We say that a mixed state ρ is *pure* if it lies in the image of Pure.

It is important to note that the mixed state $|e_0\rangle\langle e_0| + |e_1\rangle\langle e_1|$ is *not* the same as the pure, entangled (non-normalised) state $|+\rangle \coloneqq |e_0\rangle + |e_1\rangle$. To see this, consider

$$|\psi\rangle \coloneqq \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 0 \end{pmatrix}$$

which equals $|e_0\rangle$ up to normalisation – so measuring $|\psi\rangle$ gives 0 with certainty. But the mixed state $|e_0\rangle\langle e_0| + |e_1\rangle\langle e_1|$ represents state $|e_0\rangle$ with 50% probability, and state $|e_1\rangle$ with 50% probability, so that $\left(\begin{smallmatrix} 1 & 1 \\ 1 & -1 \end{smallmatrix}\right)(|e_0\rangle\langle e_0| + |e_1\rangle\langle e_1|)$ represents either $|+\rangle$ or $|-\rangle \coloneqq |e_0\rangle - |e_1\rangle$ – thus measuring gives 0 or 1 with even odds.

Density matrices are of course themselves linear operators. We can characterise them as precisely the positive operators with unit trace. Evolution operators apply to density operators by conjugation, and the density matrix of a compound system is the tensor product of the density matrices of its components. It is easy to show that a density matrix ρ represents a pure state iff $\mathrm{tr}(\rho^2) = 1$.

An important property of density matrices is that they do not uniquely determine an ensemble of pure states – although the ensemble given by the eigensystem of a density operator ρ is certainly a *possible* decomposition of ρ, it need not be the only one. However, since they provide all the information necessary to determine a measurement, systems with the same density matrix cannot be physically distinguished. This property is key to the security analysis of some quantum algorithms.

Completely positive maps Since density matrices are of course matrices, they lie in the Hilbert space $\mathbb{C}^{n\times n}$. We say that a self-adjoint matrix is *positive (semidefinite)* if all of its eigenvalues are positive, and that a linear map $\Phi : \mathbb{C}^{n\times n} \to \mathbb{C}^{m\times m}$ is *positive* if it takes positivite matrices to positive matrices.

Positivity of linear maps between density matrices is insufficient for us, because it does not entail positivity when regarded as part of a larger system. That is, any such Φ naturally induces a map $I_k \otimes \Phi : \mathbb{C}^{k\times k} \otimes \mathbb{C}^{n\times n} \to \mathbb{C}^{k\times k} \otimes \mathbb{C}^{m\times m}$, and positivity of Φ does not entail positivity of $I_k \otimes \Phi$ for all k. (For example, the matrix transpose T is a positive map, but $I_2 \otimes T$ is not.) If each $I_k \otimes \Phi$ is positive, we say that Φ is a *completely positive* linear map. The intuition is that a completely positive operator takes mixed states to mixed states, even when regarded as part of an operation on a larger system. There is a well-known characterisation theorem for such maps.

1.1.8 Theorem (Choi [1975])**.** *A linear map Φ is completely positive iff the map*

$$(I_n \otimes \Phi)\left(\sum_{ij} E_{ij} \otimes E_{ij}\right) = \sum_{ij} E_{ij} \otimes \Phi(E_{ij})$$

is positive, where E_{ij} is the matrix with 1 in the ij-th entry and 0 elsewhere. □

1.2 Quantum Information

In quantum informatics we (almost) always limit ourselves to finite-dimensional Hilbert spaces; that is, we are only interested in measuring properties which can take finitely many distinct values. Indeed, we normally work with *qubits* – two-dimensional Hilbert spaces – and combinations thereof.

1.2.1 Definition. The *qubit* is the two-dimensional Hilbert space $\mathbb{C}^2$ (with the usual inner product), together with a chosen basis $|0\rangle, |1\rangle$ which we call the *computational basis*.

We use the computational basis to represent classical data: if a qubit is in state $|i\rangle$ then we consider it to represent the classical bit i. Unlike the space of classical states, however, the qubit admits all linear combinations of these states as well. For example, we may define the (non-normalised) states

$$|+\rangle \coloneqq (|0\rangle + |1\rangle)/2 \qquad\qquad |-\rangle \coloneqq (|0\rangle - |1\rangle)/2$$

A qubit in state $|+\rangle$ then represents a superposition of $|0\rangle$ and $|1\rangle$: it is in "both" states simultaneously. If we measure this state in the computational basis, by symmetry the result is either $|0\rangle$ or $|1\rangle$ with equal probabilities. So the density matrix representing the state of a qubit prepared in $|+\rangle$ and measured in the computational basis is

$$\tfrac{1}{2}|0\rangle\langle 0| + \tfrac{1}{2}|1\rangle\langle 1|.$$

Unitary maps are often known as *gates* in quantum information theory, to highlight the analogy with classical logic gates.

1.2.1 Quantum Data

Quantum data – information encoded in the state space of quantum systems – bears many similarities to classical data. But the two are by no means the same: many things which are possible with classical data are simply not possible with quantum. For instance, there is no way in general to copy an unknown quantum system; that is, there is no unitary transformation which acts as $|\psi\rangle \otimes |e\rangle \mapsto |\psi\rangle \otimes |\psi\rangle$ for all $|\psi\rangle$ and fixed $|e\rangle$. Similarly, there is no way to delete a quantum system. The proofs are easy and follow from linearity.

1.3 The Theory of Information

We consider for a moment at the characterisation theorem of Clifton, Bub, and Halvorson [2003], and its relevance to bit commitment. Clifton et al. consider a physical theory to be specified by its algebra of observables, and restrict this to comprise a C^*-algebra. (They also argue that this is not a major restriction, in light, for instance, of the fact that it suffices for "all physical theories [not necessarily quantum] that have been found to be empiricially successful to date".)

One first has to define what is meant by quantum mechanics! In other words, to prove that certain axioms entail the use of quantum mechanics we have to give a set of conditions which pin down what we consider to be quantum mechanics. Their axioms are:

(a) observables corresponding to distinct physical systems commute,

(b) not all observables commute, and

(c) there exists a physically realisable entangled state.

The final axiom is a little delicate, because the formal existence of entangled states follows from the C^*-algebraic structure itself. The point of the axiom is that the theoretical existence of entangled states does not necessarily guarantee that they can be physically prepared. Bub [2004] gives as an example Schrödinger's (now disproved) conjecture that entangled states might decay in practice as soon as their systems were physically separated, and thus that they could never occur in nature.

The main theorem of the paper gives three different (information-theoretic) axioms for quantum mechanics, and (almost) proves them in turn equivalent to their correspondents above.

1.3.1 Theorem. *In a C^*-algebraic theory:*

(α) *"no FTL information transfer" is equivalent to (a),*

(β) *"no broadcasting" is equivalent to (b), and*

(γ) *in the presence of (a) and (b), the impossibility of a secure bit commitment protocol is equivalent to (c)*

The idea of condition (α) is that if Alice and Bob each have a quantum system, then Alice, by performing some local measurement, cannot transfer any information to Bob – thus their measurement operators do not affect each other i.e. they commute.

Broadcasting – condition (β) – is a sort of generalisation of cloning to mixed states. Clifton et al. show that it implies cloning, that if any two pure states of a C^*-algebra can be cloned then they must be orthogonal, and finally that any algebra in which all pure states are orthogonal is commutative.

Finally, the intuition behind (γ) is that there are quantum commitment protocols which "nearly work": if we can trust Alice not to transmit entangled states, then a modification of the protocol given by Bennett and Brassard [1984] is easily proved secure. If no entangled states are available – but the other tools of quantum mechanics are – then secure bit commitment is possible. Conversely, if entangled states existed then Alice could use them to attack any commitment protocol.

The converse implication of (γ) – that physical nonlocality entails impossibility of bit commitment – was proved by Halvorson [2003], completing the proof of the CBH theorem. The issue was one that we'll face as well: the main tool used in the proof in Hilbert spaces is a particular decomposition demonstrated by Hughston et al. [1993], which required generalisation to arbitrary C^*-algebras.

Chapter 2

Categorical Quantum Mechanics

This section serves as a (necessarily brief) introduction to the field of categorical quantum mechanics as pioneered by Abramsky and Coecke [2004]. We have drawn concepts from the Oxford University course *Categorical Quantum Mechanics* (and its associated course notes), [Coecke, 2005] and many discussions with members of the Quantum Group. Although we aim for mathematical correctness everywhere, we have not tried to give an exhaustive account of the subject, restricting ourselves to the tools that we use later. In particular, we do not mention the formalisation of diagrams (in terms of digraphs), scalars, or any of the forms of complementarity. We refer the reader to any of the above-mentioned texts for more detail.

2.1 Graphical Calculus

A monoidal category is a universe of processes, that can be composed sequentially or in parallel. It admits a natural graphical calculus, in which we draw a process f with inputs $A_1, \ldots, A_n$ and outputs $B_1, \ldots, B_n$ as

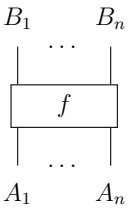

and sequential and parallel composition respectively as

Joyal and Street [1991] proved that the graphical calculus is sound and complete for monoidal categories: an equation between morphisms in monoidal categories is true iff it holds in the graphical language up to planar isotopy.

There is a host of soundness and completeness results for graphical languages which we shall henceforth ignore, comforted by the knowledge of their existence. Selinger [2011] provides a recent and thorough survey.

2.2 Additional Structure

Monoidal categories are very general universes in which to perform mathematics. They can be used to interpret physical systems and their evolution through time, data types and algorithms operating on them, algebraic structures and structure-preserving functions, or even logical propositions and proofs. We shall use them as a model for quantum mechanics. To do so, we'll need to add some more structures, which we describe below.

For the sake of brevity, we shall only mention the interpretations in passing. The reader may assume that any structure with a name lifted from Hilbert spaces is indeed a generalisation of the usual concept.

2.2.1 Symmetry

We should be able to swap the sides of diagrams; that is, for each X and Y there should be a process

$\sigma_{XY} :=$ 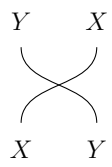

with the properties that

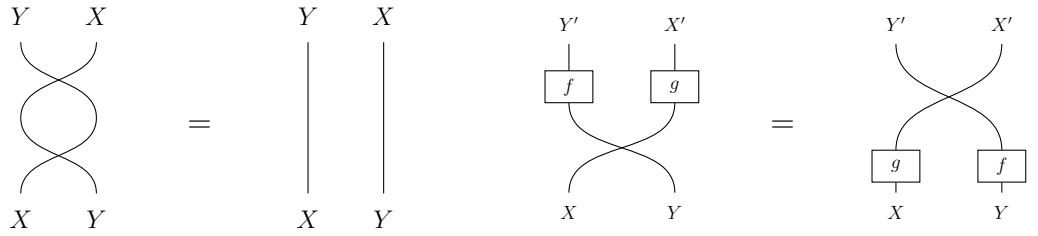

for each process $f : X \to X'$ and $g : Y \to Y'$.

2.2.2 The †-functor

The adjoint functor in Hilbert spaces generalises to a †-functor: a way to reverse morphisms. There is only one sensible way to represent this graphically: the diagram for $f^\dagger$ should be the diagram for f, reflected in a horizontal axis.

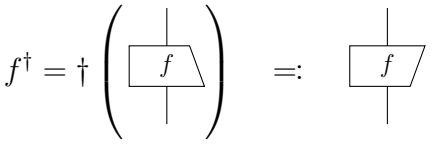

We've deliberately broken the symmetry of the boxes, so that f and $f^\dagger$ are distinct.

In formalising these diagrams, extra conditions on some of the "implementation details" in our categories are required. These manifest as further axioms on the category – but they are not drawn in the graphical language, so we don't have to worry about them here.

Sometimes the †-functor is called the *adjoint*.

2.2.3 Classical Structures

Coecke and Pavlovic [2006] introduced the notion of commutative Frobenius algebras as a way to model classical data. The axioms for a 'classical' object X are that it should be *copyable*

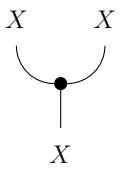

in a (co)commutative, (co)associative fashion:

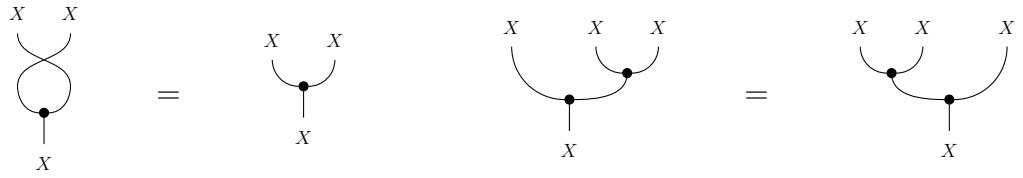

(Note that this is not a new piece of graphical notation: • is just a more concise way to write a box with no label.) It should also be *deletable*

so that deletion is a unit for copying:

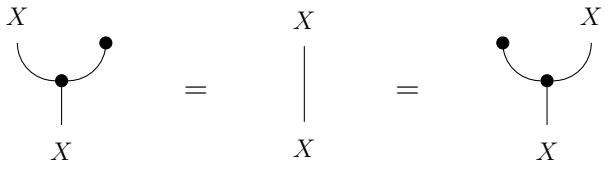

The adjoint of copying is *comparing*, and the adjoint of deletion is a *unit*.

Finally, there are two ways we may copy and compare data, and we would like them to be equal as well:

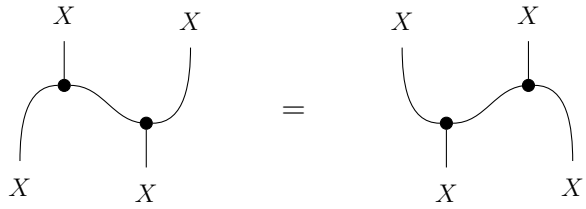

(This is called the *Frobenius* law.) We call an object X equipped with copying and comparing maps in this fashion a *classical structure*, since it represents an object holding classical data.

Using just the properties of the dots, we can prove the remarkable theorem

2.2.1 Theorem (Spider)**.** *Any morphism built from the above dots (copying, deleting, comparing, unit) is determined by its type* $X^{\otimes n} \to X^{\otimes m}$.

Proof idea. Any diagram built from • can be brought into a normal form by a particular (algorithmic) sequence of applications of the axioms. The theorem is well-known; details and further references can be found in Coecke and Duncan [2011].

The spider theorem permits us considerable laxity in how we draw the dots of the classical structure.

2.2.4 Duals

An interesting morphism that can be built out of a classical structure is

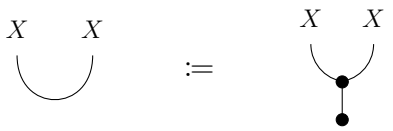

It is easy to show (from the axioms) that

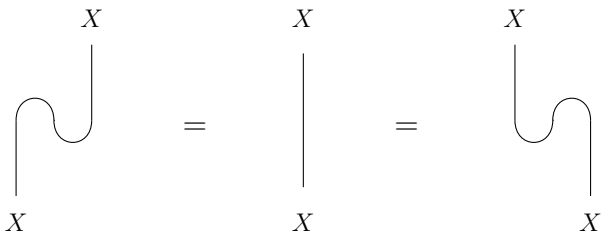

In general, any pair of morphisms $\eta : I \to A^* \otimes A$ and $\epsilon : A \otimes A^* \to I$ is called a *duality* if it satisfies the above "yank" laws. Dualities are again powerful because they allow us to ignore some specifics of the diagrams, namely, where the inputs and the outputs go. For instance, to specify a morphism f we can equivalently specify its name $\ulcorner f \urcorner$ or coname $\llcorner f \lrcorner$, defined as

respectively, since yanking either of these gives back f. (The difference between f and $\ulcorner f \urcorner$ is akin to the difference between an operator A and its matrix $[A]$.) We say the duality is a †-duality if $\eta = \epsilon^\dagger$.

It is important in general to distinguish between objects A and their duals A^*. To do so we can *direct* the wires of the calculus, adopting the convention that following a wire in reverse indicates the dual object. However, we only ever use the compact structure of the category implicitly – in the construction of $\mathrm{CPM}(\mathcal{C})$ – in contexts where it is clear in which direction the wires should go. For clarity, therefore, we elect not to indicate these directions explicitly, since we feel they clutter the diagrams with no additional benefit.

2.2.4.1 Duality as a Functor

Given a morphism $f : A \to B$ and its adjoint $f^\dagger$

we can form the two morphisms

known as f^* and f_* respectively. In particular, there are functors $-^*$ and $-_*$ acting by $A \mapsto A^*$ on objects and $f \mapsto f^*$ and f_* respectively. Note, though, that duals are only unique up to isomorphism, and that we must therefore fix a particular choice of duals upon choosing our category. Following convention, we'll take duals to "reverse the tensor product", so that $(A \otimes B)^* = B^* \otimes A^*$. This is permissible since the swap is an isomorphism, and will make some diagrams neater.

2.2.5 A Slogan

We are beginning to see a general principle when working with diagrams:

Only the topology matters!

Indeed, many of the axioms we take precisely fulfill the function of equating topologically-equal but diagrammatically-distinct diagrams: the Frobenius law allows us to exchange multiplication and comultiplication; the duality axioms mean lines may bend freely, and so on. Each of these trivial graphical manipulations corresponds to a non-trivial application of one of the algebraic axioms. The fact that we can perform such manipulations without having to hold all the axioms in our heads is one of the reasons that the graphical calculus is easier to work with than the algebraic formulation.

2.3 Semantics

We have now introduced enough graphical structure to do quantum mechanics. Before we proceed, we should briefly mention the interpretation of the language. The

semantics are well-known, and a detailed explanation can be found, for example, in the course notes.

The usual interpretation is in the †-compact category **FHilb**, with objects the finite-dimensional Hilbert spaces and morphisms the continuous linear maps. The tensor product becomes the usual tensor product of Hilbert spaces, the †-functor is given by adjunction of linear maps, and every object is dual to itself via the map

$$\eta_{\mathcal{H}} : \mathbb{C} \to \mathcal{H} \otimes \mathcal{H} :: 1 \mapsto \sum_i |i\rangle \otimes |i\rangle$$

(where the $|i\rangle$ comprise a basis of $\mathcal{H}$). Classical structures correspond bijectively to orthonormal bases [Coecke et al., 2008].

Algebraically, this generalises to the category of finitely-generated projective modules over any commutative ring.

The category **Rel** of sets and relations is another †-compact category, and one that turns out to be very useful as a toy model of quantum mechanics. There, the tensor product is the cartesian product, the †-functor is given by relational converse, and every object is dual to itself via the relation

$$\eta_A : I \to A \times A :: * \overset{A}{\sim} (a, a) \text{ for all } a \in A.$$

Classical structures on A correspond bijectively to abelian groupoids whose morphism set is A [Heunen et al., to appear]. We can think of the morphisms of **Rel** as given by boolean matrices, just as the morphisms of **FHilb** are given by complex matrices, so that **Rel** is in some sense a "possibilistic" quantum theory.

The category **Set** of sets and functions is neither a †-category nor a compact category. (The function $\emptyset : \emptyset \to \{*\}$ admits no adjoint, and since the tensor unit is terminal in **Set**, the existence of duals would imply that all conames, and hence all functions, were equal.) Since we intuitively do maths in **Set**, this is one reason that quantum mechanics can often seem counterintuitive.

2.4 Mixed States

After defining the von Neumann formalism in §1.1.1, we proceeded to "move up a level" to work not with states of Hilbert spaces but with density matrices: positive unit-trace operators representing partial information about a system. This we can also do in the graphical formalism, using a construction due to Selinger [2007] which expands a †-compact category (of "pure states") into a †-compact supercategory (of "mixed states").

Just as before, instead of considering states as (normalised) elements of Hilbert spaces we move up to their induced density operators:

The normalisation condition says that

i.e. that $\psi^\dagger$ is a retraction of ψ. Note that $\psi \circ \psi^\dagger$ is by definition positive, so it makes sense to abstract away to the following definition.

2.4.1 Definition. A *mixed state* of A is a positive morphism $\rho : A \to A$.

Just as a unitary operator is a process which takes pure states to pure states, a *completely positive* operator will be one taking mixed states to mixed states. We can use the following theorem to lift this to a definition, since by Choi's theorem on completely positive maps, the first criterion is equivalent to the usual definition of complete positivity in Hilbert spaces.

2.4.2 Theorem (Stinespring Dilation)**.** *The following are equivalent:*

- *the morphism*

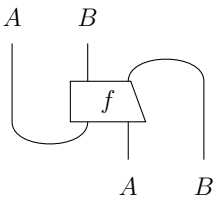

  is positive

- *there is a morphism $g : A \to C \otimes B$ such that*

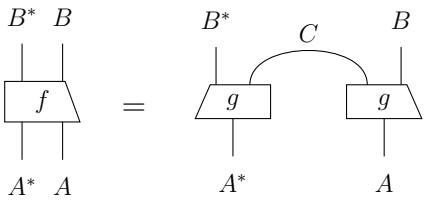

Proof. See for example Selinger [2007]. □

We call such morphisms *completely positive*.

2.4.1 Selinger's CPM construction

The completely positive morphism in fact form a †-compact category.

2.4.3 Definition. Let $\mathcal{C}$ be a †-compact category. The category CPM($\mathcal{C}$) has

- objects the objects of $\mathcal{C}$, and
- arrows $A \to B$ the completely positive $\mathcal{C}$-arrows $A^* \otimes A \to B^* \otimes B$.

2.4.4 Theorem. CPM($\mathcal{C}$) *is a* †*-compact category, with*

- *composition,* † *and* $\otimes$ *on objects as in* $\mathcal{C}$*,*
- $\otimes$ *on morphisms given by*

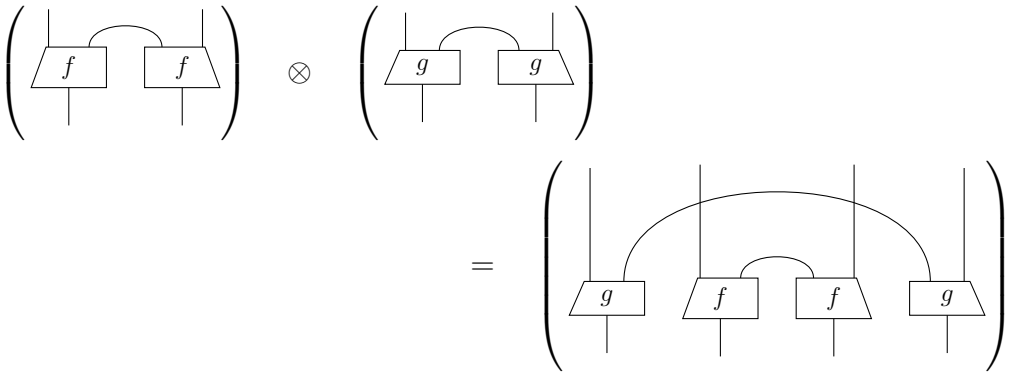

- $\sigma_{AB} : A \otimes B \to B \otimes A$ *given by*

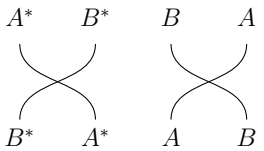

- $\epsilon_A : A^* \otimes A \to I$ *given by*

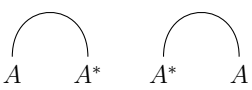

2.4.5 Proposition (Heunen and Boixo [2011])**.** WP : $\mathcal{C} \to$ CPM($\mathcal{C}$) :: $f \mapsto f_* \otimes f$ *is a symmetric strict monoidal functor that preserves* †*.*

The functor WP is nearly faithful: if $\mathrm{WP}(f) = \mathrm{WP}(g)$ then $f = \phi \bullet g$ for some scalar $\phi : I \to I$. For this reason we often say that WP "kills phases".

In fact, there is a functor $\mathrm{CPM}(\mathcal{C}) \to \mathcal{C}$, given by interpreting the CPM-diagrams in $\mathcal{C}$. That is, G sends the object X_{CPM} of $\mathrm{CPM}(\mathcal{C})$ to the object $X^* \otimes X$ of $\mathcal{C}$, and the morphism $(1 \otimes \epsilon \otimes 1) \circ (g_* \otimes g)$ of $\mathrm{CPM}(\mathcal{C})$ to the same morphism in $\mathcal{C}$. It is not too hard [Heunen and Boixo, 2011] to show that this functor is in fact symmetric strong monoidal and preserves †.

2.4.2 Environment Structures

In this section we paraphrase Coecke and Perdrix [2010].

The graphical calculus for mixed states is a little irritating because we have to draw all of the morphisms twice (precisely as we have to write ψ twice in $|\psi\rangle\langle\psi|$). It turns out that adding a *partial trace* to a category is roughly equivalent to working with its category of completely positive maps. More formally, let $\widehat{\mathcal{C}}$ be a †-compact category and $\mathcal{C}^{\mathrm{pure}}$ a †-compact subcategory with the same objects and inheriting the monoidal †-structure. (We think of $\mathcal{C}^{\mathrm{pure}}$ as the image of $\mathcal{C}$ in $\mathrm{CPM}(\mathcal{C})$ under WP.)

2.4.6 Definition (Coecke and Perdrix)**.** An *environment structure* for $\langle \mathcal{C}^{\mathrm{pure}}, \widehat{\mathcal{C}} \rangle$ is a chosen costate $\top_A$ of each object A, denoted

$\top_A :=$ 

such that

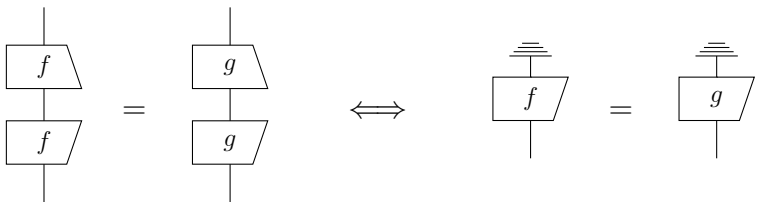

(2.1)

(2.2)

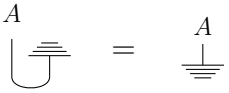

(2.3)

$\top_A$ is read as "tracing out".

If, in addition, every morphism in $\widehat{\mathcal{C}}$ can be constructed by tracing out part of the codomain of a morphism in $\mathcal{C}^{\mathrm{pure}}$, we say the environment structure has *purification*.

Working with an environment structure is equivalent to applying the CPM construction in the following sense:

2.4.7 Theorem. *Let $\mathcal{C}$ be a †-compact category.*

(i) If $\mathcal{C}$ has an environment structure then there is an invertible fully-faithful identity-on-objects monoidal functor $\xi : \widehat{\mathcal{C}} \to \mathrm{CPM}(\mathcal{C})$.

(ii) The image of $\mathcal{C} \hookrightarrow \mathrm{CPM}(\mathcal{C})$ has an environment structure induced by the cap ϵ.

Proof. See e.g., Coecke [2008]. □

2.4.3 Lifting Structure to CPM

If we wish to work with mixed states and density operators – that is, if we wish to work in $\mathrm{CPM}(\mathcal{C})$ for some †-compact category $\mathcal{C}$ – then we must first lift the structures that we used in $\mathcal{C}$ to $\mathrm{CPM}(\mathcal{C})$.

We have already seen that †-compactness lifts through CPM. The next theorem follows from the fact that WP is strict monoidal, but we write it out explicitly anyway.

2.4.8 Theorem. *Classical structures in $\mathcal{C}$ induce classical structures in* $\mathrm{CPM}(\mathcal{C})$.

Proof. Via the map

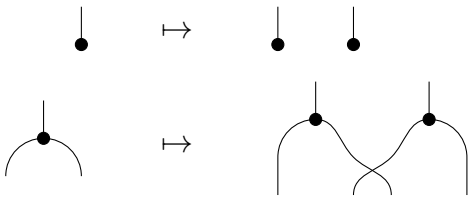

It is easy to check that this forms a classical structure in $\mathrm{CPM}(\mathcal{C})$. □

2.4.9 Corollary. Copyable states in $\mathcal{C}$ induce copyable states in $\mathrm{CPM}(\mathcal{C})$.

We say that such a classical structure in $\mathrm{CPM}(\mathcal{C})$ is *canonical*. It is not *a priori* clear that every classical structure must be canonical; that is, that the only classical structures in a category $\mathrm{CPM}(\mathcal{C})$ are induced by classical structures in $\mathcal{C}$. We shall see in §2.5.5 that every classical structure in CPM(**Rel**) is canonical, and Heunen and Boixo conjecture that this is also the case in **FHilb**.

In a similar vein, it is easy to see that a unitary morphism U in $\mathcal{C}$ induces a unitary morphism $\mathrm{WP}(U) = U_* \otimes U$ in $\mathrm{CPM}(\mathcal{C})$:

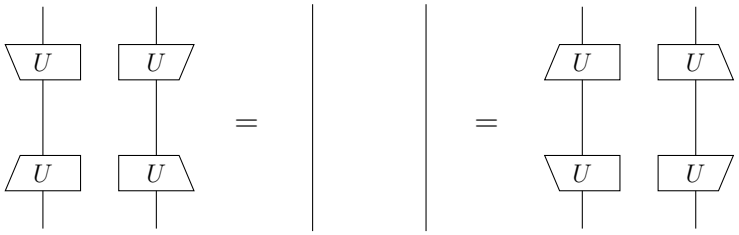

As before, let us say that a unitary morphism in $\mathrm{CPM}(\mathcal{C})$ is *canonical* if it arises as $f_* \otimes f$ for some unitary f in $\mathcal{C}$.

2.4.3.1 Non-canonical Structures in $\mathrm{CPM}(\mathcal{C})$

The fact that $\mathcal{C}$ and $\mathrm{CPM}(\mathcal{C})$ are distinct categories opens up the possibility of quantum structures existing in the latter which did not originate from the former. It is not clear how to interpret such *non-canonical* features physically; this is a topic of current research.

We'll see in §2.5 that all classical structures and all unitaries in $\mathrm{CPM}(\mathbf{Rel})$ are canonical. It follows from the work of Nayak and Sen [2007] that all trace-preserving unitaries in $\mathrm{CPM}(\mathbf{FHilb})$ are canonical, and Heunen and Boixo [2011] conjecture that all classical structures are as well.

We shall consider physical measurements to be induced by classical structures in the base category $\mathcal{C}$ – *not* arbitrary classical structures in $\mathrm{CPM}(\mathcal{C})$. Since our prime example is taken from **Rel**, this point is moot – but it is important to make the distinction if we want to interpret the definition in a general †-compact category.

2.4.4 Classical Channels

We can consider the partial trace operation as "entangling with the environment", or "broadcasting". A quantum channel which we entangle with the environment is then just a classical channel.

2.4.10 Definition (Coecke and Perdrix)**.** Let X be a classical structure. The *classical channel of type* X is the morphism $C_X : X \to X$ defined by

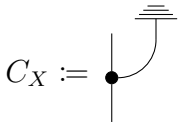

We say a copyable state ψ is *normalised* if

Classical channels are so called because they transmit normalised copyable states, called *(pure) classical states*:

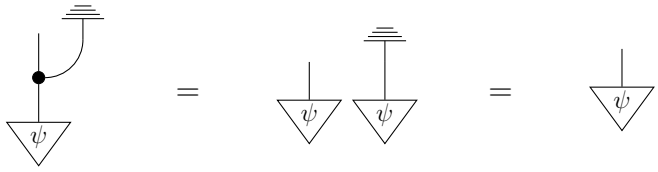

More generally, a *(mixed) classical state* of a classical channel C is a state transmitted by C.

A *measurement* takes a quantum channel as input and provides a quantum channel and a classical channel as output:

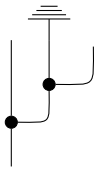

If in addition we trace out the quantum system, we get a *destructive measurement*:

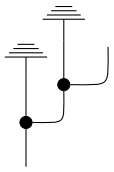

But applying the spider theorem lets us rewrite this to just the classical channel, so that passing quantum data through a classical channel is precisely performing a measurement.

It is easy to check that $\bot := \top^\dagger$ is a mixed classical point of *any* classical structure. Following Coecke [2008], we call it the *maximally-mixed* state.

2.5 A Diversion into Rel

One of the benefits of abstracting away from Hilbert spaces to †-compact categories is that we can now interpret all of the "quantum" structures we know and love in different contexts. **Rel** is a particularly interesting toy category, because of its

similarity to **Set**: although it is not too challenging to grasp relations intuitively, they still provide a model for such structures.

We have seen a few examples of graphical structures in **Rel** already. A convenient feature is that the †-functor is relational converse, which can be explicitly constructed in a way that the adjoint of Hilbert spaces cannot.

2.5.1 Classical Structures

Classical structures are given by abelian groupoids [Pavlovic, 2009, Heunen et al., to appear]. These are strange beasts – commutativity of composition is not a common property. But it is clear that it must follow from symmetry of m:

$$\begin{aligned}\sigma \circ m^\dagger &= \{\langle (h,g),(g,h)\rangle : g,h \in G\} \circ \{\langle gh,(g,h)\rangle\} \\ &= \{\langle gh,(h,g)\rangle : g,h \in G\} \\ = m^\dagger &= \{\langle gh,(g,h)\rangle : g,h \in G\}\end{aligned}$$

In particular, if $x \xrightarrow{h} y \xrightarrow{g} z$ then

$$\begin{aligned}&\langle gh,(h,g) \in \sigma \circ m^\dagger \\ \implies &\langle gh,(h,g) \in m^\dagger && (m^\dagger = \sigma \circ m^\dagger) \\ \implies &h \circ g \text{ exists and } h \circ g = g \circ h \\ \implies &x = y = z\end{aligned}$$

The inducing groupoids therefore look like

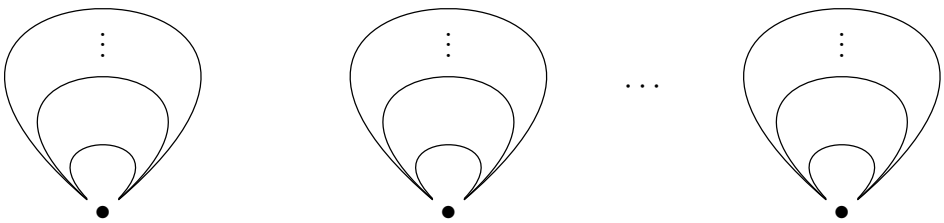

i.e. like disjoint unions of groups, with group structures chosen independently on the sets of some partition of X.

As a concrete example, consider the two-element set $\mathbb{I} \coloneqq \{0,1\}$. We can consider this as the morphism set of the trivial abelian groupoid

The composition rule is defined in the only possible way: $0 \circ 0 = 0$ and $1 \circ 1 = 1$; this yields the "copying" classical structure

$$\delta : \mathbb{I} \to \mathbb{I} \otimes \mathbb{I} :: \begin{cases} 0 \sim (0,0) \\ 1 \sim (1,1) \end{cases} \qquad \text{and} \quad \epsilon : \mathbb{I} \to I :: \begin{cases} 0 \sim * \\ 1 \sim * \end{cases}$$

(Indeed, any object X of **Rel** has a classical structure given by the trivial abelian groupoid $\bigsqcup_{x \in X} \{*\}$, whose multiplication is given by copying $x \sim (x, x)$ and for which every element is a unit.)

There is another abelian groupoid with two morphisms, namely $\mathbb{Z}^2$:

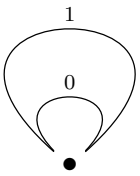

This gives rise to a different – 'switching' – classical structure on $\mathbb{I}$:

$$\delta : \mathbb{I} \to \mathbb{I} \otimes \mathbb{I} :: \begin{cases} 0 \sim \{(0,0),(1,1)\} \\ 1 \sim \{(0,1),(1,0)\} \end{cases} \qquad \text{and} \quad \epsilon : \mathbb{I} \to I :: 0 \sim *$$

2.5.2 Classical points

The abelian groupoid structure on a set induces a partition, by compatibility with respect to composition. We can write down the copyable and unbiased states of a classical structure with respect to this partition.

2.5.1 Lemma. *Let (X, m, u) be a classical structure in which X is partitioned by composability with respect to m into disjoint abelian groups as*

$$\begin{array}{llll} \{x_{11}, & x_{12}, & \ldots, & x_{1n}\} \bigsqcup \\ \{x_{21}, & x_{22}, & \ldots, & x_{2n}\} \bigsqcup \\ & \vdots & & \\ \{x_{n1}, & x_{n2}, & \ldots, & x_{nn}\} \bigsqcup \end{array}$$

Then the classical points for X are exactly the rows X_i, and the unbiased points are the sets containing precisely one element from each X_i.

Proof. See Evans, Duncan, Lang, and Panangaden [2009]. □

2.5.3 Complete Positivity

A relation R is positive iff it is $S^\dagger \circ S$ for some S. Let R be positive and aRb, so that there is some c with $aScS^\dagger b$. Then $bScS^\dagger a$, so R is *symmetric*. Similarly, R must be *semi-reflexive*: if aRb then $aScS^\dagger b$, so $aScS^\dagger a$, so aRa. In fact:

2.5.2 Proposition. *Positive morphisms in* **Rel** *are precisely semi-reflexive, symmetric relations.*

Proof. One direction is done. Conversely, let $R : X \to X$ be semi-reflexive and symmetric, and – remembering that R itself is a set of pairs – define $S : X \to R$ by

$$x \overset{S}{\sim} (x, y) \text{ and } x \overset{S}{\sim} (y, x) \quad \forall x, y \in R.$$

It is clear that S is a relation, and $S^\dagger = \{\big((x, y), x\big) : (x, y) \in R\}$ so that

$$S^\dagger \circ S = \big\{(x, y) : \exists (z, w) \in R : xS(z, w)S^\dagger y\big\}.$$

We prove that this relation is in fact equal to R.

- If $x(S^\dagger \circ S)y$, there is $(z, w) \in R$ with $xS(z, w)S^\dagger y$. Since $xS(z, w)$ we have either $x = z$ or $x = w$; since R is symmetric we may take wlog $x = z$. Since $yS(z, w)$ we have either $y = z$ or $y = w$. If $y = w$ then $(x, y) = (z, w) \in R$ and we are done. Otherwise, by semi-symmetry $(z, w) \in R \implies (z, z) \in R$, and $(x, y) = (z, z) \in R$. Either way, xRy.
- If xRy then $(x, y) \in R$, so $xS(x, y)S^\dagger y$.

This demonstrates containment in both directions, so $R = S^\dagger \circ S$ is positive. □

2.5.3 Corollary. A relation $R : A \otimes A \to B \otimes B$ in CPM(**Rel**) is completely positive just when for all a, b, c, d

$$\begin{aligned} (a, b)R(c, d) &\iff (b, a)R(d, c) \quad \text{and} \\ (a, b)R(c, d) &\implies (b, b)R(d, d) \end{aligned}$$

Given a relation R, we form the associated morphism in CPM(**Rel**) by simply doubling $R \mapsto R \otimes R$. So the partial trace

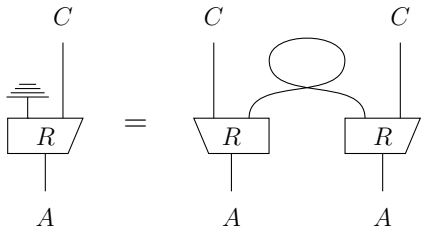

$$= \Big\{\big((a, a), (c, c)\big) : (\exists b)a \overset{R}{\sim} (b, c)\Big\},$$

as we might hope, just "forgets" one of the coordinates of the output tuple: it is (related to[1]) the canonical projection π_1 arising from the product structure of $\otimes$.

2.5.4 Classical Channels

The explicit definition of the partial trace lets us work out classical channels in **Rel**. Recall that $\bullet$ divides the underlying set into partitions by composability.

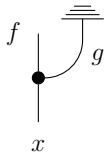

$$= \{(x, f) : (\exists g)\ fg = x\}$$

$$= \{(x, f) : x \text{ and } f \text{ are in the same partition}\}$$

Indeed, every f in the same partition as x admits a g with $fg = x$ (namely, $g \coloneqq f^{-1}x$); moreover, if $fg = x$ then by definition f is in the same partition as x.

The copyable states of a classical structure are exactly the sets of its partition. We now see that the *mixed classical states* – the states transmitted by C – are precisely disjoint unions (or coproducts) thereof. The maximally-mixed state $\bot_A = \{(*, a) : a \in A\}$ is certainly one of these, as are all of the copyable states.

2.5.5 CPM(**Rel**)

Mixed states in **Rel** – states in the category CPM(**Rel**) – are particularly pleasant to work with, due in part to the two following propositions.

2.5.4 Proposition. *In* CPM(**Rel**)*, every classical structure is canonical.*

Proof [Heunen and Boixo, 2011]. Since classical structures in **Rel** are given by multiplication of abelian groupoids, the condition for complete positivity in **Rel** becomes

$$(a, b) \cdot (c, d) = (e, f) \iff (c, d) \cdot (a, b) = (f, e)$$
$$(a, b) \cdot (c, d) = (e, f) \implies (c, d) \cdot (c, d) = (f, f)$$

Taking $(a, b) = 0$ implies that $(e, f) = (c, d) = (c, d) \cdot (a, b) = (f, e)$, so that $e = f$ – that is, the classical structure is the diagonal copy of a classical structure in **Rel**. This is precisely the condition that it be canonical. □

[1] Of course, it is not π_1 "on the nose": we must include the left unitor to remove the extra tensor unit, and take into account the doubling action of CPM. Here we mean that $B \otimes C \xrightarrow{\top \otimes 1} I \otimes C \xrightarrow{\lambda_c^{-1}} C$ is the image of $B \otimes C \xrightarrow{\pi_1} C$ under CPM.

2.5.5 Proposition. *In* CPM(**Rel**)*, every unitary is canonical.*

Proof. Recall from Theorem 2.4.2 that a completely positive morphism $R : A^* \otimes A \to B^* \otimes B$ is a morphism

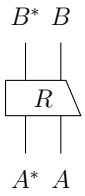

such that $(\epsilon \otimes 1 \otimes 1) \circ (1 \otimes R \otimes 1) \circ (1 \otimes 1 \otimes \eta)$ is positive. In particular, since the †-functors of CPM(**Rel**) and **Rel** coincide, a unitary morphism in CPM is also unitary in **Rel**. Unitary relations are precisely bijections, so that R is a bijection satisfying for all a, b, c, d

$$R(a,b) = (c,d) \iff R(b,a) = (d,c) \tag{2.4}$$

$$R(a,b) = (c,d) \implies R(b,b) = (d,d) \tag{2.5}$$

For any b, let $(x,y) = R(b,b)$. Then $(y,x) = R(b,b) = (x,y)$ by 2.4, so $x = y$. Writing $\pi_1 : (x,y) \mapsto x$ and $\pi_2 : (x,y) \mapsto y$, define $S(x)$ to be the common value of $\pi_1 \circ R(x,x)$ and $\pi_2 \circ R(x,x)$.

We claim $R(a,b) = \big(S(a), S(b)\big)$; that is, $R = S \times S$. Indeed, suppose $R(a,b) = (c,d)$. Then $\big(S(b), S(b)\big) = R(b,b) = (d,d)$ by definition and 2.5, so $d = S(b)$. By 2.4 $R(b,a) = (d,c)$, so that $\big(S(a), S(a)\big) = R(a,a) = (c,c)$ and $c = S(a)$. Thus $R(a,b) = (c,d) = \big(S(a), S(b)\big)$, proving the claim.

Finally, we must show that S is a bijection i.e. a unitary of **Rel**. It is clear that it is a function. Given a and b such that $S(a) = S(b)$, we have that $R(a,a) = \big(S(a), (S,a)\big) = \big(S(b), S(b)\big) = R(b,b)$, so $(a,a) = (b,b)$ and $a = b$; hence S is injective. Given a, since R is surjective there is (x,y) such that $(a,a) = R(x,y) = \big(S(x), S(y)\big)$, so $x = y$ and $S(x) = a$; hence S is surjective.

Since S is an injective, surjective function with $R = S \times S$, the proposition is proved. □

2.5.6 CPM(**FHilb**)

Similar results seem to hold for **FHilb**, though we shan't need them.

2.5.6 Theorem. *All trace-preserving unitaries in* CPM(**FHilb**) *are canonical.*

Proof. Nayak and Sen [2007] show that every completely-positive trace-preserving (CPTP) map with CPTP inverse is of the form $U(- \otimes \omega)U^\dagger$ for some unitary U and ancilla ω; the ancilla is used to introduce the necessary extra dimensions. Since a unitary map must be between spaces of the same dimension[2], their result specialises to state that every completely positive, trace preserving, unitary map is given by conjugation by some unitary: $\rho \mapsto U\rho U^\dagger$. By map-state duality, this says that every CPTP unitary in CPM(**FHilb**) is of the form $U_* \otimes U$ for some U; that is, that every unitary is canonical. □

Nayak and Sen use the operator-sum representation of a quantum operation $\mathcal{E}(\rho) = \sum_k E_k^\dagger \rho E_k$, which requires that $\sum_k E_k^\dagger E_k = I$; that is, that $\mathcal{E}$ preserves trace. The proof therefore does not directly lift to the case $\sum_k E_k^\dagger E_k \leq I$ of non-trace-preserving unitaries. It would be interesting to investigate whether the claim is indeed true in that case.

It is not yet known whether an analogue of the result on classical structures holds, although Heunen and Boixo conjecture it to be the case.

2.5.7 Conjecture. *All classical structures in* **FHilb** *are canonical.*

We have now prepared all of the background and tools that we shall need in order to formulate and prove results about commitment protocols. Without further ado, we move on.

[2] U is invertible, so has trivial kernel i.e full rank and the rank-nullity theorem applies.

Part II

Commitment

Chapter 3

Commitment Algorithms

Commitment algorithms are a cryptographic primitive akin to a "delayed transfer", motivated by the following scenario.

> Alice is a financial speculator with a secret trading algorithm. She wishes to sell Bob market predictions, but Bob is unconvinced as to her reliability. Alice presents her prediction history as evidence of her accuracy, but Bob points out that since those predictions are not useful since she could have invented them *ex post facto*. He suggests instead that Alice make him a new prediction about tomorrow's market data, but Alice refuses to reveal that information for free.
>
> Alice and Bob consult the Guru, who tells Alice to buy a safe, lock some predictions for tomorrow inside and give it to Bob. The next day, Alice gives Bob the combination to the safe, proving that her prediction was reliable. Alice and Bob go on to establish a profitable contract and begin the next housing bubble.

What has the Guru invented? Using the safe, Alice can send a message to Bob which he cannot read until Alice later sends him a key to open it. In other words, Alice can generate from any $\mathsf{message}$ a pair of tokens $\langle \mathsf{safe}, \mathsf{key} \rangle$ such that

(i) knowing safe and key reveals $\mathsf{message}$,

(ii) knowing safe without key provides no information about $\mathsf{message}$, and

(iii) the only message safe unlocks to is the original $\mathsf{message}$.

A protocol is *sound* if it satisfies (i), *concealing* if (ii) and *binding* if (iii).

Such safes (are believed to) exist as classical cryptographic algorithms. But the analogy to physical safes holds further: Alice's data is only secure as long as we

assume Bob cannot open the safe, and this assumption is based on a raft of heuristics about the hardness of steel, the availability of large drills, and so on. Indeed, if Bob had enough time he could simply try every combination and we can certainly not guard against this! We say that a protocol is *conditionally secure* if it is secure under assumptions on the computational power of an adversary.

Bennett and Brassard [1984], among others, showed that quantum mechanical systems can give rise to cryptographic algorithms that are secure even without the assumptions on Bob's resources. Indeed, the protocol for quantum key distribution (and hence secure communication) was provably secure using only the laws of physics. Communicating using such a protocol means that we do not have to worry about the computational tractability of our algorithms, or the abilities of an unknown adversary.

One might hope that a similar algorithm exists for bit commitment. Sadly, it was shown by Mayers [1997] (and independently by Lo and Chau [1997]) that it does not: although the same technique as key distribution does give rise to a potential bit commitment algorithm, it is vulnerable to an attack based on entanglement. Clifton et al. [2003] then demonstrated that – under some assumptions – the impossibility of secure bit commitment is *equivalent* to the physical existence of entangled states.

We will investigate unconditionally secure bit commitment in the context of †-compact categories. Remember that for unconditional security we have to assume that Bob knows all the details of the algorithm and that he has unlimited computational power available to him. In particular, we assume that he can implement any unitary transformation on any quantum system.

3.1 A First Try

It is worth considering a few initial ideas to see whether they work.

- Alice could simply send Bob her bit as the message.

  This is certainly sound and binding. But it is obviously not concealing.

- Alice could send Bob gibberish, and then send her bit as the key.

  This is again sound, and clearly concealing. It is not binding: Alice can invent whatever she likes to send to Bob and he cannot confirm that it was her original data.

- Alice could encrypt her bit with an information-theoretically secure cipher, and send that as the safe. The key would be the decryption key.

Shannon [1949] proved that the $\oplus$ cipher which simply takes the bitwise XOR of the message and the key is concealing: knowing the message provides no information about the key. It is not binding, though: Alice can easily invent a new key that will decrypt an arbitrary ciphertext to an arbitrary plaintext. (Indeed, this is the foundation of its security: since any ciphertext could have come from any plaintext with equal probability, there can be no reason to prefer any particular decryption.)

- Alice could encrypt her bit with an asymmetric cipher such as RSA, and send that as the safe. Again, the key would be the decryption key.

  This time we are getting somewhere: binding asymmetric ciphers do exist, in the sense that finding a key that decrypts a ciphertext to a given plaintext is as hard as breaking the cipher. Unfortunately, the problem here is that the security is no longer information-theoretic, just conditional: Bob could certainly decrypt the message given enough time, and hence reveal Alice's commitment prematurely.

- Alice can choose a number $0 \leq x \leq p-1$ as her commitment, pick a generator g of $\mathbb{Z}_p$ and send $c \coloneqq g^x$ as the safe. The key is x.

  This is another conditionally-secure protocol, this time relying on the intractability of the discrete logarithm problem. It is sound and binding but only conditionally concealing.

- Alice could encode her bit as the basis of a qubit storing a random bit, and send that to Bob. The key will be her bit.

  The data stored by the qubit must be random: since we assume that Bob knows all the details of the protocol, it cannot be fixed in advance.

  This protocol is certainly sound, and can be verified to be concealing. To do so, we compute the density matrices of the ensembles caused by Alice's random choice, and note that they are independent of her choice of bit. It certainly appears to be binding: if Alice sends a single qubit to Bob she has no way to access it later. But the existence of entanglement means that we must examine this claim more carefully!

  Two facts are important to note: firstly, the protocol takes random input as well as the committed bit, and secondly, this means that its success is probabilistic.

We can therefore not achieve unconditional security according to our definition, although we can come arbitrarily close by iterating the protocol.

3.1.1 Classical Bit Commitment

It is intuitive that unconditionally secure classical bit commitment is impossible,

> essentially because Alice can send (encrypted) information to Bob that guarantees the truth of an exclusive classical disjunction (equivalent to her commitment to a 0 or a 1) only if the information is biased towards one of the alternative disjuncts (because a classical exclusive disjunction is true if and only if one of the disjuncts is true and the other false). No principle of classical mechanics precludes Bob from extracting this information. So the security of the protocol cannot be unconditional and can only depend on issues of computational complexity. (Bub [2004])

To formalise this intuition would require us set up a framework for classical cryptography, with communication channels, messages and so forth. This would take us on rather a tangent from the main body of the dissertation.

Conditionally secure classical bit commitment is certainly possible – we've given some examples above – and is used as a cryptographic primitive in various algorithms. For example, it can be used in zero-knowledge proofs to allow the prover to submit all her information upfront, and only choose what to reveal later.

3.1.2 BB84

The latter protocol mentioned above was first proposed by Bennett and Brassard [1984]. They pointed out a flaw: if we allow Alice to cheat, we may not assume that she honestly sends as her message a qubit encoded in a basis state. If she instead sends part of an entangled system, then she can "access" Bob's qubit even after she has sent it to him, since by applying a local unitary to her ancilla she changes the joint state of the system. Although this does not inherently allow her to change the state of Bob's qubit directly – that would entail faster-than-light information transfer! – it does mean that the protocol is not obviously binding.

In this case, Alice can indeed cheat. To do so, suppose the two bases she might use are $\{|0\rangle, |1\rangle\}$ and $\{|+\rangle, |-\rangle\}$. Then she can prepare the joint state

$$|\Phi\rangle \coloneqq |00\rangle + |11\rangle = |++\rangle + |--\rangle$$

and send one of its component systems as her commitment. When she later wishes to unveil it, Alice can measure her ancilla system in either basis and produce the result as an unveil token Bob will accept.

A priori this does not seem an insurmountable flaw with quantum protocols in general – just with this particular simple one. In general, Alice cannot change the state of Bob's qubit in a controlled fashion, since this would amount to faster-than-light information transfer; she can only cause the collapse of the state, and this is a random process. The attack on BB84 therefore seems to rely on poor design of the protocol, not an intrinsic property of quantum mechanics. Could there be a more complicated scheme which avoids the entanglement attack Alice can use above?

3.2 Formalisation

Let's formally define what we mean by a commitment scheme. We present what we think is the right definition, and follow it with a discussion as to why.

3.2.1 Definition. A *(bit) commitment scheme* in a pair of †-compact categories $\langle \mathcal{C}^{\text{pure}}, \mathcal{C} \rangle$ forming part of an environment structure $\top$ (equivalently, the image under CPM of a †-compact category) is specified by

- two joint pure states $H, T : I \to A \otimes B$ in $\mathcal{C}^{\text{pure}}$

- a unitary morphism $\mathsf{unveil} : A \otimes B \to A \otimes B$ in $\mathcal{C}$, and
- a classical structure on $A \otimes B$ with distinct named classical points $\widehat{H}$ and $\widehat{T}$.

It is *sound* if

It is *concealing* if (∗), and *binding* if there is no unitary U for which (†).

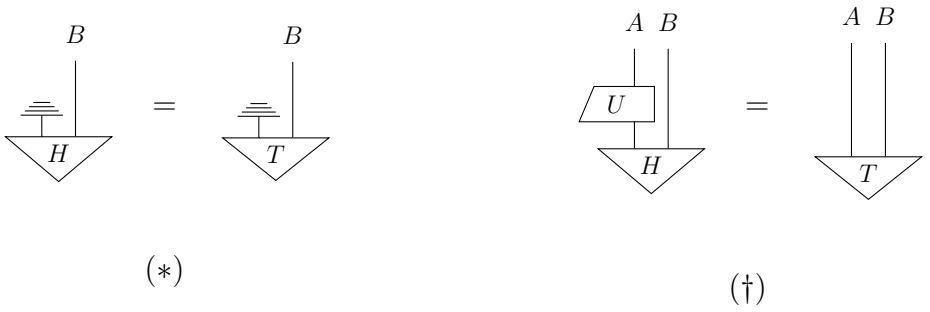

(∗) (†)

There are several things to explain here. First and foremost, such a structure can of course be used to perform bit commitment in the usual sense. To do so, Alice prepares $|H\rangle$ or $|T\rangle$ depending on the value of her bit, giving her two (possibly entangled) quantum systems. She sends the latter to Bob to perform the commitment, and the former to Bob to allow him to verify it. Bob applies the unitary unveil and measures in the basis of the specified classical structure; since we know that $\widehat{H}$ and $\widehat{T}$ are copyable states, Bob can identify them with certainty by measurement.

In FHilb Interpreting the definition in ordinary quantum mechanics is easy: the graphical language was developed exactly for this reason! States correspond to states and unitaries to unitaries, so that the axioms' interpretations become exactly what we would like. We'll look at this in a little more detail in the next section.

CPM We have phrased the definition in terms of environment structures and pure states, because we feel they are more intuitive to use. Note, however, that the use of environment structures (partial traces) means that we are implicitly working in a category CPM($\mathcal{C}$). For it to make sense there, we use the fact that the embedding WP : $\mathcal{C} \to$ CPM($\mathcal{C}$) is strict monoidal, and hence that classical points, classical structures and unitary morphisms all lift directly. This is also why we require $\mathcal{C}$ to be †-compact. (In the language of quantum mechanics, we use the fact that a pure state $|\psi\rangle$ can be regarded as the mixed state $|\psi\rangle\langle\psi|$, or "$|\psi\rangle$ with probability 1.")

Termination There is a interesting subtlety, which is that we have required there to be some point in time after the commitment has been completed but before the unveil phase has begun. That is, we require the commitment process to terminate after some finite time. It turns out [Kent, 1999] that there are in fact protocols – indeed, classical

protocols – which do offer an unconditionally secure bit commitment, in the sense that they are sound, binding and concealing; the two caveats are that a) they require the use of special relativity, and b) they require an indefinitely-long *sustain* phase in which messages are continually exchanged between the two parties. We shall not consider such protocols to be bit commitment in the usual sense, mentioning them only as an indication of other work in the area.

3.2.1 Binding

Do the categorical conditions correspond to our intuitions about the security of the protocol? Our definition of binding can be shown equivalent to the intuitive one that "Alice cannot change her commitment afterwards" – that there is no state she can prepare that allows her, by means of a local unitary, to delay choosing her committed bit until the unveil phase. For suppose there were a state cheat and unitaries R_H and R_T with the desired property:

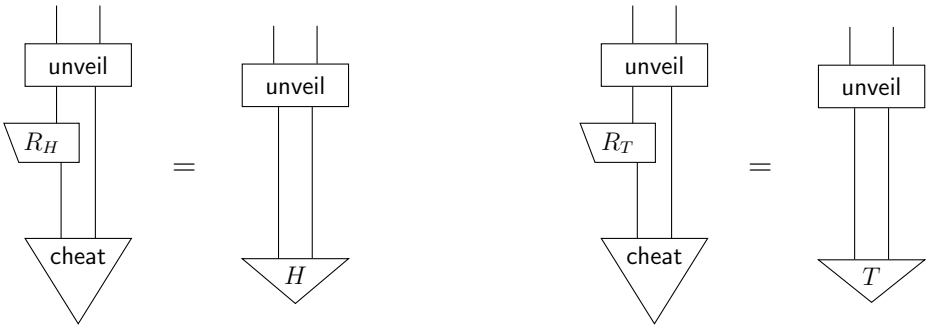

Since unveil is unitary it is monic, so

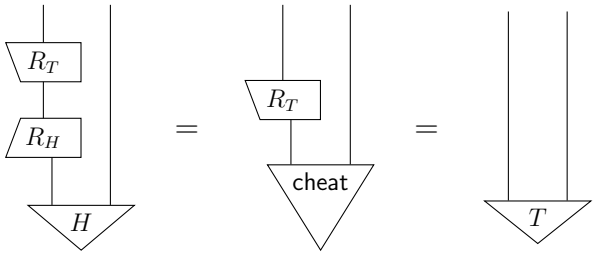

But then $R_T \circ R_H$ would be a unitary allowing Alice to change from H to T, contradicting our definition of binding.

The converse is trivial: if a protocol is not binding then there is some U contravening the definition; then Alice may choose $\mathsf{cheat} \coloneqq H$, $R_H \coloneqq \mathrm{id}$ and $R_T \coloneqq U$ to cheat in the above manner. Hence the two concepts are equivalent.

3.2.2 Concealing

Concealment implies that Bob cannot extract any information from the quantum system he is given that enables him to determine the value of Alice's bit. For the only way of extracting information from a quantum system is to perform a measurement, and the density matrix determines the value of such. Since the density matrix does not depend on Alice's choice of bit, neither can any information Bob extracts.

We might ask whether this condition corresponds to our intuition about concealment in the same way that binding does. Indeed, we could certainly formulate an intuitive definition of concealment, along the lines of "if Bob measures the the first system before receiving the second, he gets no information". This would mean, graphically, that

There are two problems with this condition, though. Firstly, this only considers measurement in the basis of $\bullet$, and we would like to prevent Bob from gaining information regardless of in which basis he measures. We should thus permit him to perform an arbitrary unitary transformation before taking the measurement.

More seriously, though, there is no reason that these two systems should be equal, because we have not done anything with Alice's unveil token. It is certainly not the case that we should require the unveil token to be independent of the committed state! We cannot delete it; there is no way to delete an arbitrary quantum system. The point is that we wish to forget it; that is, sum over all of its possible values – and *this* operation is precisely the partial trace $\top$.

Note that the condition

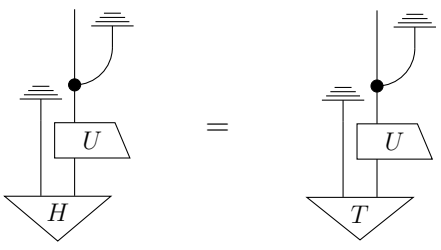

for all U

is implied by our definition. If $\bullet$ has enough nonzero-trace[1] classical points (i.e. $f = g$ as soon as $f \circ x = g \circ x$ for all classical points x with $\top \circ x \neq 0$) then this condition is equivalent to ours: classical points are copied by $\bullet$, so we may choose $U = \mathrm{id}$ to deduce

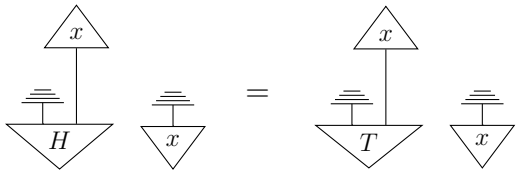

for all x. Since $\top \circ x \neq 0$ it is invertible – all nonzero scalars are – and hence it can be cancelled. It follows that $\langle x| \circ \mathrm{Tr}_A |H\rangle = \langle x| \circ \mathrm{Tr}_A |T\rangle$ for all nonzero-trace classical points, and hence $\mathrm{Tr}_A |H\rangle = \mathrm{Tr}_A |T\rangle$.

3.2.2.1 Density matrices

To see that this really does correspond to the usual density-matrix formulation, let's expand out the definition of the environment structure. By the second axiom of environment structures,

iff

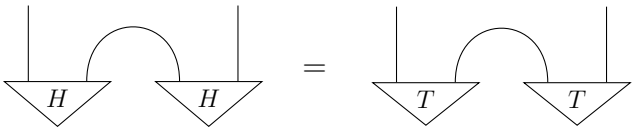

which is equivalent by a simple transposition to

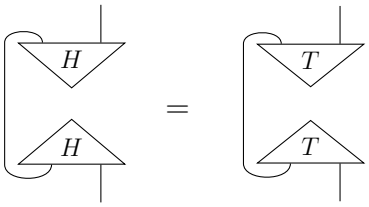

Written horizontally, this says that $\mathrm{tr}_A |H\rangle\langle H| = \mathrm{tr}_A |T\rangle\langle T|$, as we wanted.

[1]The usual definition of "all classical points" does not specifically exclude the case $\top \circ x = 0$. In **FHilb**, however, $0 = \top \circ x = \mathrm{Tr}\,|x\rangle\langle x| = \sum x_i^2$ iff $|x\rangle = 0$, and 0 does not help to distinguish states. Since the definition was abstracted from **FHilb** anyway, we take the liberty of excluding zero.

3.2.3 Commitment to an arbitrary input

We have restricted to *bit* commitment here: that is, Alice can choose either H or T as her message. It is certainly sensible to think about a commitment scheme which takes arbitrary input or even quantum input. Such a scheme could perhaps be modelled by replacing the states $|H\rangle$ and $|T\rangle$ with a morphism

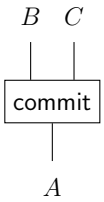

For such a scheme to be useful, Alice would need a way to encode classical data in the input quantum system – otherwise it could not actually be used for commitment! This means that we would require a classical structure on A. Moreover, the classical structure should have at least two distinct copyable states $|h\rangle$ and $|t\rangle$, since otherwise there is only one message to which Alice can commit.

In that case, defining $|H\rangle \coloneqq \mathsf{commit} \circ |h\rangle$ and $|T\rangle \coloneqq \mathsf{commit} \circ |t\rangle$ reduces such schemes to bit commitment schemes *à la* Definition 3.2.1. (Intuitively, if Alice can commit to an arbitrary non-trivial quantum system then she can encode a bit in that system, and hence commit to a bit.)

3.2.4 Commitment to a mixed state

Our definition of commitment is to a pure state. One could also consider the weaker notion of commitment to a mixed state, in which the chosen states $|H\rangle$ and $|T\rangle$ are no longer pure classical. This would allow us to formulate, for instance, Bennett and Brassard's protocol in this framework, since we may regard a pure state chosen by random coin flip as the mixed state comprising 50% of each.

Physical measurements, however, only apply to pure states: the result of measuring a mixed state is a random variable, even if it is guaranteed to be a mixture only of basis states. Thus, to incorporate this style of commitment into the framework would entail giving up determinism. Instead, our definitions of security – sound, binding, and concealing – would have to be rephrased in terms of a parameter n indicating the number of iterations of the protocol, and would require that by varying n we could cause their probability of success to approach 1.

Since we'll see that even commitment to a pure state is possible in **Rel**, this restriction is not important for the moment. Nevertheless, we mention it again in §4.1 in the context of the CBH theorem.

3.2.5 Other communication frameworks

Our model of a cryptographic protocol is about as simple as it can be: Alice prepares a joint state on two systems, sending one to commit and the other to unveil. There are certainly other models to consider. For example, Lo and Chau [1997] have both participants operate on a single system, exchanging it back and forth repeatedly.

3.3 Impossibility

It is a well known fact about quantum mechanics – in the von Neumann axiomatisation – that bit commitment is impossible. The usual proof is based on the following facts about Hilbert spaces.

3.3.1 Theorem (Singular value decomposition)**.** *Let* $f : \mathcal{H}_1 \to \mathcal{H}_2 \in \mathrm{Mor}(\mathbf{FHilb})$. *Then* $f = \mathcal{H}_1 \xrightarrow{V} \mathcal{H}_1 \xrightarrow{\Sigma} \mathcal{H}_2 \xrightarrow{U} \mathcal{H}_2$ *for some diagonal* Σ *and unitary* U *and* V.

3.3.2 Corollary (Schmidt decomposition). Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be Hilbert spaces of dimensions $m \geqslant n$ respectively. For each $|v\rangle \in \mathcal{H}_1 \otimes \mathcal{H}_2$ there exist bases $\{u_1, \ldots, u_n\} \subset \mathcal{H}_1$ and $\{v_1, \ldots, v_m\} \subset \mathcal{H}_2$ such that

$$|v\rangle = \sum_i \sqrt{p_i}(|u_i\rangle \otimes |v_i\rangle)$$

3.3.3 Theorem (Uniqueness of spectral decompositions)**.** *The decomposition of an operator into a sum of projectors onto eigenspaces is unique up to permutations of the bases of said eigenspaces. In particular, if each projector has rank 1 then the decomposition is unique.*

Proofs. Theorem 3.3.1 is a standard result of matrix analysis; a proof can be found in e.g., [Horn and Johnson, 1985, Theorem 7.3.5]. The Corollary can be deduced graphically; we do so in §3.3.1. Theorem 3.3.3 was mentioned previously. □

3.3.4 Theorem (Mayers, Lo and Chau)**.** *Bit commitment in* **FHilb** *is impossible.*

Remark. We prove that no concealing protocol can be binding. If Alice wishes to cheat, she sends Bob halves of entangled qubit pairs, and stores the other halves. The concealment condition, together with singular value decomposition, allows us to construct a unitary witnessing non-binding.

Proof. This proof was put together by Mayers [1997], and Lo and Chau [1997], although for the part we are interested in the main result is from [Hughston et al., 1993, §3.3].

Let $A \otimes B$ be the protocol's Hilbert space, containing Alice and Bob's qubits (and ancillas, if required). Let $|H\rangle$ and $|T\rangle$ be the joint states of the entire system corresponding to Alice's choice of bit. Taking the Schmidt decomposition gives

$$|H\rangle = \sum_i \sqrt{\sigma_i}|a_i\rangle \otimes |b_i\rangle \qquad\qquad |T\rangle = \sum_i \sqrt{\tau_i}|c_i\rangle \otimes |d_i\rangle.$$

Tracing out over Alice's system gives Bob's density matrices

$$\mathrm{Tr}_A\, |0\rangle\langle 0| = \sum_i \sigma_i |b_i\rangle\langle b_i| \qquad\qquad \mathrm{Tr}_A\, |1\rangle\langle 1| = \sum_i \tau_i |d_i\rangle\langle d_i|$$

For the protocol to be concealing, these density matrices must be equal.

We claim that we can choose $|d_i\rangle$ such that $\sigma_i = \tau_i$ and $|b_i\rangle = |d_i\rangle$. If the decompositions are non-degenerate – that is, if each eigenvalue has multiplicity 1 – then this follows from uniqueness of spectral decompositions. Otherwise, we note that the eigenspaces have the same dimension and that for each σ_i, $\{|d_j\rangle : \sigma_j = \sigma_i\}$ form a basis for the corresponding eigenspace, so that there is a unitary transformation $U_i :: |d_j\rangle \mapsto |b_j\rangle$ for each d_j of eigenvalue σ_i. Putting these unitaries together gives a change of basis to an alternative Schmidt decomposition in which the $|b_i\rangle$ and the $|d_i\rangle$ are equal.

Having established the claim, the local unitary $|c_i\rangle\langle a_i| :: a_i \mapsto c_i$, applied to Alice's system, changes the global state from $|H\rangle$ to $|T\rangle$. This contradicts our definition of binding. □

Remark. In the "real world" we need to consider the case that the given density matrices are only approximately equal, and bound the probability of cheating. This is possible, but a bit harder; see Mayers [1997] for details.

3.3.1 The Schmidt decomposition

This proof was phrased in the language of Hilbert spaces, though we can lift parts of it to a general †-compact category.

3.3.5 Definition. Suppose A admits a classical structure $\bullet$. We say a state $\phi : I \to A \otimes A$ is *diagonal* if there is σ such that

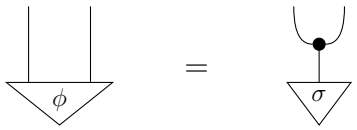

(In **FHilb**, σ would be given by $\sum_i \sigma_i |i\rangle$ for a basis $|i\rangle$ of A, and hence $\phi = \delta(\sigma) = \sum_i \sigma_i |i\rangle \otimes |i\rangle$ corresponds to a matrix which is diagonal in the induced basis on $A \otimes A$.)

3.3.6 Definition. We say a †-compact category $\mathcal{C}$ is has *singular value decompositions* if every morphism $f : A \to B$ can be written as

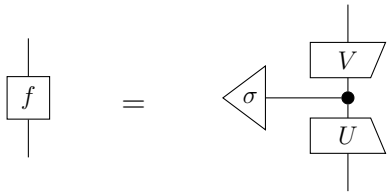

for some state σ and unitaries U and V.

3.3.7 Corollary (Schmidt decomposition). For any state $\phi : I \to A \otimes A$,

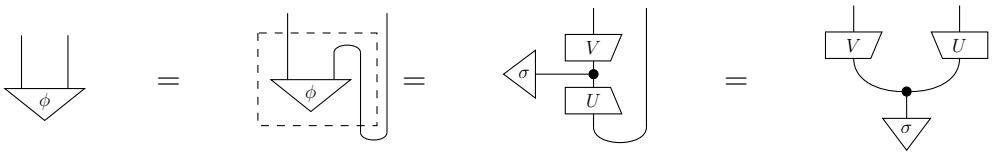

(In **FHilb**, this says that for every state $|\phi\rangle$ of a compound system there are bases $U|i\rangle$ and $V|i\rangle$ – given in terms of the basis $|i\rangle$ from our classical structure – such that

$$|\phi\rangle = (U \otimes V) \sum_i \mathsf{copy}\ |i\rangle = \sum_i U|i\rangle \otimes V|i\rangle;$$

in other words, that $|\phi\rangle$ is diagonal in those bases. So it really is the Schmidt decomposition as stated above.)

The Schmidt decomposition is easily seen to be equivalent to the singular value decomposition, since – following the above argument in reverse – an SVD of f is given by the morphism whose name is a Schmidt decomposition of $\ulcorner f \urcorner$.

The first part of the proof of Theorem 3.3.4 can now be done graphically. Suppose $|H\rangle$ and $|T\rangle$ form part of a commitment scheme in **FHilb**. Take their Schmidt decompositions

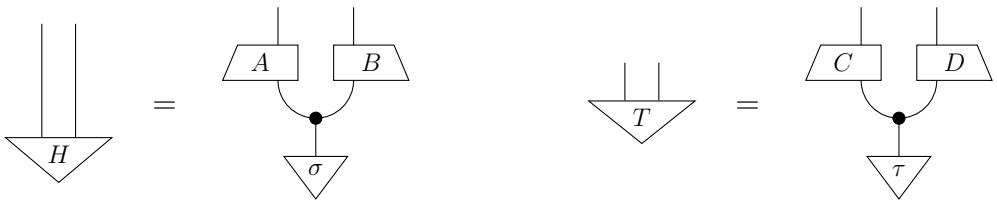

and apply the previously mentioned form of concealment

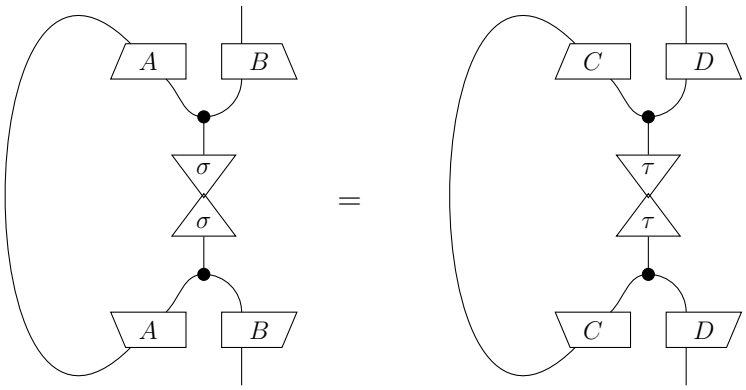

which rewrites to

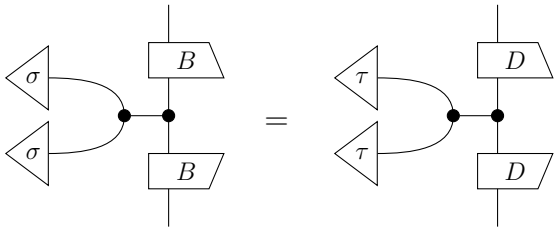

From here we would apply the uniqueness of spectral decompositions to deduce $B = D$ and $\sigma = \tau$.

The existence of singular value decompositions seems an awfully strong condition to place on a category. A good way to investigate it, therefore, is in our favourite toy quantum category.

3.4 Commitment in Rel

3.4.1 Singular Value Decompositions

Remember that classical structures in **Rel** are (the morphism sets of) abelian groupoids, so that the multiplication relates pairs (f, g) to their product fg and the comultiplication relates points h to factorisations $h = fg$. The tensor unit is the

one-point set – since tensor and Cartesian products coincide – so a state of X is a relation $R : \{*\} \to X$ i.e. a subset of X. Hence diagonal states in **Rel** are *factorisations of subsets*; that is, subsets S of $X \otimes X$ together with some $T \subseteq X$ such that $(x, y) \in S$ iff $xy \in T$.

A diagonal *relation* is of the form

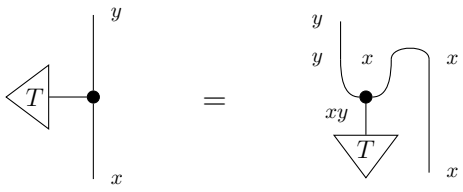

and hence is given by a subset $T \subseteq X$, relating x to y just if $yx = xy \in T$.

A singular value decomposition in **Rel** is then a decomposition of a relation into a bijection (unitary), followed by a factorisation of a subset, followed by another bijection (not necessarily the inverse of the former).

3.4.1 Proposition. *The cardinality of a diagonal relation (i.e., the number of points it relates) is $|T| \cdot |X|$, where T is the subset inducing the relation and X is the classical structure.*

Proof. Intuitively because the map $x \mapsto xy$ is a group isomorphism.

To count the number of pairs (x, y) with $xy \in T$, first fix $t \in T$. Then each $x \in X$ yields precisely one element y such that $xy = t$; namely, x/t. Moreover, if $xy = t$ then $y = t/x$, so every such pair is of this form. Thus there are $|X|$ of them for each t.

We must check that we do not double-count the pairs; that is, if two pairs arise from distinct t then they are distinct. But this is clear: if $x_1 y_1 = t_1 \neq t_2 = x_2 y_2$ then we cannot have $(x_1, y_1) = (x_2, y_2)$, since then $t_1 = t_2$. It follows that

$$|\{(x, y) : xy \in T\}| = \sum_{t \in T} |\{(x, y) : xy = t\}| = |T| \cdot |X|. \qquad \square$$

A Schmidt decomposition in **Rel** is a decomposition of a subset $S \subseteq X \otimes X$ as a factorisation of $T \subseteq X$ up to permutations of X. That is, we would like bijections $f, g : X \to X$ such that $\big(f(x), g(y)\big) \in S$ precisely when $x \cdot y \in T$, where by $- \cdot -$ we mean the multiplication induced by the classical structure (abelian groupoid) on X. It is now not hard to prove their non-existence.

3.4.2 Theorem. Rel *does not have singular value decompositions.*

Proof. Unitaries in **Rel** preserve cardinality. Hence a singular value decomposition of a morphism would imply that it had cardinality equal to that of a diagonal relation; in particular, divisible by $|X|$. This is not true in general.

(For a specific counterexample, consider the classical structure $\mathbb{Z}_2$ on the set $\{0,1\}$. The relation $\{0,1\} \to \{0,1\} :: 0 \sim 0$ has cardinality 1, which is odd, and hence cannot admit a singular value decomposition.) □

3.4.3 Corollary. CPM(**Rel**) does not have singular value decompositions.

Proof. We've already seen that classical points, classical structures and unitaries in CPM(**Rel**) are all canonical. Hence a singular value decomposition in CPM(**Rel**) induces a corresponding singular value decomposition in **Rel**, and this has just been proved impossible in general. □

3.4.2 Commitment

Theorem 3.4.2 tells us that the impossibility proof of Mayers, Lo and Chau for bit commitment fails in **Rel**. In fact, it is not too hard to see from the definition that commitment schemata do exist:

3.4.4 Theorem. *Bit commitment is possible in* **Rel**.

Sketch of proof. By construction. In **Rel** (identifying states with subsets), let

$$A \coloneqq \{1,2,3\} \qquad B \coloneqq \{a,b\}$$
$$H \coloneqq \{(1,a),(2,b),(3,b)\} \subset A \times B \qquad T \coloneqq \{(1,b),(2,a),(3,a)\} \subset A \times B.$$

Since H and T are disjoint and partition $A \times B$, the classical structure given by disjointly endowing A with $\mathbb{Z}/3\mathbb{Z}$ and B with $\mathbb{Z}/2\mathbb{Z}$ has precisely H and T as its classical points, so that choosing $\mathsf{unveil} = 1_{A\times B}$ makes $|H\rangle$ and $|T\rangle$ part of a sound commitment protocol. Concealment holds since the partial trace over A is $\{a,b\}$ for both commitments, and the protocol is binding because functions applied to a state only change the first coordinate, and there is no function f such that

$$\{(f(1),a),(f(2),b),(f(3),b)\} = \{(f(1),b),(f(2),a),(f(3),a)\}.$$ □

To formalise this proof we must lift all of the structures in **Rel** to CPM(**Rel**). Without further ado:

Proof. In **Rel**, let

$$\begin{aligned} A &\coloneqq \{1,2,3\} \\ B &\coloneqq \{a,b\} \\ H &\coloneqq \{(1,a),(2,b),(3,b)\} \subset A \times B \\ T &\coloneqq \{(1,b),(2,a),(3,a)\} \subset A \times B, \end{aligned}$$

lifting to

$$\begin{aligned} A^* \times A = \{&(1,1),(1,2),(1,3),(2,1),(2,2),\\ &(2,3),(3,1),(3,2),(3,3)\} \\ B^* \times B = \{&(a,a),(a,b),(b,a),(b,b)\} \\ H_* \times H = \Big\{&\big((a,1),(1,a)\big),\big((a,1),(2,b)\big),\big((a,1),(3,b)\big)\\ &\big((b,2),(1,a)\big),\big((b,2),(2,b)\big),\big((b,2),(3,b)\big)\\ &\big((b,3),(1,a)\big),\big((b,3),(2,b)\big),\big((b,3),(3,b)\big)\Big\} \\ T_* \times T = \Big\{&\big((b,1),(1,b)\big),\big((b,1),(2,a)\big),\big((b,1),(3,a)\big)\\ &\big((a,2),(1,b)\big),\big((a,2),(2,a)\big),\big((a,2),(3,a)\big)\\ &\big((a,3),(1,b)\big),\big((a,3),(2,a)\big),\big((a,3),(3,a)\big)\Big\} \end{aligned}$$

in CPM(**Rel**). Then

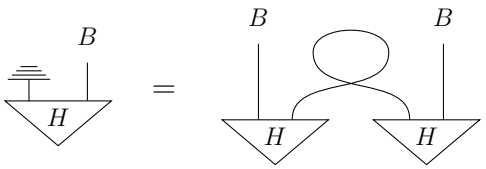

$$\begin{aligned} &= \{(\beta,\beta') \in B \times B : (\exists \alpha \in A)(\beta,\alpha) \in H_* \wedge (\alpha,\beta') \in H\} \\ &= \{(a,a),(b,b)\} \\ &= \dots \end{aligned}$$

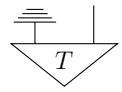

Moreover, for any function f, note that

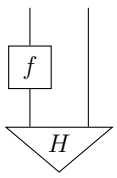

$$= \Big\{\big(f(1),a\big),\big(f(2),b\big),\big(f(3),b\big)\Big\} \neq$$

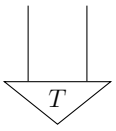

and the lift of this statement via WP hence certainly holds as well, since all unitaries are canonical.

Finally, choose unveil to be WP(1). Let $\bullet$ be the classical structure given by disjointly endowing A with $\mathbb{Z}/3\mathbb{Z}$ and B with $\mathbb{Z}/2\mathbb{Z}$, and WP($\bullet$) its lift to CPM(**Rel**). Then WP$|H\rangle$ and WP$|T\rangle$ are classical points of WP($\bullet$), completing the definition of a sound, binding, concealing commitment protocol. □

Remark. The restriction to unitary operators and measurements is second nature in quantum mechanics, but counterintuitive in **Rel**: we cannot interpret the above protocol as "send $\{a, b\}$ to commit and $\{1, 2, 3\}$ to unveil". Indeed, to do so would be to interpret the diagrams not in **Rel** but in **Set**, which is not even a †-compact category. To understand this as a commitment protocol we must have faith in the abstract definition.

We have shown that **Rel** has bit commitment schemes. What does this mean for us? The next chapter attempts to answer that question.

Chapter 4

Where Now?

We see that bit commitment *à la* Definition 3.2.1, though impossible in the category **FHilb**, can be implemented in **Rel**. This is not entirely an agreeable state of affairs, in light of Clifton et al.'s characterisation theorem. If we would like to treat †-compact categories as an axiomatisation of quantum mechanics, we must quibble with one of our postulates.

(1) We could strengthen our definition of a bit commitment scheme.

This seems wrong for two reasons. First, the axioms we have given correspond well with our intuition and with the usual classical definitions. But more importantly, they are sufficient for the proof in **FHilb** as they stand, so it is hard to make a case that they are not strong enough.

(2) We could strengthen the axioms for our categories. We've already seen the beginnings of one such strengthening; namely, that adding singular value decompositions gets us a long way[1] towards proving impossibility. This axiom, though, does not seem justified: singular value decompositions are a consequence of the linear structure of Hilbert spaces, which is not part of the categorical formulation.

Moreover, †-compact categories are already widely used as a sufficient model for quantum phenomena, so that strengthening them would be a major change.

Only one avenue remains to us.

(3) We could dispute the characterisation of quantum theories.

For Clifton et al., a physical theory is specified by its C^*-algebra of observables. Their theorem then states that any such theory is quantum iff it satisfies three

[1] The other result that we'd need is uniqueness of spectral decompositions, which it is not clear how to categorify.

information-theoretic axioms, one of which is the impossibility of secure bit commitment.

If we wish to apply their theorem, we must show two things: firstly, that our definition of bit commitment includes their particular protocol, and secondly that the "generalised" quantum theories given by †-compact categories fit into this class of physical theories. In other words, for it to apply to the quantum theory of a †-compact category, we would have to give a sensible definition of the "space of observables" of a †-compact category, and show that it comprises a C^*-algebra.

4.1 Weaker concepts of security

Let's look at the CBH commitment algorithm in a little more detail, in order to see how it could fit in our framework. They first use non-commutativity of the C^*-algebra (which, by their theorem, is equivalent to no-broadcasting) to deduce the existence of distinct pure states $\omega_{1,2}$ and $\omega_\pm$ such that

$$\tfrac{1}{2}(\omega_1 + \omega_2) = \tfrac{1}{2}(\omega_+ + \omega_-).$$

It follows that the mixed states

$$\rho_H \coloneqq \tfrac{1}{2}\big(|\omega_1\rangle \otimes |\omega_1\rangle + |\omega_2\rangle \otimes |\omega_2\rangle\big) \qquad \rho_T \coloneqq \tfrac{1}{2}\big(|\omega_+\rangle \otimes |\omega_+\rangle + |\omega_-\rangle \otimes |\omega_-\rangle\big)$$

have identical traces over each subsystem.

To commit to H, Alice flips a fair coin and produces either $\omega_0 \otimes \omega_0$ or $\omega_1 \otimes \omega_1$ depending on the result; to commit to T she does the same with $\omega_+ \otimes \omega_+$ and $\omega_- \otimes \omega_-$. She sends one half of the system to Bob and stores the other half.

To unveil the commitment, Alice reveals b and Bob performs a measurement in the $\omega_{0,1}$ or $\omega_\pm$ bases depending on its value. Finally, Alice performs the same measurement on her half of the system and sends the result to Bob, who verifies that they agree.

As stated this does not fit into our definition of a commitment protocol, since Alice performs a measurement after unveiling to verify her honesty. We can modify it slightly, by including Alice's measurement into the unveil phase. In the modified protocol, Alice simply sends her half of the system to Bob as well as the bit b. Bob then measures both halves in the basis corresponding to b; if they differ then Alice must have been cheating.

However, if Alice cheats – for example, by preparing $\omega_1\omega_+ + \omega_2\omega_-$ – there is still a chance that Bob's measurement outcomes will coincide, since measuring ω_+

in the $\{1, 2\}$ basis will yield ω_1 half of the time. Thus Bob cannot be certain of Alice's honesty after concluding the protocol. We get around this problem by noting that although in a single iteration Alice might be able to cheat, by conducting the protocol repeatedly the chance of doing so successfully approaches zero. This is a different notion of security from that which we stated before: it is (much) stronger than conditional security, but not quite as strong as what we called unconditional.

Since we want all of the commit stages to be completed before any unveil stage begins, to iterate the protocol in practise Alice prepares a long string of photons and sends them all as her commitment.

Mayers [1997] distinguishes between *perfectly* and *unconditionally* secure protocols. He defines perfect security to be what we have called unconditional, and calls a protocol unconditionally secure if it has an implicit parameter n and can be made secure with probability $p(n) \to 1$ as $n \to \infty$. The usual impossibility results in fact generalise to the probabilistic case; intuitively, if the partial traces over Alice are far apart from each other then Bob receives too much information, and if they are too close together then Alice's cheating strategy works with high probability. Clifton et al.'s above protocol is unconditionally secure as defined by Mayers, but of course not perfectly secure.

It is easy to modify our definition of commitment to allow $|H\rangle$ and $|T\rangle$ to be mixed states. However, since measurements of mixed states – even mixtures of classical points – are inherently probabilistic, we can no longer expect our security conditions to hold. In particular, we expect

- perfect soundness: if both Alice and Bob are honest then the measurements will be correlated with certainty
- perfect concealment: $\mathrm{tr}_A(\rho_H) = {}^1\!/_2(\omega_1 + \omega_2) = {}^1\!/_2(\omega_+ + \omega_-) = \mathrm{tr}_A(\rho_T)$
- unconditional binding: Alice may attempt to cheat, but then her measurement will differ from Bob's with $p \geq 0.5$ independently for each qubit

In order completely to characterise the type of commitment used in the CBH theorem, we would need to generalise our definition of security to allow for probabilistic success as per Mayers's definition. Having done so, we would hope to formalise the antecedent of the third condition of CBH in terms of the "abstract C^*-algebra of observables" of our category, and hence progress towards a categorical version of the theorem.

4.2 Conclusion

The search for the right language in which to conduct quantum mechanics is ongoing. Part of this search involves deciding when a new formalism pins down quantum mechanics as we understand it, and bit commitment, via the CBH theorem, is one piece of this. Categorical quantum mechanics is one such formalism, particularly convenient for describing algorithms because they can be written in terms of the sequence of operations from which they are built and omitting the 'implementation details'. But it is more than this: by abstracting away from Hilbert spaces to categories of processes, we provide a true generalisation from ordinary quantum mechanics. After this abstraction, we can still pin down some properties of quantum mechanics as we know it, in the spirit of Clifton et al.. In this vein, we ask whether the impossibility of bit commitment is a property specific to Hilbert spaces or whether it is true in any †-compact category.

In this dissertation, I have argued the former, via a definition of bit commitment in terms not of Hilbert spaces but of interacting processes. This definition reduces to the usual notion of commitment in **FHilb** but something rather different in **Rel**. The fact that commitment schemata exist there demonstrates that their impossibility requires more than just †-compactness, and indeed we have seen that the usual impossibility proof uses in addition the linear structure of Hilbert spaces.

We are led naturally to the question "what *does* forbid bit commitment in **FHilb**?" In CBH's C^*-algebraic formulation, the axiom they use is the physical realisability of entangled states. The phrasing of this axiom is a little delicate, because the formal existence of entangled states follows from the algebraic structure. Nevertheless, carefully stated, it is equivalent to the impossibility of bit commitment.

A natural direction for future work, linking in with the CP^* construction of Coecke, Heunen and Kissinger, is to investigate this equivalence in terms of the abstract C^*-algebras of a general †-compact category $\mathcal{C}$. Success would lift it to a result about commitment protocols in $\mathrm{CP}^*(\mathcal{C})$, of the same flavour as CBH's but entirely at a categorical level.

References

Samson Abramsky and Bob Coecke. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science*, LICS '04, pages 415–425. IEEE Computer Society, 2004. arXiv:quant-ph/0402130v5.

Charles Bennett and Gilles Brassard. *Quantum cryptography: Public key distribution and coin tossing*, pages 175–179. Bangalore, India, 1984.

Béla Bollobás. *Linear Analysis*. Cambridge University Press, 1990.

Jeffrey Bub. Why the quantum? *Studies in History and Philosophy of Modern Physics*, 35B: 241–266, 2004. arXiv:quant-ph/0402149v1.

Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and Its Applications*, pages 285–290, 1975. doi:10.1016/0024-3795(75)90075-0.

Rob Clifton, Jeffrey Bub, and Hans Halvorson. Characterizing quantum theory in terms of information-theoretic constraints. *Foundations of Physics*, 33:1561–1591, 2003. arXiv:quant-ph/0211089v2.

Bob Coecke. Quantum computer science. Unpublished course notes. Available at `http://www.cs.ox.ac.uk/people/bob.coecke/QCS.pdf`.

Bob Coecke. Kindergarten quantum mechanics. In *Proceedings of QTRF-III (AIP proceedings)*, 2005. Lecture notes of invited talks at Quantum Information, Computation & Logic (Perimeter Institute), QTRF-III (Vaxjo), Google (Silicon Valley) and Kestrel Institute (Silicon Valley).

Bob Coecke. Axiomatic description of mixed states from Selinger's CPM-construction. *Electronic Notes in Theoretical Computer Science*, 210:3–13, 2008.

Bob Coecke and Ross Duncan. Interacting quantum observables: categorical algebra and diagrammatics. *New Journal of Physics*, 13(4):043016, 2011. arXiv:0906.4725v3 [quant-ph].

Bob Coecke and Dusko Pavlovic. Quantum mechanics without sums. *The Mathematics of Quantum Computation and Technology*, pages 559–596, 2006. arXiv:quant-ph/0608035v2.

Bob Coecke and Simon Perdrix. Environment and classical channels in categorical quantum mechanics. In *Proceedings of the 19th EACSL Annual Conference on Computer Science Logic*, pages 230–244. Springer-Verlag, 2010. arXiv:1004.1598v4 [quant-ph].

Bob Coecke, Dusko Pavlovic, and Jamie Vicary. A new description of orthonormal bases. *Mathematical Structures in Computer Science*, 2008. arXiv:0810.0812v1 [quant-ph].

Julia Evans, Ross Duncan, Alex Lang, and Prakash Panangaden. Classifying all mutually unbiased bases in rel. arXiv:0909.4453v2 [quant-ph], Sep 2009.

Hans Halvorson. Remote preparation of arbitrary ensembles and quantum bit commitment. *J.Math.Phys.*, 45:4920–4931, 2003. arXiv:quant-ph/0310001v2.

Chris Heunen and Sergio Boixo. Completely positive classical structures and sequentializable quantum protocols. In *Proceedings of the 8th International Workshop on Quantum Physics and Logic*, 2011. To appear.

Chris Heunen, Ivan Contreras, and Alberto S. Cattaneo. Relative frobenius algebras are groupoids. *J. Pure and Applied Algebra*, to appear. arXiv:1112.1284v2 [math.CT].

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

Lane P. Hughston, Richard Jozsa, and William K. Wootters. A complete classification of quantum ensembles having a given density matrix. *Physics Letters A*, 183(1):14–18, Nov 1993.

André Joyal and Ross Street. The geometry of tensor calculus I. *Advances in Mathematics*, 88: 55–113, 1991.

Adrian Kent. Unconditionally secure bit commitment. *Phys. Rev. Lett.*, 83:1447–1450, 1999. arXiv:quant-ph/9810068v4.

Hoi-Kwong Lo and H.F. Chau. Is quantum bit commitment really possible? *Phys.Rev.Lett.*, 78: 3410, 1997. arXiv:quant-ph/9603004v2.

Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Phys. Rev. Lett.*, 78:3414–3417, Apr 1997. arXiv:quant-ph/9605044v2.

Ashwin Nayak and Pranab Sen. Invertible quantum operations and perfect encryption of quantum states. *Quantum Information and Computation*, pages 103–110, 2007. arXiv:quant-ph/0605041v4.

Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000. ISBN 9780521635035.

Dusko Pavlovic. Quantum and classical structures in nondeterministic computation. In Peter Bruza, Don Sofge, and Keith van Rijsbergen, editors, *Proceedings of Quantum Interaction*, pages 143–158. Springer-Verlag, 2009. arXiv:0812.2266v3 [quant-ph].

Peter Selinger. Dagger compact closed categories and completely positive maps. *Electronic Notes in Theoretical Computer Science*, 170:139–163, March 2007. doi:10.1016/j.entcs.2006.12.018.

Peter Selinger. A survey of graphical languages for monoidal categories. *Springer Lecture Notes in Physics*, 813:289–355, 2011. arXiv:0908.3347v1 [math.CT].

Claude E. Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28: 656–715, 1949. URL `http://netlab.cs.ucla.edu/wiki/files/shannon1949.pdf`.