

Quantum Strategies and Local Operations

by

Gustav Gutoski

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, Canada, 2009

© Gustav Gutoski 2009

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

This thesis is divided into two parts. In Part I we introduce a new formalism for quantum strategies, which specify the actions of one party in any multi-party interaction involving the exchange of multiple quantum messages among the parties. This formalism associates with each strategy a single positive semidefinite operator acting only upon the tensor product of the input and output message spaces for the strategy. We establish three fundamental properties of this new representation for quantum strategies and we list several applications, including a quantum version of von Neumann’s celebrated 1928 Min-Max Theorem for zero-sum games and an efficient algorithm for computing the value of such a game.

In Part II we establish several properties of a class of quantum operations that can be implemented locally with shared quantum entanglement or classical randomness. In particular, we establish the existence of a ball of local operations with shared randomness lying within the space spanned by the no-signaling operations and centred at the completely noisy channel. The existence of this ball is employed to prove that the weak membership problem for local operations with shared entanglement is strongly NP-hard. We also provide characterizations of local operations in terms of linear functionals that are positive and “completely” positive on a certain cone of Hermitian operators, under a natural notion of complete positivity appropriate to that cone. We end the thesis with a discussion of the properties of no-signaling quantum operations.

Acknowledgements

First and foremost, I thank my supervisor John Watrous for his support. He is an outstanding supervisor and I am very lucky to have been given the opportunity to benefit from his guidance for the past six years. I also thank the rest of my committee—Andrew Childs, Richard Cleve, Alex Russell, and Levent Tunçel—for taking the time to read this thesis and for their helpful suggestions. Finally, I thank my parents for a lifetime of encouragement and support.

Contents

1	Introduction	1
1.1	Overview	1
1.1.1	Quantum strategies	1
1.1.2	Local operations with shared entanglement	5
1.2	Mathematical preliminaries	9
1.2.1	Linear algebra	10
1.2.2	Convexity	20
1.2.3	Quantum information	22
I	Quantum Strategies	25
2	Introduction to Quantum Strategies	26
2.1	Operational formalism	26
2.1.1	Formal definitions	26
2.1.2	Immediate observations	29
2.2	New formalism	32
2.2.1	Formal definitions	32
2.2.2	Immediate observations	34
3	Three Important Properties	38
3.1	Interaction output probabilities	39
3.2	Characterization of strategies	44
3.3	Maximum output probabilities	48
3.3.1	Proof by convex polarity	49
3.3.2	Proof by semidefinite optimization duality	52
3.4	Other properties of strategies	58

4 Applications	60
4.1 Theory of zero-sum quantum games	60
4.1.1 Min-max theorem for zero-sum quantum games	61
4.1.2 Efficient algorithm to compute the value	64
4.2 Quantum interactive proofs	72
4.3 Kitaev’s bound for strong coin-flipping	78
5 Distance Measures	81
5.1 A new norm for strategies	81
5.2 Distinguishing convex sets of strategies	89
5.3 The dual of the diamond norm	95
II Local Operations with Shared Entanglement	102
6 Introduction to Local Operations	103
7 Ball Around the Completely Noisy Channel	109
7.1 General results on separable operators	110
7.2 Ball around the completely noisy channel	115
8 Recognizing LOSE Operations is NP-hard	117
8.1 Co-operative quantum games with shared entanglement	117
8.2 Strategies and weak validity	119
8.3 The Yudin-Nemirovskii Theorem and weak membership	120
9 Characterizations of Local Operations	122
9.1 Characterization of local operations with shared randomness	122
9.2 Characterization of local operations with shared entanglement	123
10 No-Signaling Operations	126
10.1 Formal definition of a no-signaling operation	126
10.2 Two characterizations of no-signaling operations	127
10.3 A separable no-signaling operation that is not a LOSE operation	131

11 Conclusion	133
11.1 Quantum strategies	133
11.2 Local operations with shared entanglement	134
References	135

Chapter 1

Introduction

This thesis investigates two distinct topics of interest within the discipline of quantum information theory: quantum strategies and local operations with shared entanglement. The discussion on quantum strategies is contained in Part I, while the discussion on local operations with shared entanglement is given in Part II.

This introductory chapter provides a broad overview of the results of the thesis in Section 1.1. A review of relevant background material from linear algebra, convex analysis, and quantum information is provided in Section 1.2.

1.1 Overview

In this section we provide a summary of the main contributions of the present thesis. Mathematics and quantum formalism are invoked only informally so as to facilitate a broad description of results without getting bogged down in detail.

Except where otherwise noted, the content of this thesis is drawn from existing literature as follows:

- Part I, excluding Chapter 5, first appeared in preliminary form in Ref. [GW07].
- Chapter 5 is otherwise unpublished and due solely to the present author.
- Part II first appeared in Ref. [Gut09].

1.1.1 Quantum strategies

In Part I of this thesis we propose a new mathematical formalism for quantum strategies, prove several fundamental properties of this formalism, and provide several applications.

Informally, a *quantum strategy* is a complete specification of the actions of one party in any multi-party interaction involving the exchange of one or more quantum messages among the parties. Due to the generality of this notion, the potential for application of this formalism is very broad. Indeed, our formalism should in principle apply to any framework that incorporates the exchange of quantum information among multiple entities, such as quantum cryptography, computational complexity, communication complexity, and distributed computation.

In this introductory section concerning quantum strategies, it is convenient to avoid cluttering discussion with historical background and citations. Instead, the necessary background and references for each topic are covered in detail as they appear in the main body of the thesis.

Three properties of the new formalism

Chapter 2 is devoted to formal definitions of quantum strategies and a discussion thereof. Under our new formalism, such a strategy is represented by a positive semidefinite operator S , the dimensions of which depend only upon the size of the messages exchanged in the interaction and *not* upon the size of any memory workspace maintained by the strategy between messages. (This distinction is important, as it permits us to consider strategies that call for an arbitrarily large memory workspace.) We prove in Chapter 3 that the set of all positive semidefinite operators which are valid representations of quantum strategies is characterized by a simple and efficiently-verifiable collection of linear equality conditions.

In order to extract useful classical information from such an interaction, a strategy will often call for one or more quantum measurements throughout the interaction. In this case, the strategy is instead represented by a set $\{S_a\}$ of positive semidefinite operators indexed by all the possible combinations of outcomes of the measurements. These strategies are called *measuring strategies* and satisfy $\sum_a S_a = S$ for some ordinary (non-measuring) strategy S . (By comparison, an ordinary POVM-type quantum measurement $\{P_a\}$ satisfies $\sum_a P_a = I$.)

We also prove in Chapter 3 that the relationship between measuring strategies and other strategies is analogous to that between ordinary quantum measurements and quantum states. In particular, basic quantum formalism tells us that for any ordinary quantum measurement $\{P_a\}$ with outcomes indexed by a and any quantum state ρ it holds that the probability with which the measurement $\{P_a\}$ yields a particular outcome a when applied to a quantum system in state ρ is given by the inner product

$$\Pr[\{P_a\} \text{ yields outcome } a \text{ on } \rho] = \langle P_a, \rho \rangle = \text{Tr}(P_a \rho).$$

Similarly, we show that the probability with which a measuring strategy $\{S_a\}$ yields outcome a after an interaction with a compatible quantum strategy T is given by

$$\Pr[\{S_a\} \text{ yields outcome } a \text{ when interacting with } T] = \langle S_a, T \rangle = \text{Tr}(S_a T).$$

Finally, we establish a convenient formula for computing the maximum probability with which a given measuring strategy $\{S_a\}$ can be forced to produce a given outcome a . This probability is given by the *minimum* real number λ for which there exists an ordinary (non-measuring) strategy Q with the property that the operator $\lambda Q - S_a$ is still positive semidefinite.

Applications

These three properties of quantum strategies—their linear characterization, inner product relationship, and formula for maximum output probability—open the door to a variety of new applications, several of which are presented in Chapter 4.

First and foremost, these properties pave the way for the first fully general theory of two-player zero-sum quantum games. In particular, we prove a quantum analogue of von Neumann’s famous 1928 Min-Max Theorem for zero-sum games. We also show that the value of such a game can be expressed as the value of a semidefinite optimization problem and can therefore be efficiently approximated to arbitrary precision by standard algorithms for semidefinite optimization.

We then apply this newfound algorithm to computational complexity theory, establishing that the fundamental class EXP of decision problems that admit deterministic exponential-time classical solutions coincides with the exotic class QRG of decision problems that admit a quantum interactive proof with two competing provers. That is,

$$\text{QRG} = \text{EXP}.$$

This equivalence is a rare characterization of a fundamental classical complexity class by a purely quantum complexity class. As problems in EXP also admit *classical* interactive proofs with competing provers, we obtain as a corollary the fact that quantum interactive proofs with competing provers provably contain no additional expressive power beyond that of classical interactive proofs with competing provers. By contrast, it is widely believed, but not proven, that polynomial-time quantum computers are strictly more powerful than polynomial-time classical computers.

Elsewhere within the domain of complexity theory, we employ our new formalism to prove that many-message quantum interactive proofs with one prover or with two competing provers may be repeated multiple times in parallel so as to decrease the probability of error without increasing the number of messages exchanged among the parties in the interaction.

Finally, the aforementioned properties of quantum strategies are applied to yield an alternate and simplified proof of Kitaev’s bound for strong quantum coin-flipping protocols. Coin-flipping is a fundamental primitive arising in the study of cryptography in the context of secure two-party computation. A *coin-flipping* protocol is an interaction between two mutually untrusting parties who wish to agree on a random bit (a coin flip) via remote communication. A *strong coin-flipping protocol* with bias ε has the property that two honest parties always produce a perfectly

random coin toss, yet a dishonest party who attempts to force a given outcome upon an honest party can succeed with probability no more than $1/2 + \varepsilon$. Kitaev showed that any strong coin-flipping protocol in which the parties exchange and process quantum information must have bias at least $1/\sqrt{2} - 1/2 \approx 0.207$. Kitaev’s original proof of this fact relied upon the powerful machinery of semidefinite optimization duality. In our proof, the complication of semidefinite optimization duality is successfully encapsulated in the properties of quantum strategies; what remains is a simple calculation that fits easily into half of a page.

It is noteworthy that a fourth application of the formalism of quantum strategies—beyond quantum game theory, complexity theory, and coin-flipping—appears in Chapter 8 of Part II of this thesis. In particular, the inner product relationship for quantum measuring strategies is employed to establish the NP-hardness of weak membership testing for local operations with shared entanglement.

In addition to our results, other authors have used the formalism of quantum strategies in other areas, as we now describe.

Independent development of the new formalism

Our formalism and some of its properties were independently re-discovered by Chiribella, D’Ariano, and Perinotti [CDP08e, CDP09a]. What we call a “quantum strategy,” they call a “quantum comb.” In their initial publication on the subject, these authors prove an analogue of our Theorem 3.2 (Characterization of strategies). Moreover, our Theorem 3.1 (Interaction output probabilities) is established in Refs. [CDP08a, CDP08d].

These and other authors have provided several additional applications of quantum strategies to such problems as optimization of quantum circuits architecture [CDP08e], cloning and learning of unitary operations [CDP08b, BCD⁺09a], and an impossibility proof for quantum bit commitment [CDP⁺09b], among others [CDP08c, BCD⁺09b].

Distance measures

After establishing useful properties of quantum strategies in Chapter 3 and then applying those properties in Chapter 4, we return in Chapter 5 to basic formalism for quantum strategies. We define a new norm that captures the distinguishability of quantum strategies in the same sense that the trace norm for operators captures the distinguishability of quantum states or the diamond norm for super-operators captures the distinguishability of quantum operations. Whereas the trace norm $\|\rho - \sigma\|_{\text{Tr}}$ for quantum states ρ, σ is given by

$$\|\rho - \sigma\|_{\text{Tr}} = \max \{ \langle P_0 - P_1, \rho - \sigma \rangle : \{P_0, P_1\} \text{ is a quantum measurement} \},$$

we define the *strategy r -norm* $\|Q - R\|_{\diamond_r}$ for quantum strategies Q, R by

$$\|Q - R\|_{\diamond_r} = \max \{ \langle S_0 - S_1, Q - R \rangle : \{S_0, S_1\} \text{ is a measuring strategy} \}.$$

Here the subscript r denotes the number of rounds of messages in the protocol for which Q, R are strategies. In particular, each positive integer r induces a different strategy norm. Our choice of notation is inspired by the fact that this norm is shown to coincide with the diamond norm for the case $r = 1$.

Our primary application of the strategy norm is a generalization of a result of Ref. [GW05], which states that for any two convex sets $\mathbf{A}_0, \mathbf{A}_1$ of quantum states there exists a *fixed* quantum measurement that distinguishes *any* two states chosen from these sets with probability that varies according to the minimal trace norm distance between the sets \mathbf{A}_0 and \mathbf{A}_1 . In other words, this measurement can be used to distinguish *any* choices of states from $\mathbf{A}_0, \mathbf{A}_1$ at least as well as any measurement could distinguish the two *closest* states from those sets.

Accordingly, our new result states that for any two convex sets $\mathbf{S}_0, \mathbf{S}_1$ of r -round quantum strategies, there exists a fixed measuring strategy that distinguishes any choices of strategies from those sets with probability according to the minimal distance between the sets \mathbf{S}_0 and \mathbf{S}_1 as measured by the new strategy r -norm.

We conclude Chapter 5 with a discussion of the dual of the diamond norm for super-operators and its relation to the strategy norms. The dual of the diamond norm $\|\Phi\|_{\diamond}^*$ of a super-operator Φ is defined by

$$\|\Phi\|_{\diamond}^* = \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle|$$

for some appropriate notion of inner product between super-operators.

While the diamond norm plays a fundamental role in the theory of quantum information, its dual has never been studied. Hence, we establish several basic facts about this norm. For example, the maximum in the definition of $\|\Phi\|_{\diamond}^*$ is achieved by a Hermitian-preserving super-operator whenever Φ is Hermitian-preserving, and by a completely positive super-operator whenever Φ is completely positive. These facts are employed to show that a variant of the strategy 1-norm coincides with the dual of the diamond norm. Thus, the strategy r -norms are shown to generalize both the diamond norm *and* its dual.

1.1.2 Local operations with shared entanglement

In Part II we prove several properties of local operations with shared entanglement or randomness. Informally, a *local operation* is a quantum operation that can be implemented by distinct parties, each acting only upon his or her own portion of an overall quantum system.

In a *local operation with shared randomness (LOSR)*, the parties are permitted to share random bits—say, common knowledge of an integer sampled at random according to some fixed probability distribution. The parties may use their knowledge of this shared randomness to correlate their distinct operations.

In a *local operation with shared entanglement (LOSE)*, rather than randomness, the parties are permitted to share among them distinct portions of some quantum

system, the overall state of which may be entangled across the different parties. The existence of such a shared state permits the parties to achieve correlations among their local operations that could not otherwise be achieved with randomness alone.

By contrast with Section 1.1.1, it is expedient to include a discussion of historical background and citations in this introductory section concerning local operations with shared entanglement.

Background

LOSE operations are of particular interest in the quantum information community in part because any physical operation jointly implemented by spatially separated parties who obey both quantum mechanics and relativistic causality must necessarily be of this form. Moreover, it is notoriously difficult to say anything meaningful about this class of operations, despite its simple definition.

One source of such difficulty stems from the fact that there exist two-party LOSE operations with the fascinating property that they cannot be implemented with any finite amount of shared entanglement [LTW08]. This difficulty has manifested itself quite prominently in the study of two-player co-operative games: classical games characterize NP,¹ whereas quantum games with shared entanglement are not even known to be computable. Within the context of these games, interest has focused largely on special cases of LOSE operations [KM03, CHTW04, Weh06, CSUU08, CGJ09, KRT08], but progress has been made recently in the general case [KKM⁺08, KKMV08, LTW08, DLTW08, NPA08, IKM09]. In the physics literature, LOSE operations are often discussed in the context of *no-signaling* operations [BGNP01, ESW02, PHHH06].

In the present thesis some light is shed on the general class of multi-party LOSE operations, as well as the sub-class of LOSR operations. Several distinct results are established, many of which mirror some existing result pertaining to separable quantum states. What follows is a brief description of each result together with its analogue from the literature on separable states where appropriate.

Ball around the identity

Prior work on separable quantum states. If A is a Hermitian operator acting on a d -dimensional bipartite space and whose Frobenius norm at most 1 then the perturbation $I \pm A$ of the identity represents an (unnormalized) bipartite separable quantum state. In other words, there is a ball of (normalized) bipartite separable states with radius $\frac{1}{d}$ in Frobenius norm centred at the completely mixed state $\frac{1}{d}I$ [Gur02, GB02].

¹ As noted in Ref. [KKM⁺08], this characterization follows from the PCP Theorem [ALM⁺98, AS98].

A similar ball exists in the multipartite case, but with smaller radius. In particular, there is a ball of m -partite separable d -dimensional states with radius $\Omega(2^{-m/2}d^{-1})$ in Frobenius norm centred at the completely mixed state [GB03]. Subsequent results offer some improvements on this radius [Sza05, GB05, Hil05].

Present work on local quantum operations. An analogous result is proven in Chapter 7 for multi-party LOSE and LOSR operations. Specifically, if A is a Hermitian operator acting on an n -dimensional m -partite space and whose Frobenius norm scales as $O(2^{-m}n^{-3/2})$ then $I \pm A$ is the Choi-Jamiołkowski representation of an (unnormalized) m -party LOSR operation. As the unnormalized completely noisy channel

$$\Delta : X \mapsto \text{Tr}(X)I$$

is the unique quantum operation whose Choi-Jamiołkowski representation equals the identity, it follows that there is a ball of m -party LOSR operations (and hence also of LOSE operations) with radius $\Omega(2^{-m}n^{-3/2}d^{-1})$ in Frobenius norm centred at the completely noisy channel $\frac{1}{d}\Delta$. (Here the normalization factor d is the dimension of the output system.)

The perturbation A must lie in the space spanned by Choi-Jamiołkowski representations of the no-signaling operations. Conceptual implications of this technicality are discussed in the concluding remarks of Chapter 7. No-signaling operations are discussed in Chapter 10 of the present thesis, as summarized below.

Comparison of proof techniques. Existence of this ball of LOSR operations is established via elementary linear algebra. By contrast, existence of the ball of separable states was originally established via a delicate combination of

- (i) the fundamental characterizations of separable states in terms of positive super-operators (described in more detail below), together with
- (ii) nontrivial norm inequalities for these super-operators.

Moreover, the techniques presented herein for LOSR operations are of sufficient generality to immediately imply a ball of separable states without the need for the aforementioned characterizations or their accompanying norm inequalities. This simplification comes in spite of the inherently more complicated nature of LOSR operations as compared to separable states. It should be noted, however, that the ball of separable states implied by the present work is smaller than the ball established in prior work by a factor of $2^{-m/2}d^{-3/2}$.

Weak membership problems are NP-hard

Prior work on separable quantum states. The weak membership problem for separable quantum states asks,

“Given a description of a quantum state ρ and an accuracy parameter ε , is ρ within distance ε of a separable state?”

This problem was proven strongly NP-complete under oracle (Cook) reductions by Gharibian [Gha08], who built upon the work of Gurvits [Gur02] and Liu [Liu07]. In this context, “strongly NP-complete” means that the problem remains NP-complete even when the accuracy parameter $\varepsilon = 1/s$ is given in unary as 1^s .

NP-completeness of the weak membership problem was originally established by Gurvits [Gur02]. The proof consists of an NP-completeness result for the weak *validity* problem—a decision version of linear optimization over separable states—followed by an application of the Yudin-Nemirovskii Theorem [YN76, GLS88], which provides an oracle-polynomial-time reduction from weak validity to weak membership for general convex sets.

As a precondition of the Yudin-Nemirovskii Theorem, the convex set in question (in our case, the set of separable quantum states) must contain a sufficiently large ball. In particular, this NP-completeness result relies crucially upon the existence of the aforementioned ball of separable quantum states.

For *strong* NP-completeness, it is necessary to employ a specialized “approximate” version of the Yudin-Nemirovskii Theorem due to Liu [Liu07, Theorem 2.3].

Present work on local quantum operations. In Chapter 8 it is proved that the weak membership problems for LOSE and LOSR operations are both strongly NP-hard under oracle reductions. The result for LOSR operations follows trivially from Gharibian (just take the input spaces to be empty). But it is unclear how to obtain the result for LOSE operations without invoking the contributions of the present thesis.

The proof begins by observing that the weak validity problem for LOSE operations is merely a two-player quantum game in disguise and hence is strongly NP-hard [KKM⁺08]. The hardness result for the weak *membership* problem is then obtained via a Gurvits-Gharibian-style application of Liu’s version of the Yudin-Nemirovskii Theorem, which of course depends upon the existence of the ball revealed in Section 7.2.

Characterization in terms of positive super-operators

Prior work on separable quantum states. A quantum state ρ of a bipartite system $\mathcal{X}_1 \otimes \mathcal{X}_2$ is separable if and only if the operator

$$(\Phi \otimes \mathbb{1}_{\mathcal{X}_2})(\rho)$$

is positive semidefinite whenever the super-operator Φ is positive. This fundamental fact was first proven in 1996 by Horodecki *et al.* [HHH96].

The multipartite case reduces inductively to the bipartite case: the state ρ of an m -partite system is separable if and only if $(\Phi \otimes \mathbb{1})(\rho)$ is positive semidefinite whenever the super-operator Φ is positive on $(m - 1)$ -partite separable operators [HHH01].

Present work on local quantum operations. In Chapter 9 it is proved that a multi-party quantum operation Λ is a LOSE operation if and only if

$$\varphi(J(\Lambda)) \geq 0$$

whenever the linear functional φ is “completely” positive on a certain cone of separable Hermitian operators, under a natural notion of complete positivity appropriate to that cone. A characterization of LOSR operations is obtained by replacing *complete* positivity of φ with mere positivity on that *same* cone. Here $J(\Lambda)$ denotes the Choi-Jamiołkowski representation of the super-operator Λ .

The characterizations presented in Chapter 9 do not rely upon any of the prior discussion in this thesis. This independence contrasts favorably with prior work on separable quantum states, wherein the existence of the ball around the completely mixed state (and the subsequent NP-hardness result) relied crucially upon the characterization of separable states in terms of positive super-operators.

No-signaling operations

A quantum operation is *no-signaling* if it cannot be used by spatially separated parties to violate relativistic causality. By definition, every LOSE operation is a no-signaling operation. Moreover, there exist no-signaling operations that are not LOSE operations. Indeed, it is noted in Chapter 10 that the standard nonlocal box of Popescu and Rohrlich [PR94] is an example of a no-signaling operation that is separable (in the sense of Rains [Rai97]), yet it is not a LOSE operation.

Two characterizations of no-signaling operations are also discussed in Chapter 10. These characterizations were first established somewhat implicitly for the bipartite case in Beckman *et al.* [BGNP01]. The present thesis generalizes these characterizations to the multi-party setting and recasts them more explicitly in terms of the Choi-Jamiołkowski representation for quantum super-operators.

1.2 Mathematical preliminaries

In this section we summarize the background mathematical knowledge upon which the work in this thesis rests. While extensive, by no means is this summary intended to be comprehensive. Instead, the purpose of this section is only to review existing concepts so that we may fix terminology and notation throughout the thesis.

1.2.1 Linear algebra

Vectors, operators, and inner products

The vector space \mathbb{C}^n of all n -tuples of complex numbers is called a *complex Euclidean space*. Complex Euclidean spaces are denoted by capital script letters such as \mathcal{X} , \mathcal{Y} , and \mathcal{Z} so as to better facilitate discussions involving multiple distinct spaces.

Vectors in a complex Euclidean space are denoted by lowercase Roman letters such as u , v , and w . The *standard orthonormal basis* of each n -dimensional complex Euclidean space is typically written $\{e_1, \dots, e_n\}$ where e_i denotes the n -tuple whose i th component equals 1 with all other components equal to 0. The *standard inner product* between two vectors $u, v \in \mathcal{X}$ is denoted $\langle u, v \rangle$. In this thesis, this inner product is conjugate linear in the first argument and linear in the second argument. In particular, if

$$\begin{aligned}u &= (\alpha_1, \dots, \alpha_n) \\v &= (\beta_1, \dots, \beta_n)\end{aligned}$$

then

$$\langle u, v \rangle = \sum_{i=1}^n \overline{\alpha_i} \beta_i.$$

The *standard Euclidean norm* of a vector u is denoted $\|u\|$ and is given by

$$\|u\| = \sqrt{\langle u, u \rangle}.$$

Each vector $u \in \mathcal{X}$ induces a *dual vector* u^* , which is a linear function $u^* : \mathcal{X} \rightarrow \mathbb{C}$ defined by

$$u^* : x \mapsto \langle u, x \rangle.$$

We sometimes use the alternate notation

$$u^*v = \langle u, v \rangle.$$

The vector $y \in \mathcal{Y}$ obtained by applying a linear operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ to a vector $x \in \mathcal{X}$ is denoted by the simple juxtaposition $y = Ax$. Similarly, the operator $C : \mathcal{X} \rightarrow \mathcal{Z}$ obtained by composing the linear operators $A : \mathcal{X} \rightarrow \mathcal{Y}$ and $B : \mathcal{Y} \rightarrow \mathcal{Z}$ is denoted by the juxtaposition $C = BA$.

Each pair of vectors $x \in \mathcal{X}, y \in \mathcal{Y}$ induces an operator $yx^* : \mathcal{X} \rightarrow \mathcal{Y}$ defined by

$$(yx^*)u = y(x^*u) = \langle x, u \rangle y$$

for all $u \in \mathcal{X}$. By analogy with the inner product, the operator yx^* is sometimes called the *outer product* of x and y , but we will not use that terminology in this thesis.

The *adjoint* of a linear operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ is the unique operator $A^* : \mathcal{Y} \rightarrow \mathcal{X}$ satisfying

$$\langle y, Ax \rangle = \langle A^*y, x \rangle$$

for every $x \in \mathcal{X}, y \in \mathcal{Y}$. The *standard inner product* for operators is given by

$$\langle A, B \rangle = \text{Tr}(A^*B).$$

Matrix representation of vectors and operators

It is sometimes convenient to think of a linear operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ as an $m \times n$ matrix and of an element $x \in \mathcal{X}$ as a $n \times 1$ column vector, so that the $m \times 1$ column vector $y \in \mathcal{Y}$ obtained by applying A to x is given by standard matrix multiplication:

$$\begin{bmatrix} \delta_1 \\ \vdots \\ \delta_m \end{bmatrix} = y = Ax = \begin{bmatrix} a_{1,1} & \dots & a_{1,n} \\ \vdots & \ddots & \vdots \\ a_{m,1} & \dots & a_{m,n} \end{bmatrix} \begin{bmatrix} \gamma_1 \\ \vdots \\ \gamma_n \end{bmatrix}$$

In this view, the asterisk superscript indicates the conjugate-transpose matrix operation. In particular, the *conjugate* operator $\bar{A} : \mathcal{X} \rightarrow \mathcal{Y}$ and *transpose* operator $A^\top : \mathcal{Y} \rightarrow \mathcal{X}$ are the linear operators whose matrix representations are given by

$$\bar{A} = \begin{bmatrix} \overline{a_{1,1}} & \dots & \overline{a_{1,n}} \\ \vdots & \ddots & \vdots \\ \overline{a_{m,1}} & \dots & \overline{a_{m,n}} \end{bmatrix}, \quad A^\top = \begin{bmatrix} a_{1,1} & \dots & a_{m,1} \\ \vdots & \ddots & \vdots \\ a_{1,n} & \dots & a_{m,n} \end{bmatrix}$$

and the matrix representation of the adjoint operator A^* is given by

$$A^* = (\bar{A})^\top = \overline{(A^\top)}.$$

The asterisk notation for matrices is consistent with the definition of the dual vector. Indeed, matrix multiplication can be used to compute the vector inner product:

$$u^*v = \begin{bmatrix} \overline{\alpha_1} & \dots & \overline{\alpha_n} \end{bmatrix} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_n \end{bmatrix} = \sum_{i=1}^n \overline{\alpha_i} \beta_i.$$

Matrix multiplication is also employed to compute the matrix representation of the operator $yx^* : \mathcal{X} \rightarrow \mathcal{Y}$:

$$yx^* = \begin{bmatrix} \delta_1 \\ \vdots \\ \delta_m \end{bmatrix} \begin{bmatrix} \overline{\gamma_1} & \dots & \overline{\gamma_n} \end{bmatrix} = \begin{bmatrix} \delta_1 \overline{\gamma_1} & \dots & \delta_1 \overline{\gamma_n} \\ \vdots & \ddots & \vdots \\ \delta_m \overline{\gamma_1} & \dots & \delta_m \overline{\gamma_n} \end{bmatrix}.$$

Similarly, matrix multiplication can be used to compute the matrix inner product—the composition A^*B of linear operators is computed by matrix multiplication of A^* and B , and the familiar trace function $\text{Tr}(A^*B)$ equals the sum of the diagonal entries of the resulting matrix A^*B .

Tensor products

The tensor product is a mathematical concept that is so fundamental to quantum information that discussion always implicitly assumes a working knowledge. The concept is important because it is the mechanism by which two separate quantum systems are viewed as a single system.

For any two complex Euclidean spaces \mathcal{X}, \mathcal{Y} and any two vectors $x \in \mathcal{X}, y \in \mathcal{Y}$, the *tensor product* associates a third vector $x \otimes y$, which is an element of a third vector space $\mathcal{X} \otimes \mathcal{Y}$ of dimension $\dim(\mathcal{X}) \dim(\mathcal{Y})$. Specifically, letting

$$\begin{aligned} \{e_0, \dots, e_{\dim(\mathcal{X})-1}\} &\subset \mathcal{X}, \\ \{f_0, \dots, f_{\dim(\mathcal{Y})-1}\} &\subset \mathcal{Y}, \\ \{g_0, \dots, g_{\dim(\mathcal{X} \otimes \mathcal{Y})-1}\} &\subset \mathcal{X} \otimes \mathcal{Y} \end{aligned}$$

denote the standard bases of \mathcal{X} , \mathcal{Y} , and $\mathcal{X} \otimes \mathcal{Y}$, respectively, the vector $x \otimes y$ is specified in terms of standard basis elements as follows:

$$e_i \otimes f_j = g_{i \dim(\mathcal{Y}) + j}.$$

The complex Euclidean space $\mathcal{X} \otimes \mathcal{Y}$ is called the *tensor product* of the spaces \mathcal{X} and \mathcal{Y} . (Indeed, any complex Euclidean space \mathcal{Z} whose dimension is not a prime number may be viewed as a tensor product $\mathcal{Z} = \mathcal{X} \otimes \mathcal{Y}$ of two nontrivial spaces \mathcal{X} and \mathcal{Y} .)

From this definition, it is possible to derive many of the widely known basic properties of the tensor product. For example, the bilinear mapping defined by

$$(x, y) \mapsto x \otimes y$$

is *universal*, meaning that *any* bilinear transformation $\psi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{V}$ for some complex Euclidean space \mathcal{V} can alternately be written as a linear operator $A_\psi : \mathcal{X} \otimes \mathcal{Y} \rightarrow \mathcal{V}$ satisfying

$$\psi(x, y) = A_\psi(x \otimes y).$$

As the set of all linear operators from one complex Euclidean space to another is itself a complex vector space, the tensor product may be extended in the obvious way to linear operators. From there it is possible to derive all the widely known properties of the tensor product for operators. We shall not list these properties here.

The tensor product of two or more vectors, spaces, or operators is defined inductively, as suggested by the expression

$$\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z} = \mathcal{X} \otimes (\mathcal{Y} \otimes \mathcal{Z}).$$

As the tensor product operation is associative, there is no ambiguity in writing simply $\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z}$.

Under the matrix representation, properties of the tensor product may be derived from the following straightforward definition. For 2×2 matrices

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \quad P = \begin{bmatrix} p & q \\ r & s \end{bmatrix}$$

we have

$$A \otimes P = \begin{bmatrix} aP & bP \\ cP & dP \end{bmatrix} = \begin{bmatrix} a \begin{bmatrix} p & q \\ r & s \end{bmatrix} & b \begin{bmatrix} p & q \\ r & s \end{bmatrix} \\ c \begin{bmatrix} p & q \\ r & s \end{bmatrix} & d \begin{bmatrix} p & q \\ r & s \end{bmatrix} \end{bmatrix} = \begin{bmatrix} ap & aq & bp & bq \\ ar & as & br & bs \\ cp & cq & dp & dq \\ cr & cs & dr & ds \end{bmatrix}.$$

This definition extends in the obvious way to arbitrary matrices of any dimension, including column vectors and non-square matrices. In this context, the tensor product might also be called the *Kronecker product*.

Often in this thesis our discussion involves tensor products of finite sequences of vectors, spaces, operators, and so on. As such, it is convenient to adopt the following shorthand for such a product: if X_1, \dots, X_m are arbitrary operators then we define

$$X_{i\dots j} \stackrel{\text{def}}{=} X_i \otimes \dots \otimes X_j$$

for integers $1 \leq i \leq j \leq m$. A similar notation shall be used for vectors and complex Euclidean spaces.

Sets of operators

A linear operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ is called an *isometry* if it holds that $\|Ax\| = \|x\|$ for all $x \in \mathcal{X}$. This condition can only be met when $\dim(\mathcal{Y}) \geq \dim(\mathcal{X})$. When \mathcal{X} and \mathcal{Y} have equal dimension, an isometry is also called a *unitary* operator.

The (complex) vector space of linear operators of the form $A : \mathcal{X} \rightarrow \mathcal{X}$ is denoted $\mathbf{L}(\mathcal{X})$. The identity operator in $\mathbf{L}(\mathcal{X})$ is denoted $I_{\mathcal{X}}$ and the subscript is dropped whenever the space \mathcal{X} is clear from the context.

An element A of $\mathbf{L}(\mathcal{X})$ is *Hermitian* (or *self-adjoint*) if $A^* = A$. The set of all Hermitian operators within $\mathbf{L}(\mathcal{X})$ forms a (real) vector space, which we denote by $\mathbf{H}(\mathcal{X})$.

An operator $A \in \mathbf{L}(\mathcal{X})$ is *positive semidefinite* if u^*Au is a nonnegative real number for each vector $u \in \mathcal{X}$. Every positive semidefinite operator is also Hermitian, and the set of all positive semidefinite operators within $\mathbf{H}(\mathcal{X})$ is denoted $\mathbf{H}^+(\mathcal{X})$. In general, the use of bold font is reserved for sets of operators. For an arbitrary set $\mathbf{S} \subset \mathbf{H}(\mathcal{X})$ of Hermitian operators, we let

$$\mathbf{S}^+ = \mathbf{S} \cap \mathbf{H}^+(\mathcal{X})$$

denote the set of positive semidefinite elements in \mathbf{S} .

We adopt the notation $P \succeq 0$ to indicate that the operator P is positive semidefinite. As suggested by this notation, the *semidefinite partial ordering* on Hermitian operators is defined so that $P \succeq Q$ if and only if $P - Q$ is positive semidefinite. For each positive semidefinite operator P there exists a unique positive semidefinite operator \sqrt{P} called the *square root* of P with the property that $(\sqrt{P})^2 = P$.

Super-operators

A *super-operator* is a linear operator of the form $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$. The identity super-operator from $\mathbf{L}(\mathcal{X})$ to itself is denoted $\mathbb{1}_{\mathcal{X}}$ and the subscript is dropped at will.

Whereas the application of an operator to a vector is denoted by simple juxtaposition, the operator $Y \in \mathbf{L}(\mathcal{Y})$ obtained by applying the super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ to an operator $X \in \mathbf{L}(\mathcal{X})$ is always denoted with parentheses: $Y = \Phi(X)$. Similarly, whereas operator composition is denoted by simple juxtaposition, the super-operator $\Gamma : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Z})$ obtained by composing the linear operators $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ and $\Psi : \mathbf{L}(\mathcal{Y}) \rightarrow \mathbf{L}(\mathcal{Z})$ is denoted $\Gamma = \Psi \circ \Phi$.

The *standard inner product* for super-operators $\Phi, \Psi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is given by

$$\langle \Phi, \Psi \rangle = \sum_{i,j=1}^{\dim(\mathcal{X})} \langle \Phi(e_i e_j^*), \Psi(e_i e_j^*) \rangle$$

where $\{e_1, \dots, e_{\dim(\mathcal{X})}\} \subset \mathcal{X}$ is the standard basis for \mathcal{X} . That this definition is a natural extension of the operator inner product can be argued from the fact that the operator inner product satisfies

$$\langle A, B \rangle = \sum_{i=1}^{\dim(\mathcal{X})} \langle A e_i, B e_i \rangle$$

for any two operators $A, B : \mathcal{X} \rightarrow \mathcal{Y}$.

Just as the definition of the tensor product is extended in a natural way to operators, so too can it be extended to super-operators.

A super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is said to be

- *Hermitian-preserving* if $\Phi(X)$ is Hermitian whenever X is Hermitian.
- *positive on \mathbf{K}* if $\Phi(X)$ is positive semidefinite whenever $X \in \mathbf{K}$.
- *positive* if Φ is positive on $\mathbf{H}^+(\mathcal{X})$.
- *completely positive* if $\Phi \otimes \mathbb{1}_{\mathcal{Z}}$ is positive for every choice of complex Euclidean space \mathcal{Z} .
- *trace-preserving* if $\text{Tr}(\Phi(X)) = \text{Tr}(X)$ for all X .

While the definition of complete positivity might seem awkward at first, it is significantly simplified by the observation that Φ is completely positive if and only if $\Phi \otimes \mathbb{1}_{\mathcal{Z}}$ is positive for a space \mathcal{Z} with $\dim(\mathcal{Z}) = \dim(\mathcal{X})$. In particular, there is no need to verify positivity of $\Phi \otimes \mathbb{1}_{\mathcal{Z}}$ for infinitely many spaces \mathcal{Z} .

An important example of a completely positive and trace-preserving super-operator is the *partial trace*. For any complex Euclidean spaces \mathcal{X}, \mathcal{Y} , this super-operator has the form $\text{Tr}_{\mathcal{X}} : \mathbf{L}(\mathcal{X} \otimes \mathcal{Y}) \rightarrow \mathbf{L}(\mathcal{Y})$. It is most easily specified by its actions upon product operators $X \otimes Y$ for $X \in \mathbf{L}(\mathcal{X}), Y \in \mathbf{L}(\mathcal{Y})$ by the expression

$$\text{Tr}_{\mathcal{X}}(X \otimes Y) = \text{Tr}(X)Y.$$

The importance of the partial trace stems from the fact that this super-operator is used to compute the quantum state of a portion of some larger system whose quantum state is already known.

Just as with operators, the *adjoint* of a super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is the unique super-operator $\Phi^* : \mathbf{L}(\mathcal{Y}) \rightarrow \mathbf{L}(\mathcal{X})$ satisfying

$$\langle Y, \Phi(X) \rangle = \langle \Phi^*(Y), X \rangle$$

for every $X \in \mathbf{L}(\mathcal{X}), Y \in \mathbf{L}(\mathcal{Y})$.

Every super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ may be expressed in the *operator-sum* notation, whereby there exist operators $A_1, \dots, A_k, B_1, \dots, B_k : \mathcal{X} \rightarrow \mathcal{Y}$ such that

$$\Phi(X) = \sum_{i=1}^k A_i X B_i^*$$

for all X , from which it follows that

$$\Phi^*(Y) = \sum_{i=1}^k A_i^* Y B_i$$

for all Y . It holds that Φ is completely positive if and only if it has a symmetric operator-sum decomposition, so that

$$\Phi(X) = \sum_{i=1}^k A_i X A_i^*$$

for all X .

Each super-operator also has a *Stinespring representation*, whereby there exists a complex Euclidean space \mathcal{Z} and operators $A, B : \mathcal{X} \rightarrow \mathcal{Y} \otimes \mathcal{Z}$ such that

$$\Phi(X) = \text{Tr}_{\mathcal{Z}}(A X B^*)$$

for all X . By analogy to the operator-sum representation, Φ is completely positive if and only if it has a symmetric Stinespring representation, so that

$$\Phi(X) = \text{Tr}_{\mathcal{Z}}(A X A^*)$$

for all X . If, in addition, Φ is also trace-preserving then the operator A must be an isometry.

The operator-vector and Choi-Jamiolkowski isomorphisms

In this thesis we make use of an unconventional but useful isomorphism vec that associates with each operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ a unique vector $\text{vec}(A)$ in the complex Euclidean space $\mathcal{Y} \otimes \mathcal{X}$. Letting $\{e_1, \dots, e_{\dim(\mathcal{X})}\} \subset \mathcal{X}$ and $\{f_1, \dots, f_{\dim(\mathcal{Y})}\} \subset \mathcal{Y}$ denote the standard bases of \mathcal{X} and \mathcal{Y} , this isomorphism defined by

$$\text{vec}(f_j e_i^*) = f_j \otimes e_i.$$

Note that this correspondence is basis-dependent, and we have chosen the standard basis in our definition.

In the matrix representation, the column vector $\text{vec}(A)$ is obtained from the entries of the matrix A by taking each row of A , transposing that row to form a column vector, and then stacking each of these column vectors so as to form one large column vector. For example, the vec mapping acts as follows on 2×2 matrices:

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \quad \text{vec}(A) = \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix}.$$

The vec isomorphism has many convenient properties, some of which we list here. Each of these identities may be verified by straightforward calculation.

Proposition 1.1 (Properties of the vec isomorphism). *The following hold:*

1. *The vec mapping preserves the standard inner products of operators and vectors. That is, for each $A, B : \mathcal{X} \rightarrow \mathcal{Y}$ it holds that*

$$\langle A, B \rangle = \langle \text{vec}(A), \text{vec}(B) \rangle.$$

2. *For any operators A, B , and X for which the composition AXB is defined it holds that*

$$(A \otimes B^\top) \text{vec}(X) = \text{vec}(AXB).$$

3. *For each $A, B : \mathcal{X} \rightarrow \mathcal{Y}$ it holds that*

$$\begin{aligned} \text{Tr}_{\mathcal{X}}(\text{vec}(A) \text{vec}(B)^*) &= AB^*, \\ \text{Tr}_{\mathcal{Y}}(\text{vec}(A) \text{vec}(B)^*) &= (B^*A)^\top. \end{aligned}$$

4. *For each $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ it holds that*

$$\text{vec}(yx^*) = y \otimes \bar{x}.$$

The *Choi-Jamiołkowski isomorphism* associates with each super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ a unique operator $J(\Phi) \in \mathbf{L}(\mathcal{Y} \otimes \mathcal{X})$. Letting $\{e_1, \dots, e_{\dim(\mathcal{X})}\} \subset \mathcal{X}$ denote the standard basis of \mathcal{X} , this isomorphism defined by

$$J(\Phi) = \sum_{i,j=1}^{\dim(\mathcal{X})} \Phi(e_i e_j^*) \otimes e_i e_j^*.$$

As with the vec isomorphism, the Choi-Jamiołkowski isomorphism is basis-dependent and it is always defined with respect to the standard basis.

Given that $I_{\mathcal{X}} = \sum_{i=1}^{\dim(\mathcal{X})} e_i e_i^*$, we obtain the following alternate characterization of the Choi-Jamiołkowski isomorphism:

$$J(\Phi) = (\Phi \otimes \mathbb{1}_{\mathcal{X}}) (\text{vec}(I_{\mathcal{X}}) \text{vec}(I_{\mathcal{X}})^*).$$

Like the vec isomorphism, the Choi-Jamiołkowski isomorphism has many convenient properties, some of which we list here.

Proposition 1.2 (Properties of the Choi-Jamiołkowski isomorphism). *The following hold for all super-operators $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$:*

1. *The Choi-Jamiołkowski isomorphism preserves the standard inner product of super-operators and operators. That is, for each $\Phi, \Psi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ it holds that*

$$\langle \Phi, \Psi \rangle = \langle J(\Phi), J(\Psi) \rangle.$$

2. *For each $X \in \mathbf{L}(\mathcal{X})$ it holds that*

$$\Phi(X) = \text{Tr}_{\mathcal{X}} ((I_{\mathcal{Y}} \otimes X^T) J(\Phi)).$$

3. *If Φ has operator-sum and Stinespring representations given by*

$$\Phi(X) = \sum_{i=1}^k A_i X B_i^* \quad \text{and} \quad \Phi(X) = \text{Tr}_{\mathcal{Z}}(A X B^*)$$

for all X then it holds that

$$J(\Phi) = \sum_{i=1}^k \text{vec}(A_i) \text{vec}(B_i)^* \quad \text{and} \quad J(\Phi) = \text{Tr}_{\mathcal{Z}}(\text{vec}(A) \text{vec}(B)^*),$$

respectively.

4. *Φ is Hermitian-preserving if and only if $J(\Phi)$ is Hermitian.*
5. *Φ is completely positive if and only if $J(\Phi)$ is positive semidefinite.*
6. *Φ is trace-preserving if and only if $\text{Tr}_{\mathcal{Y}}(J(\Phi)) = I_{\mathcal{X}}$.*

A generalization of item 2 is proven in Proposition 9.4 of Chapter 9. An additional identity involving the Choi-Jamiołkowski isomorphism appears later in this section in Proposition 1.3.

Operator decompositions

There are two operator decompositions that are so fundamental to quantum information that, like the tensor product, they are often used implicitly. Indeed, in this thesis we make little explicit mention of these decompositions, yet our discussion always assumes a working knowledge of their existence.

The *Singular Value Theorem* states that every operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ has at least one *singular value decomposition*, whereby there exist orthonormal sets $\{x_1, \dots, x_r\} \subset \mathcal{X}$ and $\{y_1, \dots, y_r\} \subset \mathcal{Y}$ and positive real numbers $s_1, \dots, s_r \in \mathbb{R}$ such that

$$A = \sum_{i=1}^r s_i y_i x_i^*.$$

Here r is the rank of the operator A . The real numbers s_1, \dots, s_r are called the *singular values* of A . Sometimes, the vectors y_1, \dots, y_r are called the *left singular vectors* of A , whereas x_1, \dots, x_r are called the *right singular vectors* of A .

The *Spectral Theorem* implies that an operator $A \in \mathbf{L}(\mathcal{X})$ is Hermitian if and only if there exists at least one *spectral decomposition* whereby there is an orthonormal set $\{x_1, \dots, x_r\} \subset \mathcal{X}$ and real numbers $\lambda_1, \dots, \lambda_r$ such that

$$A = \sum_{i=1}^r \lambda_r x_i x_i^*.$$

Again, r is the rank of A . The real numbers $\lambda_1, \dots, \lambda_r$ are called the *eigenvalues* of A and the vectors x_1, \dots, x_r are called the *eigenvectors* of A .

It is easy to see that a Hermitian operator A is positive semidefinite if and only if each of its eigenvalues is nonnegative. It follows immediately from the Spectral Theorem that every Hermitian operator A has a *Jordan decomposition* whereby there exist positive semidefinite operators P and Q with the property that

$$A = P - Q$$

and $PQ = 0$ —that is, P and Q act on orthogonal subspaces. The *absolute value* $|A|$ of A is a positive semidefinite operator defined via the Jordan decomposition by

$$|A| = P + Q.$$

A positive semidefinite operator P is called a *projection* if each of its eigenvalues is either zero or one.

Norms of vectors and operators

For each real number $p \geq 1$ and each vector $u = (\alpha_1, \dots, \alpha_n)$ in some n -dimensional complex Euclidean space, the *vector p -norm* $\|u\|_p$ of u is defined by

$$\|u\|_p = \left(\sum_{i=1}^n |\alpha_i|^p \right)^{1/p}$$

with $\|u\|_\infty$ given by

$$\|u\|_\infty = \lim_{p \rightarrow \infty} \|u\|_p = \max_i |\alpha_i|.$$

These norms satisfy $\|u\|_p \leq \|u\|_q$ whenever $p \geq q$, and it also holds that

$$\|u\|_1 \leq \sqrt{n} \|u\|_2 \leq n \|u\|_\infty.$$

The standard Euclidean norm $\|u\|$ is an instance of the vector p -norm with $p = 2$, so that $\|u\| = \|u\|_2$ for all u .

Each vector p -norm induces a norm on operators via the singular value decomposition. In particular, the *Schatten p -norm* $\|A\|_p$ of an operator A is defined as the vector p -norm of the singular values of A . As such, the Schatten p -norms inherit many properties from the vector p -norms. For example, $\|A\|_p \leq \|A\|_q$ whenever $p \geq q$ and

$$\|A\|_1 \leq \sqrt{n} \|A\|_2 \leq n \|A\|_\infty$$

for operators $A \in \mathbf{L}(\mathcal{X})$ with $\dim(\mathcal{X}) = n$.

In this thesis we are interested only in the Schatten p -norms for the values $p = 1, 2, \infty$. The Schatten p -norm for $p = 1$ is also called the *trace norm* and is alternately denoted $\|A\|_{\text{Tr}}$.

The Schatten p -norm for $p = 2$ is also called the *Frobenius norm* and is alternately denoted $\|A\|_{\text{F}}$. The Frobenius norm is merely the standard Euclidean norm for complex Euclidean spaces applied to the complex vector space of linear operators:

$$\|A\|_{\text{F}} = \sqrt{\langle A, A \rangle} = \sqrt{\langle \text{vec}(A), \text{vec}(A) \rangle} = \|\text{vec}(A)\|.$$

The Schatten p -norm for $p = \infty$ is also called the *standard operator norm* and is alternately denoted without any subscript by $\|A\|$. The operator norm is *induced* from the standard Euclidean norm for vectors:

$$\|A\| = \max_{\|u\|=1} \|Au\|.$$

In this thesis we prefer the notation $\|A\|_{\text{Tr}}$, $\|A\|_{\text{F}}$, and $\|A\|$ to $\|A\|_1$, $\|A\|_2$, and $\|A\|_\infty$.

We now prove a simple identity involving the trace norm of the Choi-Jamiołkowski representation of a completely positive and trace-preserving super-operator.

Proposition 1.3. *For any completely positive and trace-preserving super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ it holds that $\|J(\Phi)\|_{\text{Tr}} = \|J(\Phi^*)\|_{\text{Tr}} = \dim(\mathcal{X})$.*

Proof. By definition, we have

$$\begin{aligned} J(\Phi) &= (\Phi \otimes \mathbb{1}_{\mathcal{X}}) (\text{vec}(I_{\mathcal{X}}) \text{vec}(I_{\mathcal{X}})^*), \\ J(\Phi^*) &= (\Phi^* \otimes \mathbb{1}_{\mathcal{Y}}) (\text{vec}(I_{\mathcal{Y}}) \text{vec}(I_{\mathcal{Y}})^*). \end{aligned}$$

As $J(\Phi)$ is positive semidefinite, so too must be $J(\Phi^*)$, from which we obtain

$$\begin{aligned}\|J(\Phi)\|_{\text{Tr}} &= \text{Tr}(J(\Phi)), \\ \|J(\Phi^*)\|_{\text{Tr}} &= \text{Tr}(J(\Phi^*)).\end{aligned}$$

As Φ is trace-preserving, it holds that

$$\text{Tr}(J(\Phi)) = \text{vec}(I_{\mathcal{X}})^* \text{vec}(I_{\mathcal{X}}) = \dim(\mathcal{X}).$$

Moreover, it is easy to verify that $\Phi^*(I_{\mathcal{Y}}) = I_{\mathcal{X}}$, from which it follows that

$$\text{Tr}(J(\Phi^*)) = \text{Tr}(\Phi^*(I_{\mathcal{Y}})) = \text{Tr}(I_{\mathcal{X}}) = \dim(\mathcal{X}).$$

□

1.2.2 Convexity

Carathéodory's Theorem

A subset C of a real vector space \mathbb{R}^n is called *convex* if for every $x, y \in C$ and every real number $\alpha \in [0, 1]$ it holds that

$$\alpha x + (1 - \alpha)y \in C.$$

For vectors $v_1, \dots, v_m \in \mathbb{R}^n$ and nonnegative real numbers $\alpha_1, \dots, \alpha_m \geq 0$ with

$$\sum_{i=1}^m \alpha_i = 1$$

the vector

$$\sum_{i=1}^m \alpha_i v_i$$

is called a *convex combination* of v_1, \dots, v_m . For an arbitrary set $S \subset \mathbb{R}^n$, the *convex hull* of S is the subset of \mathbb{R}^n consisting of all convex combinations of elements in S . *Carathéodory's Theorem* is a useful result that bounds the number of terms in the sum of a given convex combination.

Fact 1.4 (Carathéodory's Theorem). *Let $S \subset \mathbb{R}^n$ be an arbitrary set. Every element x of the convex hull of S can be written as a convex combination of no more than $n + 1$ elements of S .*

Proofs of Carathéodory's Theorem can be found, for example, in Rockafellar [Roc70] or Barvinok [Bar02].

Separation Theorem

The fundamental *Separation Theorem* tells us that every pair of disjoint convex sets may be separated by a hyperplane. There are many variants of the Separation Theorem, each of which differs only slightly based upon properties of the sets in question such as whether or not they are open, closed, bounded, or cones. (A subset $K \subset \mathbb{R}^n$ is called a *cone* if for every $x \in K$ and every positive real number $\lambda > 0$ it holds that $\lambda x \in K$.) The following variant of the Separation Theorem serves all our needs in this thesis.

Fact 1.5 (Separation Theorem). *Let $C, D \subset \mathbb{R}^n$ be nonempty disjoint convex sets such that D is open. There exists a vector $h \in \mathbb{R}^n$ and a real number $\alpha \in \mathbb{R}$ with the property that*

$$\begin{aligned}\langle h, x \rangle &\geq \alpha \text{ for all } x \in C, \\ \langle h, y \rangle &< \alpha \text{ for all } y \in D.\end{aligned}$$

Moreover, if C is a cone then we may take $\alpha = 0$.

Proof. It follows from Theorem 11.3 of Rockafellar [Roc70] that there exists a nonzero vector $h \in \mathbb{R}^n$ and a real number $\alpha \in \mathbb{R}$ with $\langle h, x \rangle \geq \alpha \geq \langle h, y \rangle$ for all $x \in C$ and $y \in D$. Suppose toward a contradiction that there is a $y' \in D$ with $\langle h, y' \rangle = \alpha$. Choose any vector b with $\langle h, b \rangle > 0$, so that $\langle h, y' + b \rangle > \alpha$. By rescaling b and using the fact that D is open, we may assume $y' + b \in D$, which contradicts the fact that $\langle h, y \rangle \leq \alpha$ for all $y \in D$.

Next, suppose that C is a cone and let α' denote the infimum of $\langle h, x \rangle$ over all $x \in C$, so that

$$\langle h, x \rangle \geq \alpha' \geq \alpha > \langle h, y \rangle$$

for all $x \in C$ and $y \in D$. Our proof that $\alpha' = 0$ follows that of Theorem 11.7 of Rockafellar [Roc70]. First, suppose toward a contradiction that $\alpha' < 0$ and choose $x \in C$ with $\langle h, x \rangle < 0$. Then $\langle h, \lambda x \rangle$ can be made into an arbitrarily large negative number by an appropriately large choice of λ , contradicting the lower bound α on the infimum α' . Conversely, suppose toward a contradiction that $\alpha' > 0$ and choose λ small enough so that $\langle h, \lambda x \rangle < \alpha'$, contradicting the definition of α' . \square

Convexity and Hermitian operators

As noted in Section 1.2.1, the set $\mathbf{H}(\mathcal{X})$ of all Hermitian operators is a real vector space of dimension $\dim(\mathcal{X})^2$. As such, the formalism of convexity translates without complication to Hermitian operators. For example, the set $\mathbf{H}^+(\mathcal{X})$ of all positive semidefinite operators within $\mathbf{H}(\mathcal{X})$ is a closed convex cone. Moreover, Carathéodory's Theorem and the Separation Theorem have natural analogues in the context of $\mathbf{H}(\mathcal{X})$. Indeed, it is worth restating Fact 1.5 (Separation Theorem) in terms of Hermitian operators.

Fact 1.6 (Separation Theorem, Hermitian operator version). *Let $\mathbf{C}, \mathbf{D} \subset \mathbf{H}(\mathcal{X})$ be nonempty disjoint convex sets such that \mathbf{D} is open. There exists a Hermitian operator $H \in \mathbf{H}(\mathcal{X})$ and a real number $\alpha \in \mathbb{R}$ with the property that*

$$\begin{aligned}\langle H, X \rangle &\geq \alpha \text{ for all } X \in \mathbf{C}, \\ \langle H, Y \rangle &< \alpha \text{ for all } Y \in \mathbf{D}.\end{aligned}$$

Moreover, if \mathbf{C} is a cone then we may take $\alpha = 0$.

1.2.3 Quantum information

Quantum states

Associated with each d -level physical system is a d -dimensional complex Euclidean space \mathcal{X} . The *quantum state* of such a system at some fixed point in time is described uniquely by a positive semidefinite operator $\rho \in \mathbf{H}^+(\mathcal{X})$ with trace equal to one. Such an operator might also be called a *density operator*.

A quantum state ρ called *pure* if ρ has rank one—that is $\rho = uu^*$ for some unit vector $u \in \mathcal{X}$. It follows immediately from the Spectral Theorem that every quantum state may be written as a convex combination (or probabilistic ensemble) of orthogonal pure states.

Two distinct physical systems with associated complex Euclidean spaces \mathcal{X} and \mathcal{Y} may be viewed as a single larger system with associated complex Euclidean space $\mathcal{X} \otimes \mathcal{Y}$. If $\rho \in \mathbf{H}^+(\mathcal{X} \otimes \mathcal{Y})$ is the state of this larger system then the state $\sigma \in \mathbf{H}^+(\mathcal{X})$ of the subsystem associated only with \mathcal{X} is given by the partial trace

$$\sigma = \text{Tr}_{\mathcal{Y}}(\rho).$$

Conversely, if the systems with associated spaces \mathcal{X} and \mathcal{Y} are in states σ and σ' respectively then it must hold that the state $\rho \in \mathbf{H}^+(\mathcal{X} \otimes \mathcal{Y})$ of the larger system satisfies $\sigma = \text{Tr}_{\mathcal{Y}}(\rho)$ and $\sigma' = \text{Tr}_{\mathcal{X}}(\rho)$. Any state ρ meeting these conditions is said to be *consistent* with σ and σ' . If there is no correlation or entanglement between these two subsystems then ρ is given by the tensor product

$$\rho = \sigma \otimes \sigma'.$$

Such a state is called a *product state*. For each σ there are many states ρ consistent with σ other than product states. For example, each state $\sigma \in \mathbf{H}^+(\mathcal{X})$ may be *purified*. In other words, there always exists a complex Euclidean space \mathcal{Z} with $\dim \mathcal{Z} = \text{rank}(\sigma)$ and a pure state $u \in \mathcal{X} \otimes \mathcal{Z}$ such that

$$\sigma = \text{Tr}_{\mathcal{Z}}(uu^*).$$

The state uu^* is called a *purification* of σ .

Quantum operations

A *quantum operation* is a physically realizable discrete-time mapping (at least in an ideal sense) that takes as input a quantum state $\rho \in \mathbf{H}^+(\mathcal{X})$ of some system \mathcal{X} and produces as output a quantum state $\sigma \in \mathbf{H}^+(\mathcal{Y})$ of some system \mathcal{Y} . Every quantum operation is described uniquely by a completely positive and trace-preserving super-operator $\Psi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$, so that, for each input state ρ , the corresponding output state σ is given by

$$\sigma = \Psi(\rho).$$

As noted in Section 1.2.1, each quantum operation Ψ has a Stinespring representation whereby there exists a complex Euclidean space \mathcal{Z} with $\dim(\mathcal{Z}) = \text{rank}(J(\Psi))$ and an isometry $A : \mathcal{X} \rightarrow \mathcal{Y} \otimes \mathcal{Z}$ such that

$$\Phi : X \mapsto \text{Tr}_{\mathcal{Z}}(AXA^*).$$

Naturally, two quantum operations $\Phi_b : \mathbf{L}(\mathcal{X}_b) \rightarrow \mathbf{L}(\mathcal{Y}_b)$ for $b = 1, 2$ acting upon distinct physical systems may be viewed as one large quantum operation acting upon one large system via the tensor product:

$$\Phi_1 \otimes \Phi_2 : \mathbf{L}(\mathcal{X}_1 \otimes \mathcal{X}_2) \rightarrow \mathbf{L}(\mathcal{Y}_1 \otimes \mathcal{Y}_2).$$

Using our shorthand notation, this tensor product can alternately be written

$$\Phi_{1\dots 2} : \mathbf{L}(\mathcal{X}_{1\dots 2}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots 2}).$$

Quantum measurements

Many applications of quantum information necessitate the ability to extract classical information from a quantum state of some physical system \mathcal{X} . This extraction is accomplished via *quantum measurement*. Each measurement is uniquely specified by a finite set Γ of *measurement outcomes* and a finite set $\{P_a\}_{a \in \Gamma} \subset \mathbf{H}^+(\mathcal{X})$ of positive semidefinite *measurement operators* obeying the condition

$$\sum_{a \in \Gamma} P_a = I_{\mathcal{X}}.$$

(These measurement operators are sometimes called *POVM elements* for historical reasons—we avoid that terminology in this thesis.) For any state ρ of the system \mathcal{X} , the probability with which the measurement $\{P_a\}$ yields a particular outcome $a \in \Gamma$ is given by the inner product

$$\text{Pr}[\{P_a\} \text{ yields outcome } a \text{ on } \rho] = \langle P_a, \rho \rangle.$$

It is convenient to adopt the convention that the physical system associated with \mathcal{X} is destroyed when the measurement is applied. There is a simple and

common modification of the formalism of quantum measurements that admits non-destructive measurements that do not destroy the system upon which they act, but we do not require such a formalism in this thesis.

A measurement is called *projective* if each of its measurement operators is a projection. An arbitrary measurement $\{P_a\}_{a \in \Gamma} \subset \mathbf{H}^+(\mathcal{X})$ may be “simulated” by a projective measurement in the sense that there always exists a complex Euclidean space \mathcal{Z} , an isometry $A : \mathcal{X} \rightarrow \mathcal{X} \otimes \mathcal{Z}$ and a projective measurement $\{\Pi_a\}_{a \in \Gamma} \subset \mathbf{H}^+(\mathcal{X} \otimes \mathcal{Z})$ with the property that

$$\langle P_a, \rho \rangle = \langle \Pi_a, A\rho A^* \rangle = \langle A^* \Pi_a A, \rho \rangle$$

for every quantum state $\rho \in \mathbf{H}^+(\mathcal{X})$.

Part I

Quantum Strategies

Chapter 2

Introduction to Quantum Strategies

In this chapter we develop two distinct formalisms for quantum strategies. We begin with a naive operational formalism in Section 2.1, and we provide a new formalism in Section 2.2.

2.1 Operational formalism

In this section we develop an intuitive, operational formalism for quantum strategies and discuss some properties and problems associated with that formalism.

2.1.1 Formal definitions

At a high level, a *strategy* is a complete description of one party's actions in a multiple-round interaction involving the exchange of quantum information with one or more other parties. For convenience, let us call this party *Alice*. As we are only concerned for the moment with Alice's actions during the interaction, it is convenient to bundle the remaining parties into one party, whom we call *Bob*.

From Alice's point of view, every finite interaction decomposes naturally into a finite number r of *rounds*. In a typical round a message comes in, the message is processed, and a reply is sent out. Naturally, this reply might depend upon messages exchanged during previous rounds of the interaction. To account for such a dependence, we allow for a memory workspace to be maintained between rounds.

In order to facilitate discussion of the distinct strategies available to Alice for a given interaction, it is convenient to adopt the convention that the interaction specifies the number r of rounds and the size of each message in the interaction. In particular, individual strategies for Alice are *not* free to specify messages whose number or size deviates from those dictated by the interaction. By contrast, an

interaction does *not* dictate the size of the memory workspace to be used by each party. In particular, different strategies for Alice may call for different amounts of memory between rounds, and there is *no limit* on the amount of memory workspace a strategy may use. (However, we shall see in Chapter 3 that every possible strategy can be implemented with a reasonable fixed amount of memory workspace.)

The complex Euclidean spaces corresponding to the incoming and outgoing messages in an arbitrary round i shall be denoted \mathcal{X}_i and \mathcal{Y}_i , respectively. The space corresponding to the memory workspace to be stored for the next round shall be denoted \mathcal{Z}_i . As the size of each message is fixed by the interaction, the dimension of the spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ are also fixed by the interaction. Conversely, the dimension of the spaces $\mathcal{Z}_1, \dots, \mathcal{Z}_r$ may be arbitrarily large and these dimensions may vary among different strategies for the same interaction.

Thus, in a typical round i of the quantum interaction, Alice's actions are faithfully represented by a quantum operation

$$\Phi_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{Z}_{i-1}) \rightarrow \mathbf{L}(\mathcal{Y}_i \otimes \mathcal{Z}_i).$$

The first round of the interaction is a special case: there is no need for an incoming memory space for this round, so the quantum operation Φ_1 has the form

$$\Phi_1 : \mathbf{L}(\mathcal{X}_1) \rightarrow \mathbf{L}(\mathcal{Y}_1 \otimes \mathcal{Z}_1).$$

The final round of the interaction is also a special case: there is no immediate need for an outgoing memory space for this round. However, the presence of this final memory space better facilitates the forthcoming discussion of strategies involving measurements. Thus, the quantum operation Φ_r representing Alice's actions in the final round of the interaction has the same form as those from previous rounds:

$$\Phi_r : \mathbf{L}(\mathcal{X}_r \otimes \mathcal{Z}_{r-1}) \rightarrow \mathbf{L}(\mathcal{Y}_r \otimes \mathcal{Z}_r).$$

In order to extract classical information from the interaction, it suffices to permit Alice to perform a single quantum measurement on her final memory workspace. Sufficiency of a single measurement at the end of the interaction follows from foundational results on mixed state quantum computations [AKN98], which tell us that any quantum operation calling for one or more intermediate measurements can be efficiently simulated by an operation with a single measurement at the end.

Discussion thus far in this section is aptly summarized by the following formal definition.

Definition 2.1 (Operational definition of a strategy). Let $\mathcal{X}_1, \dots, \mathcal{X}_r$ and $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ be complex Euclidean spaces. An *operational r -round non-measuring strategy* for an interaction with *input spaces* $\mathcal{X}_1, \dots, \mathcal{X}_r$ and *output spaces* $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ consists of:

1. complex Euclidean spaces $\mathcal{Z}_1, \dots, \mathcal{Z}_r$, called *memory spaces*, and

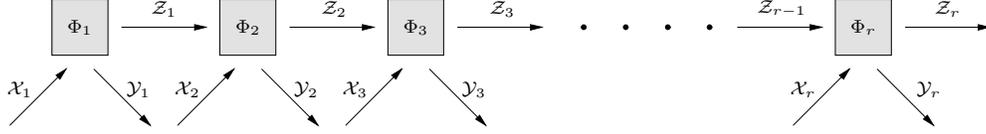


Figure 2.1: An r -round strategy.

2. an r -tuple of quantum operations (Φ_1, \dots, Φ_r) of the form

$$\begin{aligned} \Phi_1 &: \mathbf{L}(\mathcal{X}_1) \rightarrow \mathbf{L}(\mathcal{Y}_1 \otimes \mathcal{Z}_1) \\ \Phi_i &: \mathbf{L}(\mathcal{X}_i \otimes \mathcal{Z}_{i-1}) \rightarrow \mathbf{L}(\mathcal{Y}_i \otimes \mathcal{Z}_i) \quad (2 \leq i \leq r). \end{aligned}$$

An *operational r -round measuring strategy* with outcomes indexed by a consists of items 1 and 2 above, as well as:

3. a measurement $\{P_a\}$ on the last memory space \mathcal{Z}_r .

Figure 2.1 illustrates an r -round non-measuring strategy. □

Having formalized the notion of a strategy for Alice, it is straightforward to formalize the notion of a compatible strategy for Bob. Essentially, Bob must provide to Alice incoming messages with corresponding spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and accept her outgoing replies with corresponding spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$.

Definition 2.2 (Operational definition of a co-strategy). Let $\mathcal{X}_1, \dots, \mathcal{X}_r$ and $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ be complex Euclidean spaces. An *operational r -round non-measuring co-strategy* for an interaction with *input spaces* $\mathcal{X}_1, \dots, \mathcal{X}_r$ and *output spaces* $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ consists of:

1. complex Euclidean *memory spaces* $\mathcal{W}_0, \dots, \mathcal{W}_r$,
2. a quantum state $\rho_0 \in \mathbf{H}^+(\mathcal{X}_1 \otimes \mathcal{W}_0)$, and
3. a r -tuple of quantum operations (Ψ_1, \dots, Ψ_r) of the form

$$\begin{aligned} \Psi_i &: \mathbf{L}(\mathcal{Y}_i \otimes \mathcal{W}_{i-1}) \rightarrow \mathbf{L}(\mathcal{X}_{i+1} \otimes \mathcal{W}_i) \quad (1 \leq i \leq r-1) \\ \Psi_r &: \mathbf{L}(\mathcal{Y}_r \otimes \mathcal{W}_{r-1}) \rightarrow \mathbf{L}(\mathcal{W}_r). \end{aligned}$$

An *operational r -round measuring co-strategy* with outcomes indexed by b consists of items 1, 2 and 3 above, as well as:

4. a measurement $\{Q_b\}$ on the last memory space \mathcal{W}_r .

Figure 2.2 depicts the interaction between a r -round strategy and co-strategy. □

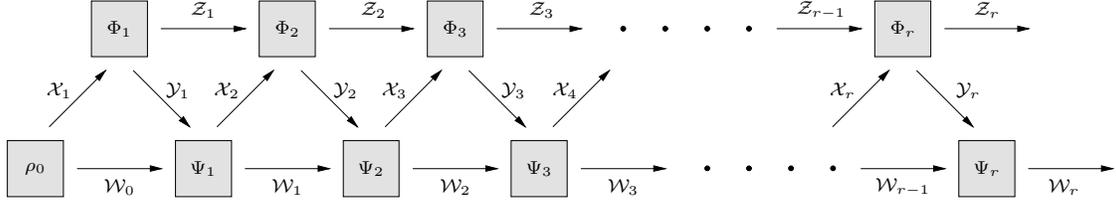


Figure 2.2: An interaction between a r -round strategy and co-strategy.

2.1.2 Immediate observations

Generality of the definition

Our definition of a strategy allows for trivial input or output spaces with dimension one—such spaces correspond to empty messages. Hence, simple actions such as the preparation of a quantum state or performing a measurement without producing a quantum output can be viewed as special cases of strategies.

While we require that Alice receive the first message and send the last message, any interaction that deviates from this format can easily be recast to fit this mold at the possible expense of one additional round of messages, some of which might be empty.

Indeed, our definition is flexible enough that any r -round co-strategy may equivalently be viewed as a $(r + 1)$ -round strategy with input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_r$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r, \mathbb{C}$. It is only for later convenience that we have chosen to define strategies and co-strategies as distinct objects.

Restriction to finite interactions

In this thesis we restrict our attention to interactions in which the size and number of messages exchanged is fixed. As we shall see in Section 2.2, this restriction permits us the luxury of representing every conceivable strategy for a given interaction by a single positive semidefinite operator of fixed finite dimension. Many of the results proven in this thesis rely crucially upon the finiteness of these representations of strategies.

At first, this restriction might seem overly constraining. For example, one might reasonably wish to study strategies for, say, coin-flipping protocols wherein two parties continue exchanging messages until they reach an agreement. With each round of messages, this agreement might be reached only probabilistically according to some quantum measurement. Such an interaction is infinite in the sense that it admits strategies with the property that, for every number r of rounds, there is a nonzero probability that it will not terminate in r rounds or fewer. (Indeed, this scenario also admits silly strategies in which the interaction *never* terminates.) To

forbid interactions such as this might seem careless, as our study aspires to absolute generality for quantum strategies.

While it would certainly be desirable to include in our formalism potentially infinite interactions, the generality that we sacrifice in the name of finiteness is less of a liability than it might at first seem. Many “infinite” interactions of any interest—such as the previous coin-flipping example—will terminate after a finite number of messages with probability 1. Interactions such as this can be approximated to arbitrary precision with only a finite number of messages simply by truncating the interaction appropriately.

Admittedly, the precise extent to which generality is sacrificed by our restriction to finite interactions remains to be seen. But it cannot be denied that the set of finite interactions encompasses a very wide swath of interesting quantum interactions. In the real world, all interactions are finite.

Restriction to isometric quantum operations

When convenient, we shall assume without loss of generality that each of Alice’s quantum operations Φ_i are actually linear isometries, meaning that

$$\Phi_i : X \mapsto A_i X A_i^*$$

for some linear isometry A_i . Similarly, we may assume that Bob’s initial state ρ_0 is actually a pure state, and that each of Bob’s quantum operations Ψ_i have the form

$$\Psi_i : X \mapsto B_i X B_i^*$$

for some linear isometry B_i . Moreover, if Alice’s or Bob’s strategy is a measuring strategy then we may also assume that the measurements $\{P_a\}$ for Alice and $\{Q_b\}$ for Bob are actually projective measurements, meaning that each P_a and Q_b is actually a projection operator.

These assumptions follow from the Stinespring representation for quantum operations mentioned in Section 1.2.1. Specifically, an arbitrary quantum operation $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ may be written $\Phi : X \mapsto \text{Tr}_{\mathcal{H}}(UXU^*)$ for some auxiliary space \mathcal{H} and some isometry $U : \mathcal{X} \rightarrow \mathcal{Y} \otimes \mathcal{H}$. As the memory spaces for strategies may be arbitrary, the auxiliary space \mathcal{H} for a given round i may be “absorbed” into the private memory space \mathcal{Z}_i of Alice’s quantum operation Φ_i , thus removing the need for the partial trace over \mathcal{H} and leaving us with $\Phi_i : X \mapsto A_i X A_i^*$ as desired.

The justification for projective measurements is similar: Naimark’s Theorem implies that any measurement on the final memory space \mathcal{Z}_r may be simulated by a projective measurement on $\mathcal{Z}_r \otimes \mathcal{H}$ for some auxiliary space \mathcal{H} . As before, \mathcal{H} may be “absorbed” into the definition of \mathcal{Z}_r , leaving us with only a projective measurement on this space.

Undesirable properties of the operational formalism

Definition 2.1 (Operational definition of a strategy) is very natural in the sense that it is clear that any conceivable actions taken by Alice during an interaction can be represented by a strategy of that form. Unfortunately, this definition has several undesirable mathematical properties that make this formalization cumbersome to use.

Picking an example arbitrarily, it is easy to see that the set of all strategies for a given interaction lacks a convenient distributive property for probabilistic combinations of strategies. In particular, suppose Alice plays according to $\Phi = (\Phi_1, \dots, \Phi_r)$ with probability p , otherwise she plays according to $\Phi' = (\Phi'_1, \dots, \Phi'_r)$. What is the r -tuple of quantum operations that describes this probabilistic combination of strategies? In an ideal world, this r -tuple would be given by the simple convex combination

$$p\Phi + (1 - p)\Phi'$$

where the i th component of the resulting r -tuple is given in the usual way by

$$p\Phi_i + (1 - p)\Phi'_i.$$

Alas, this identity does not always hold. Indeed, unless (Φ_1, \dots, Φ_r) and $(\Phi'_1, \dots, \Phi'_r)$ happen to agree on the choice of memory spaces $\mathcal{Z}_1, \dots, \mathcal{Z}_r$, the above convex combination is not even well-defined! Of course, this issue is easily overcome by “padding” the smaller memory spaces so that their dimensions agree with those of the larger spaces. But this frivolous observation ignores the larger problem.

Indeed, it is easy to exhibit a pair of strategies for which the above identity does not hold, even when the two strategies agree on the dimension of their memory spaces. The simple example we present here is a toy classical strategy expressed in the quantum formalism.

Example 2.3 (Lack of distributive property). Consider a two-round interaction wherein Alice’s incoming messages are empty. Letting ρ_α, ρ_β denote two distinct quantum states, the 2-tuples $(\rho_\alpha, \rho_\alpha)$ and (ρ_β, ρ_β) each specify a silly strategy for Alice that uses no memory space and returns the same state in both rounds. Suppose Alice plays one of $(\rho_\alpha, \rho_\alpha)$ and (ρ_β, ρ_β) uniformly at random. At the end of this interaction, Bob has received from Alice a combined state

$$\frac{1}{2}\rho_\alpha \otimes \rho_\alpha + \frac{1}{2}\rho_\beta \otimes \rho_\beta.$$

But the 2-tuple resulting from the naive computation

$$\frac{1}{2}(\rho_\alpha, \rho_\alpha) + \frac{1}{2}(\rho_\beta, \rho_\beta) = \left(\frac{\rho_\alpha + \rho_\beta}{2}, \frac{\rho_\alpha + \rho_\beta}{2} \right)$$

results in strategy for Alice in which Bob is instead left with the state

$$\frac{1}{4}\rho_\alpha \otimes \rho_\alpha + \frac{1}{4}\rho_\beta \otimes \rho_\beta + \frac{1}{4}\rho_\alpha \otimes \rho_\beta + \frac{1}{4}\rho_\beta \otimes \rho_\alpha$$

at the end of the interaction. As these final states for Bob are not equal, it follows that the naive convex combination $1/2(\rho_\alpha, \rho_\alpha) + 1/2(\rho_\beta, \rho_\beta)$ does not denote the desired probabilistic combination of the strategies described by $(\rho_\alpha, \rho_\alpha)$ and (ρ_β, ρ_β) .

□

This lack of a convenient distributive property suggests a cumbersome nonlinear dependence of probabilistic combinations $p\Phi + (1 - p)\Phi'$ of strategies upon their constituent strategies Φ, Φ' . Other examples of undesirable properties include:

Strategies are not unique. There exist pairs of strategies that differ in each component, yet they specify the same actions for an interaction in the sense that no interacting co-strategy could possibly distinguish the two.

This lack of uniqueness suggests that our operational formalism is carrying around some unnecessary information, which typically leads to unnecessary complication.

Outcomes depend nonlinearly on strategies. For a multi-round interaction between Alice and Bob, the probability of obtaining a given measurement outcome at the end of the interaction depends multilinearly upon *each component* Φ_i of Alice's strategy and Ψ_i of Bob's strategy. But the dependence of this probability as a function of the two r -tuples (Φ_1, \dots, Φ_r) and (Ψ_1, \dots, Ψ_r) is highly nonlinear.

As a consequence, existing algorithms for standard linear and semidefinite optimization problems cannot be employed to efficiently compute some important properties of strategies, such as the maximum probability with which a given strategy can be forced to yield a given measurement outcome.

2.2 New formalism

In the previous section we saw that the operational representation for quantum strategies has several drawbacks. In this section, we develop a new representation for quantum strategies that rectifies these problems.

2.2.1 Formal definitions

We begin with several new definitions for non-measuring and measuring quantum strategies and co-strategies.

Definition 2.4 (New formalism for non-measuring strategies). Let (Φ_1, \dots, Φ_r) be an operational representation of a non-measuring strategy for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$. Our new representation for quantum

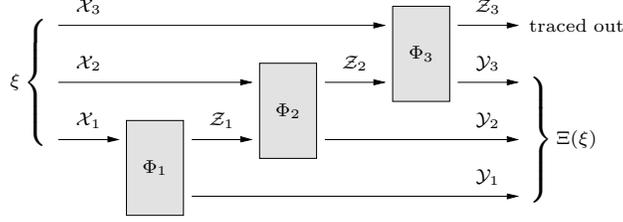


Figure 2.3: The super-operator Ξ associated with a three-round strategy.

strategies associates with every such operational strategy a single positive semidefinite operator

$$Q \in \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$$

via the following construction. Let

$$\Xi : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$$

denote the super-operator composed of Φ_1, \dots, Φ_r as suggested by Figure 2.3. Specifically, Ξ is the r -fold composition of Φ_1, \dots, Φ_r on the memory spaces $\mathcal{Z}_1, \dots, \mathcal{Z}_{r-1}$ followed by a partial trace on \mathcal{Z}_r . With some abuse of notation, this composition may be expressed succinctly as

$$\Xi = \text{Tr}_{\mathcal{Z}_r} \circ \Phi_r \circ \dots \circ \Phi_1.$$

(Here a tensor product with the identity super-operator $\mathbb{1}$ on the appropriate spaces is implicitly inserted where necessary in order for this composition to make sense.)

The r -round non-measuring strategy is given by $Q \stackrel{\text{def}}{=} J(\Xi)$. An operator Q is a valid representation of an r -round non-measuring strategy if and only if it has this form. \square

Definition 2.5 (New formalism for measuring strategies). With each operational representation $(\Phi_1, \dots, \Phi_r, \{P_a\})$ of an r -round measuring strategy for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ we associate a finite set

$$\{Q_a\} \subset \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$$

of positive semidefinite operators via the following construction. For each measurement outcome a let

$$\Xi_a = \Gamma_a \circ \Phi_r \circ \dots \circ \Phi_1$$

where the super-operator Γ_a is given by

$$\Gamma_a : X \mapsto \text{Tr}_{\mathcal{Z}_r} ((P_a \otimes I_{\mathcal{Y}_{1\dots r}}) X).$$

(Compare: for non-measuring strategies we defined Ξ via the partial trace $\text{Tr}_{\mathcal{Z}_r}$. For measuring strategies we define Ξ_a via Γ_a instead of the partial trace.)

The r -round measuring strategy $\{Q_a\}$ is given by $Q_a \stackrel{\text{def}}{=} J(\Xi_a)$ for each a . A set $\{Q_a\}$ of operators is a valid representation of an r -round measuring strategy if and only if it has this form. \square

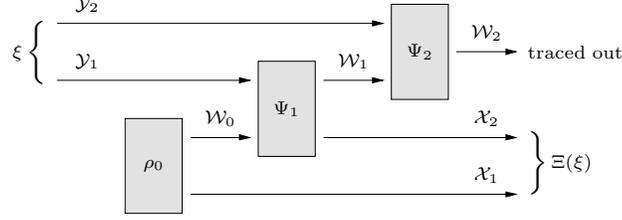


Figure 2.4: The super-operator Ξ associated with a two-round co-strategy.

Definition 2.6 (New formalism for co-strategies). Non-measuring co-strategies are defined in the same manner as non-measuring strategies, except the operator Q is given by $Q \stackrel{\text{def}}{=} J(\Xi^*)$ instead of $J(\Xi)$. Similarly, measuring co-strategies are defined in the same manner as measuring strategies, except the operators $\{Q_a\}$ are given by $Q_a \stackrel{\text{def}}{=} J(\Xi_a^*)$ instead of $J(\Xi_a)$ for each a .

Let us clarify this new representation for co-strategies. Let $(\rho_0, \Psi_1, \dots, \Psi_r)$ denote an operational representation of an r -round non-measuring co-strategy for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$, so that

$$\begin{aligned} \rho_0 &\in \mathbf{H}^+(\mathcal{X}_1 \otimes \mathcal{W}_0) \\ \Psi_i &: \mathbf{L}(\mathcal{Y}_i \otimes \mathcal{W}_{i-1}) \rightarrow \mathbf{L}(\mathcal{X}_{i+1} \otimes \mathcal{W}_i) \quad (1 \leq i \leq r-1) \\ \Psi_r &: \mathbf{L}(\mathcal{Y}_r \otimes \mathcal{W}_{r-1}) \rightarrow \mathbf{L}(\mathcal{W}_r). \end{aligned}$$

The induced super-operator $\Xi : \mathbf{L}(\mathcal{Y}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{X}_{1\dots r})$ is depicted in Figure 2.4 for the case $r = 2$. The adjoint super-operator has the form $\Xi^* : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$. In particular, we have

$$\begin{aligned} J(\Xi) &\in \mathbf{H}^+(\mathcal{X}_{1\dots r} \otimes \mathcal{Y}_{1\dots r}), \\ J(\Xi^*) &\in \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}). \end{aligned}$$

By taking $Q = J(\Xi^*)$, we ensure that strategies and co-strategies for the same input and output spaces lie within the same space $\mathbf{H}(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$ of Hermitian operators—with $\mathcal{Y}_{1\dots r}$ occurring *before* $\mathcal{X}_{1\dots r}$ in the tensor product.

Incidentally, it is instructive to note that $J(\Xi)$ is a valid representation of an $(r+1)$ -round *strategy* for input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_r$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r, \mathbb{C}$ according to Definition 2.4.

Conversely, if $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$ is a super-operator for which $J(\Lambda)$ is an r -round strategy for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ then $J(\Lambda^*)$ denotes a valid $(r+1)$ -round co-strategy for input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_r$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r, \mathbb{C}$. \square

2.2.2 Immediate observations

Some basic properties of this new representation for quantum strategies may be pointed out immediately. However, a more comprehensive discussion of basic prop-

erties of strategies must wait until Section 3.4, at which point the results of Chapter 3 are employed to establish with ease some additional basic properties.

Zero-round strategies

It is convenient to adopt the convention that an r -round non-measuring strategy or co-strategy Q denotes the scalar 1 when $r = 0$. Similarly, if $\{Q_a\}$ is a zero-round measuring strategy or co-strategy then $\{Q_a\}$ denotes a finite set of nonnegative real numbers that sum to one.

Representation in terms of isometries

Let (Φ_1, \dots, Φ_r) be an operational description of an r -round strategy. As discussed in Section 2.1.2, we may assume without loss of generality that each of the quantum operations Φ_i is actually an isometric operation, so that $\Phi_i : X \mapsto A_i X A_i^*$ for some isometry A_i . In this case, the super-operator Ξ in Definition 2.4 (New formalism for non-measuring strategies) has the form

$$\Xi : X \mapsto \text{Tr}_{\mathcal{Z}_r}(A X A^*)$$

where $A : \mathcal{X}_{1\dots r} \rightarrow \mathcal{Y}_{1\dots r} \otimes \mathcal{Z}_r$ is an isometry defined as the composition of A_1, \dots, A_r on the memory spaces $\mathcal{Z}_1, \dots, \mathcal{Z}_{r-1}$. With some abuse of notation, this composition may be expressed succinctly as

$$A = A_r A_{r-1} \cdots A_2 A_1$$

where a tensor product with the identity operator I on the appropriate spaces is inserted where necessary in order for this product to make sense.

Under the new formalism for strategies, the operator Q representing (Φ_1, \dots, Φ_r) has the special form

$$Q = J(\Xi) = \text{Tr}_{\mathcal{Z}_r}(\text{vec}(A) \text{vec}(A)^*).$$

For a measuring strategy with measurement $\{P_a\}$, each element of $\{Q_a\}$ may be written

$$Q_a = \text{Tr}_{\mathcal{Z}_r}((P_a \otimes I_{\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}}) \text{vec}(A) \text{vec}(A)^*).$$

If Q is a co-strategy, rather than a strategy, then by Definition 2.6 (New formalism for co-strategies) we have

$$Q = J(\Xi^*) = \text{Tr}_{\mathcal{Z}_r}(\text{vec}(A^*) \text{vec}(A^*)^*)$$

for non-measuring co-strategies and

$$Q_a = \text{Tr}_{\mathcal{Z}_r}((P_a \otimes I_{\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}}) \text{vec}(A^*) \text{vec}(A^*)^*)$$

for measuring co-strategies.

Measuring strategy elements sum to a non-measuring strategy

We conclude this chapter with an exercise that should help to wet our feet before diving into Chapter 3.

Proposition 2.7. *A set $\{Q_a\}$ of positive semidefinite operators is a measuring strategy or co-strategy if and only if $\sum_a Q_a$ is a non-measuring strategy or co-strategy, respectively.*

Proof. The proof for co-strategies is completely symmetric to the proof for strategies, so we only prove the proposition for strategies.

The “only if” portion of the proof is straightforward. As $\{Q_a\}$ is a measuring strategy, there is an r -tuple (A_1, \dots, A_r) of isometries and a measurement $\{P_a\}$ with the property that each Q_a is given by

$$Q_a = \text{Tr}_{\mathcal{Z}_r} ((P_a \otimes I_{\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}}) \text{vec}(A) \text{vec}(A)^*)$$

where the isometry A is given by the r -fold composition $A = A_r \cdots A_1$ is as noted earlier in this section. Then by linearity we have

$$\begin{aligned} \sum_a Q_a &= \text{Tr}_{\mathcal{Z}_r} \left(\left(\left(\sum_a P_a \right) \otimes I_{\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}} \right) \text{vec}(A) \text{vec}(A)^* \right) \\ &= \text{Tr}_{\mathcal{Z}_r} (\text{vec}(A) \text{vec}(A)^*), \end{aligned}$$

which is an r -round non-measuring strategy as required.

We now proceed to the “if” portion of the proof. As $\sum_a Q_a$ is a non-measuring strategy, there is an r -tuple (A_1, \dots, A_r) of isometries with the property that

$$\sum_a Q_a = \text{Tr}_{\mathcal{Z}_r} (\text{vec}(A) \text{vec}(A)^*)$$

where again the isometry

$$A : \mathcal{X}_{1\dots r} \rightarrow \mathcal{Y}_{1\dots r} \otimes \mathcal{Z}_r$$

is given by $A = A_r \cdots A_1$. Let

$$S : \mathcal{Z}_r \rightarrow \mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}$$

be the operator obtained from A by swapping the spaces $\mathcal{X}_{1\dots r}$ and \mathcal{Z}_r . (In other words, S is the image of A under the mapping $(y \otimes z)x^* \mapsto (y \otimes x)z^*$ on standard basis states.) We may thus write

$$\sum_a Q_a = \text{Tr}_{\mathcal{Z}_r} (\text{vec}(S) \text{vec}(S)^*) = SS^*.$$

Let T denote the Moore-Penrose pseudo-inverse of S , so that

$$ST = T^*S^* = \Pi_S$$

where Π_S denotes the projection onto the image of S . For each outcome a we let $P_a = TQ_aT^*$. As Q_a is positive semidefinite, it is clear that P_a is also positive semidefinite. Observe that $Q_a = SP_aS^*$, which follows from

$$SP_aS^* = STQ_aT^*S^* = \Pi_S Q_a \Pi_S = Q_a$$

where the final equality follows from the fact that the image of Q_a is contained in the image of S . Then

$$\begin{aligned} Q_a &= SP_aS^* \\ &= \text{Tr}_{Z_r}(\text{vec}(SP_a) \text{vec}(S)^*) \\ &= \text{Tr}_{Z_r}((P_a \otimes I_{Y_{1\dots r} \otimes X_{1\dots r}}) \text{vec}(S) \text{vec}(S)^*) \\ &= \text{Tr}_{Z_r}((P_a \otimes I_{Y_{1\dots r} \otimes X_{1\dots r}}) \text{vec}(A) \text{vec}(A)^*). \end{aligned}$$

As $\sum_a Q_a = \text{Tr}_{Z_r}(\text{vec}(A) \text{vec}(A)^*)$, it follows that

$$\sum_a P_a = I_{Z_r}.$$

As each P_a is positive semidefinite, the set $\{P_a\}$ is a quantum measurement and hence $\{Q_a\}$ is a measuring strategy. \square

Chapter 3

Three Important Properties

In this chapter we establish three fundamental and powerful properties of our new representation for quantum strategies. For ease of reference, we begin with rigorous statements for each of these properties. Their proofs appear in the subsequent sections, and the chapter concludes with a short list of other basic properties of strategies.

The first of these three properties establishes a bilinear dependence of the probability of a given measurement outcome upon the interacting strategy and co-strategy. (By contrast, recall that the operational representation for strategies led to a nonlinear dependence of outcomes on strategies in Section 2.1.2.)

Theorem 3.1 (Interaction output probabilities). *Let $\{Q_a\}$ be a measuring strategy and let $\{R_b\}$ be a compatible measuring co-strategy. For each pair (a, b) of measurement outcomes, the probability with which the interaction between $\{Q_a\}$ and $\{R_b\}$ yields (a, b) is given by the inner product $\langle Q_a, R_b \rangle$.*

The second property provides a recursive characterization of r -round strategies in terms of $(r - 1)$ -round strategies.

Theorem 3.2 (Characterization of strategies). *Let $Q \in \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$ be an arbitrary positive semidefinite operator. The following hold:*

1. *Q is an r -round strategy if and only if there exists an $(r - 1)$ -round strategy R with the property that $\text{Tr}_{\mathcal{Y}_r}(Q) = R \otimes I_{\mathcal{X}_r}$. Moreover, R is obtained from Q by terminating that strategy after $r - 1$ rounds.*
2. *Q is an r -round co-strategy if and only if there exists an operator R for which $Q = R \otimes I_{\mathcal{Y}_r}$ and $\text{Tr}_{\mathcal{X}_r}(R)$ is an $(r - 1)$ -round co-strategy. Moreover, $\text{Tr}_{\mathcal{X}_r}(R)$ is obtained from Q by terminating that co-strategy after $r - 1$ rounds.*
3. *Every r -round strategy or co-strategy Q may be described by isometries in such a way that the final memory space has dimension equal to $\text{rank}(Q)$.*

The recursive characterization of Theorem 3.2 may be equivalently expressed as an explicit list of linear constraints on positive semidefinite operators. These linear constraints are efficiently checkable and hence amenable to standard algorithms for semidefinite optimization problems.

Theorem 3.2 (Characterization of strategies, alternate version). *An operator $Q \in \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$ is an r -round strategy if and only if there exist operators*

$$Q_k \in \mathbf{H}^+(\mathcal{Y}_{1\dots k} \otimes \mathcal{X}_{1\dots k})$$

for $k = 1, \dots, r - 1$ such that

$$\begin{aligned} \text{Tr}_{\mathcal{Y}_{k\dots r}}(Q) &= Q_{k-1} \otimes I_{\mathcal{X}_{k\dots r}} \quad (2 \leq k \leq r), \\ \text{Tr}_{\mathcal{Y}_{1\dots r}}(Q) &= I_{\mathcal{X}_{1\dots r}}. \end{aligned}$$

Moreover, each Q_k is obtained from Q by terminating that strategy after k rounds.

It follows immediately from this characterization that the sets of strategies and co-strategies are compact and convex. For completeness, we provide a formal proof of this important fact as Proposition 3.15 in Section 3.4.

Our third property of strategies provides a formula for the maximum probability with which some co-strategy can force a given measuring strategy to output a given measurement outcome. Whereas the previous two properties are fundamental—both in their statements and their proofs—this third property is more advanced. Its proof relies crucially upon the previous two properties, as well as more advanced ideas from analysis such as convex polarity or semidefinite optimization duality.

Theorem 3.3 (Maximum output probability). *Let $\{Q_m\}$ be a measuring strategy. For each outcome m , the maximum probability with which $\{Q_m\}$ can be forced to output m by a compatible co-strategy is given by*

$$\min \{p \in [0, 1] : Q_m \preceq pR \text{ for some strategy } R\}.$$

An analogous result holds for co-strategies.

3.1 Interaction output probabilities

This section is devoted to a proof of Theorem 3.1 (Interaction output probabilities). We preface the proof with a discussion of a super-operator called the *contraction*, whose useful properties will be employed in the proof that follows.

The contraction operation

Informally speaking, the contraction operation is defined so that a composition of operators

$$B_r A_r B_{r-1} \cdots B_1 A_1$$

is the image under the contraction of the tensor product

$$B_r \otimes \cdots \otimes B_1 \otimes A_r \otimes \cdots \otimes A_1.$$

The ability to “unravel” operator compositions in this fashion is useful for our purpose because it allows us to isolate the actions of Alice (A_1, \dots, A_r) from those of Bob (B_1, \dots, B_r) in some r -round interaction.

Still speaking informally, the existence of an operation such as the contraction follows from the fact that the composition BA is multilinear in the operators A, B . Of course, the tensor product $A \otimes B$ is also multilinear in these operators. Indeed, as noted in Section 1.2.1, the tensor product possesses a special *universality* property whereby any multilinear mapping on A, B could equivalently be expressed as a linear mapping on the tensor product $A \otimes B$. In our case, that linear mapping is the contraction operation.

Let us formally define this operation.

Definition 3.4 (Contraction operation). The *contraction operation* in its full generality is more easily defined as a linear functional on vectors. For any complex Euclidean space \mathcal{V} we define

$$\text{contract} : \mathcal{V} \otimes \mathcal{V} \rightarrow \mathbb{C} : u \mapsto \left(\sum_{i=1}^{\dim(\mathcal{V})} e_i^* \otimes e_i^* \right) u$$

where $\{e_1, \dots, e_{\dim(\mathcal{V})}\}$ denotes the standard basis for \mathcal{V} .

Just as the trace function is tensored with the identity to yield the partial trace, the contraction is often tensored with the identity to yield the *partial contraction over* \mathcal{V} :

$$\text{contract}[\mathcal{V}] : \mathcal{V} \otimes \mathcal{A} \otimes \mathcal{V} \rightarrow \mathcal{A} : u \mapsto \left(\sum_{i=1}^{\dim(\mathcal{V})} e_i^* \otimes I_{\mathcal{A}} \otimes e_i^* \right) u.$$

(The exact ordering of the three spaces in the above tensor product $\mathcal{V} \otimes \mathcal{A} \otimes \mathcal{V}$ is immaterial—the contraction is defined similarly for other orderings.)

The partial contraction is often viewed as a super-operator as follows. If $u = \text{vec}(X)$ for some operator $X : \mathcal{V} \rightarrow \mathcal{V} \otimes \mathcal{A}$ then we may dispense with the $\text{vec}(\cdot)$ notation and simply write

$$\text{contract}[\mathcal{V}](X) = \sum_{i=1}^{\dim(\mathcal{V})} (e_i^* \otimes I_{\mathcal{A}}) X e_i.$$

Finally, we write $\text{contract}[\mathcal{V}, \mathcal{W}]$ as shorthand for the composition $\text{contract}[\mathcal{V}] \circ \text{contract}[\mathcal{W}]$. \square

Two useful properties of the contraction operation—including the ability to “unravel” operator compositions—are noted in the following proposition. Each item in this proposition is proven by a straightforward but tedious exercise in “index gymnastics.”

Proposition 3.5. *The following hold:*

1. *For any operators $A : \mathcal{X} \rightarrow \mathcal{A} \otimes \mathcal{Y}$ and $B : \mathcal{B} \otimes \mathcal{Y} \rightarrow \mathcal{Z}$ we have*

$$\text{contract}[\mathcal{Y}](A \otimes B) = (B \otimes I_{\mathcal{A}}) (A \otimes I_{\mathcal{B}}).$$

In particular, if $\mathcal{A} = \mathcal{B} = \mathbb{C}$ then $\text{contract}[\mathcal{Y}](A \otimes B) = BA$.

2. *For any operators $A : \mathcal{X} \rightarrow \mathcal{Y}$ and $B : \mathcal{Y} \rightarrow \mathcal{X}$ we have*

$$\text{contract}[\mathcal{X}, \mathcal{Y}](A \otimes B) = \text{Tr}(AB) = \langle A^*, B \rangle.$$

Proof. We begin with item 1. Throughout the proof it might be helpful to remember that $A \otimes B$ and its contraction take the following forms:

$$\begin{aligned} A \otimes B &: \mathcal{X} \otimes \mathcal{B} \otimes \mathcal{Y} \rightarrow \mathcal{A} \otimes \mathcal{Y} \otimes \mathcal{Z}, \\ \text{contract}[\mathcal{Y}](A \otimes B) &: \mathcal{X} \otimes \mathcal{B} \rightarrow \mathcal{A} \otimes \mathcal{Z}. \end{aligned}$$

We denote the standard bases of \mathcal{A} , \mathcal{B} , \mathcal{X} , \mathcal{Y} , and \mathcal{Z} by

$$\begin{aligned} \{a_1, \dots, a_{\dim(\mathcal{A})}\} &\subset \mathcal{A}, & \{x_1, \dots, x_{\dim(\mathcal{X})}\} &\subset \mathcal{X}, \\ \{b_1, \dots, b_{\dim(\mathcal{B})}\} &\subset \mathcal{B}, & \{y_1, \dots, y_{\dim(\mathcal{Y})}\} &\subset \mathcal{Y}, \\ & & \{z_1, \dots, z_{\dim(\mathcal{Z})}\} &\subset \mathcal{Z}. \end{aligned}$$

Then for some complex numbers $\alpha_{i,j,k}, \beta_{l,m,n} \in \mathbb{C}$ we may write

$$\begin{aligned} A &= \sum_{i=1}^{\dim(\mathcal{Y})} \sum_{j=1}^{\dim(\mathcal{X})} \sum_{k=1}^{\dim(\mathcal{A})} \alpha_{i,j,k} (a_k \otimes y_i) x_j^*, \\ B &= \sum_{l=1}^{\dim(\mathcal{Y})} \sum_{m=1}^{\dim(\mathcal{B})} \sum_{n=1}^{\dim(\mathcal{Z})} \beta_{l,m,n} z_n (b_m \otimes y_l)^*. \end{aligned}$$

By the definitions of the contraction and of matrix multiplication, we obtain

$$\begin{aligned} \text{contract}[\mathcal{Y}](A \otimes B) &= \sum_{o=1}^{\dim(\mathcal{Y})} (I_{\mathcal{A} \otimes \mathcal{Z}} \otimes y_o^*) (A \otimes B) (I_{\mathcal{X} \otimes \mathcal{B}} \otimes y_o) \\ &= \sum_{j,k,m,n} \left(\sum_{o=1}^{\dim(\mathcal{Y})} \alpha_{o,j,k} \beta_{o,m,n} \right) (a_k \otimes z_n) (x_j \otimes b_m)^* \\ &= (B \otimes I_{\mathcal{A}}) (A \otimes I_{\mathcal{B}}). \end{aligned}$$

For item 2, we denote the standard bases of \mathcal{X} and \mathcal{Y} by

$$\begin{aligned} \{x_1, \dots, x_{\dim(\mathcal{X})}\} &\subset \mathcal{X}, \\ \{y_1, \dots, y_{\dim(\mathcal{Y})}\} &\subset \mathcal{Y}. \end{aligned}$$

Then for some complex numbers $\alpha_{i,j}, \beta_{l,m} \in \mathbb{C}$ we may write

$$\begin{aligned} A &= \sum_{i=1}^{\dim(\mathcal{X})} \sum_{j=1}^{\dim(\mathcal{Y})} \alpha_{j,i} y_j x_i^*, \\ B &= \sum_{l=1}^{\dim(\mathcal{X})} \sum_{m=1}^{\dim(\mathcal{Y})} \beta_{l,m} x_l y_m^*. \end{aligned}$$

By the definitions of the contraction, matrix multiplication, and the trace we obtain

$$\begin{aligned} \text{contract}[\mathcal{X}, \mathcal{Y}](A \otimes B) &= \sum_{o=1}^{\dim(\mathcal{X})} \sum_{p=1}^{\dim(\mathcal{Y})} (y_p \otimes x_o)^* (A \otimes B) (x_o \otimes y_p) \\ &= \sum_{o=1}^{\dim(\mathcal{X})} \sum_{p=1}^{\dim(\mathcal{Y})} \alpha_{p,o} \beta_{o,p} \\ &= \text{Tr}(AB) = \langle A^*, B \rangle \end{aligned}$$

as desired. □

Proof of Theorem 3.1 (Interaction output probabilities)

We are now ready to provide the promised proof. The theorem is restated here for convenience.

Theorem 3.1 (Interaction output probabilities). *Let $\{Q_a\}$ be a measuring strategy and let $\{R_b\}$ be a compatible measuring co-strategy. For each pair (a, b) of measurement outcomes, the probability with which the interaction between $\{Q_a\}$ and $\{R_b\}$ yields (a, b) is given by the inner product $\langle Q_a, R_b \rangle$.*

Proof. Suppose $\{Q_a\}$ is described by isometries A_1, \dots, A_r and a projective measurement $\{\Pi_a\}$. Similarly, suppose $\{R_b\}$ is described by a pure state u_0 , isometries B_1, \dots, B_r , and a projective measurement $\{\Delta_b\}$. For each pair (a, b) let $v_{a,b} \in \mathcal{Z}_r \otimes \mathcal{W}_r$ denote the vector obtained by applying the measurement operator $\Pi_a \otimes \Delta_b$ to the pure state of the entire system at the end of the interaction. That is, $v_{a,b}$ is given by

$$v_{a,b} \stackrel{\text{def}}{=} (\Pi_a \otimes \Delta_b) B_r A_r B_{r-1} \cdots B_1 A_1 u_0.$$

(Our notation here suppresses the numerous tensors with identity.) Then the desired probability is equal to $\|v_{a,b}\|^2$. The remainder of this proof is dedicated to proving that

$$\|v_{a,b}\|^2 = \langle Q_a, R_b \rangle.$$

We use the contraction operation to pull apart the composition of $v_{a,b}$ and express it as a contraction of the tensor product

$$(\Pi_a \otimes I_{\mathcal{Y}_r}) A_r \otimes A_{r-1} \otimes \cdots \otimes A_1 \otimes \Delta_b B_r \otimes B_{r-1} \otimes \cdots \otimes B_1 \otimes u_0$$

over every space except \mathcal{Z}_r and \mathcal{W}_r . For convenience, this expression for $v_{a,b}$ is written

$$v_{a,b} = \text{contract}[\mathcal{X}_1, \dots, \mathcal{X}_r, \mathcal{Y}_1, \dots, \mathcal{Y}_r] (A \otimes B)$$

where the operators

$$\begin{aligned} A &: \mathcal{X}_{1\dots r} \rightarrow \mathcal{Y}_{1\dots r} \otimes \mathcal{Z}_r, \\ B &: \mathcal{Y}_{1\dots r} \rightarrow \mathcal{X}_{1\dots r} \otimes \mathcal{W}_r \end{aligned}$$

are given by

$$\begin{aligned} A &\stackrel{\text{def}}{=} \text{contract}[\mathcal{Z}_1, \dots, \mathcal{Z}_{r-1}] ((\Pi_a \otimes I_{\mathcal{Y}_r}) A_r \otimes A_{r-1} \otimes \cdots \otimes A_1), \\ B &\stackrel{\text{def}}{=} \text{contract}[\mathcal{W}_0, \dots, \mathcal{W}_{r-1}] (\Delta_b B_r \otimes B_{r-1} \otimes \cdots \otimes B_1 \otimes u_0). \end{aligned}$$

With an eye toward the end of the proof, we observe that

$$\begin{aligned} Q_a &= \text{Tr}_{\mathcal{Z}_r} (\text{vec}(A) \text{vec}(A)^*), \\ R_b &= \text{Tr}_{\mathcal{W}_r} (\text{vec}(B^*) \text{vec}(B^*)^*) \end{aligned}$$

as per Definition 2.4 (New formalism for non-measuring strategies).

Let $\{e_1, \dots, e_{\dim(\mathcal{Z}_r)}\}$ and $\{f_1, \dots, f_{\dim(\mathcal{W}_r)}\}$ denote the standard bases of \mathcal{Z}_r and \mathcal{W}_r , respectively, and for each i, j define the operators

$$\begin{aligned} A_{(i)} &: \mathcal{X}_{1\dots r} \rightarrow \mathcal{Y}_{1\dots r}, \\ B_{(j)} &: \mathcal{Y}_{1\dots r} \rightarrow \mathcal{X}_{1\dots r} \end{aligned}$$

by

$$\begin{aligned} A_{(i)} &\stackrel{\text{def}}{=} (e_i^* \otimes I_{\mathcal{Y}_{1\dots r}}) A, \\ B_{(j)} &\stackrel{\text{def}}{=} (f_j^* \otimes I_{\mathcal{X}_{1\dots r}}) B. \end{aligned}$$

Again, with an eye toward the end of the proof we observe that

$$\begin{aligned} \sum_{i=1}^{\dim(\mathcal{Z}_r)} \text{vec}(A_{(i)}) \text{vec}(A_{(i)})^* &= \text{Tr}_{\mathcal{Z}_r} (\text{vec}(A) \text{vec}(A)^*) = Q_a, \\ \sum_{j=1}^{\dim(\mathcal{W}_r)} \text{vec}(B_{(j)}^*) \text{vec}(B_{(j)}^*)^* &= \text{Tr}_{\mathcal{W}_r} (\text{vec}(B^*) \text{vec}(B^*)^*) = R_b. \end{aligned}$$

The (i, j) th component of $v_{a,b}$ (in the standard basis) is given by

$$(e_i^* \otimes f_j^*) v_{a,b} = \text{contract}[\mathcal{X}_1, \dots, \mathcal{X}_r, \mathcal{Y}_1, \dots, \mathcal{Y}_r] (A_{(i)} \otimes B_{(j)}) = \langle B_{(j)}^*, A_{(i)} \rangle.$$

Employing the cyclic property of the trace, we find that the modulus squared of the (i, j) th component of $v_{a,b}$ is

$$\langle B_{(j)}^*, A_{(i)} \rangle \langle A_{(i)}, B_{(j)}^* \rangle = \langle \text{vec}(A_{(i)}) \text{vec}(A_{(i)})^*, \text{vec}(B_{(j)}^*) \text{vec}(B_{(j)}^*)^* \rangle.$$

Then the desired norm $\|v_{a,b}\|^2$ is the sum of these moduli squared:

$$\begin{aligned} \|v_{a,b}\|^2 &= \sum_{i,j} \langle \text{vec}(A_{(i)}) \text{vec}(A_{(i)})^*, \text{vec}(B_{(j)}^*) \text{vec}(B_{(j)}^*)^* \rangle \\ &= \left\langle \sum_{i=1}^{\dim(\mathcal{Z}_r)} \text{vec}(A_{(i)}) \text{vec}(A_{(i)})^*, \sum_{j=1}^{\dim(\mathcal{W}_r)} \text{vec}(B_{(j)}^*) \text{vec}(B_{(j)}^*)^* \right\rangle \\ &= \langle Q_a, R_b \rangle. \end{aligned}$$

□

3.2 Characterization of strategies

This section is devoted to a proof of Theorem 3.2 (Characterization of strategies). Before providing the proof, it is appropriate to comment on the relationship between our characterization of strategies and prior work on so-called “no-signaling” quantum operations.

Relationship between strategies and no-signaling operations

The content of Theorem 3.2 was originally established in 2002 within the context of no-signaling operations via the combined work of Beckman *et al.* [BGNP01] and Eggeling, Schlingemann, and Werner [ESW02]. The proof presented in this thesis was developed by the present author and Watrous in 2007 [GW07] within the context of quantum strategies and without any knowledge of this prior work.

Let us elaborate upon the connection between strategies and no-signaling operations. Simply put, Theorem 3.2 states that Q is an r -round strategy if and only if it obeys the partial trace condition

$$\text{Tr}_{\mathcal{Y}_r}(Q) = R \otimes I_{\mathcal{X}_r}$$

for some $(r - 1)$ -round strategy R . This partial trace condition also appears in Beckman *et al.* [BGNP01, Theorem 8], wherein it was established that the above condition also captures the *one-directional no-signaling* property of certain quantum operations.

For quantum strategies, this property means that it is impossible to send information (a “signal”) from the system \mathcal{X}_i to any of the systems $\mathcal{Y}_1, \dots, \mathcal{Y}_{i-1}$ for each i . Intuitively, we should expect quantum strategies to obey a no-signaling

condition of this sort. Any “strategy” that disobeys this condition could be used to communicate backwards in time—from future rounds of interaction to previous rounds. A world that permitted such clairvoyant strategies would be an interesting world indeed! Alas, such a world is not causally consistent.

While Beckman *et al.* showed that the partial trace condition of Theorem 3.2 is a *necessary* condition for quantum strategies, it was Eggeling, Schlingemann, and Werner who established that this condition is also *sufficient*. In particular, they showed that any quantum operation that forbids signaling in *one* direction may be implemented by two separate quantum operations, possibly with communication in the *other* direction.

For quantum strategies, signaling is forbidden in the future-past direction, but permitted in the past-future direction. In this context, such past-future communication is better known as “memory” and the two separate quantum operations can be taken to represent the actions of the strategy in two distinct rounds of the interaction. In this way, a quantum strategy is constructed from an operation that obeys the one-directional no-signaling operation, just as in Theorem 3.2.

Proof of Theorem 3.2 (Characterization of strategies)

Let us present the promised proof. The theorem is restated here for convenience.

Theorem 3.2 (Characterization of strategies). *Let $Q \in \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$ be an arbitrary positive semidefinite operator. The following hold:*

1. *Q is an r -round strategy if and only if there exists an $(r - 1)$ -round strategy R with the property that $\text{Tr}_{\mathcal{Y}_r}(Q) = R \otimes I_{\mathcal{X}_r}$. Moreover, R is obtained from Q by terminating that strategy after $r - 1$ rounds.*
2. *Q is an r -round co-strategy if and only if there exists an operator R for which $Q = R \otimes I_{\mathcal{Y}_r}$ and $\text{Tr}_{\mathcal{X}_r}(R)$ is an $(r - 1)$ -round co-strategy. Moreover, $\text{Tr}_{\mathcal{X}_r}(R)$ is obtained from Q by terminating that co-strategy after $r - 1$ rounds.*
3. *Every r -round strategy or co-strategy Q may be described by isometries in such a way that the final memory space has dimension equal to $\text{rank}(Q)$.*

Proof. We begin with a proof of item 1. Along the way, we will also prove item 3. Item 2 follows from item 1—a fact we establish at the end of this proof.

Suppose first that Q is an r -round strategy and let R denote the $(r - 1)$ -round strategy obtained from Q by terminating that strategy after the first $r - 1$ rounds. We will prove that $\text{Tr}_{\mathcal{Y}_r}(Q) = R \otimes I_{\mathcal{X}_r}$.

Toward that end, let Ξ_r, Ξ_{r-1} be the quantum operations satisfying $Q = J(\Xi_r)$ and $R = J(\Xi_{r-1})$. As illustrated in Figure 3.1, is clear that the super-operators $(\text{Tr}_{\mathcal{Y}_r} \circ \Xi_r)$ and $(\Xi_{r-1} \otimes \text{Tr}_{\mathcal{X}_r})$ are equal. We have

$$\text{Tr}_{\mathcal{Y}_r}(Q) = \text{Tr}_{\mathcal{Y}_r}(J(\Xi_r)) = J(\text{Tr}_{\mathcal{Y}_r} \circ \Xi_r) = J(\Xi_{r-1} \otimes \text{Tr}_{\mathcal{X}_r}) = R \otimes I_{\mathcal{X}_r}$$

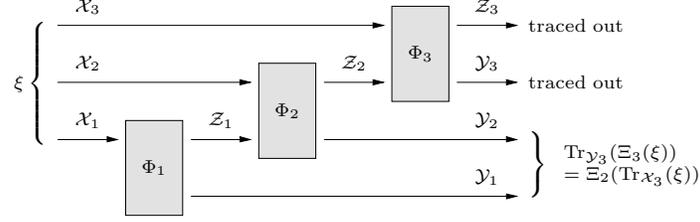


Figure 3.1: For a three-round strategy, the super-operators $(\text{Tr}_{\mathcal{Y}_3} \circ \Xi_3)$ and $(\Xi_2 \otimes \text{Tr}_{\mathcal{X}_3})$ are equal.

as desired.

Next, assume that Q satisfies $\text{Tr}_{\mathcal{Y}_r}(Q) = R \otimes I_{\mathcal{X}_r}$ for some $(r-1)$ -round strategy R . We will prove that Q is an r -round strategy. Along the way, we will also establish item 3 in the statement of the theorem.

This portion of the proof is by induction on the number of rounds r . We begin with the base case $r = 1$. It is clear that Q is a one-round strategy if and only if $Q = J(\Phi_1)$ for some quantum operation $\Phi_1 : \mathbf{L}(\mathcal{X}_1) \rightarrow \mathbf{L}(\mathcal{Y}_1)$. As Φ_1 is a quantum operation, it must be the case that $\text{Tr}_{\mathcal{Y}_1}(Q) = I_{\mathcal{X}_1}$. By convention, the zero-round strategy R that we seek must be represented by the scalar 1. The desired expression $\text{Tr}_{\mathcal{Y}_1}(Q) = R \otimes I_{\mathcal{X}_1}$ follows from the simple observation that $I_{\mathcal{X}_1} = 1 \otimes I_{\mathcal{X}_1}$. Item 3 for this case follows from the usual Stinespring representation: there exists a space \mathcal{Z}_1 with $\dim(\mathcal{Z}_1) = \text{rank}(Q)$ and an isometry $A_1 : \mathcal{X}_1 \rightarrow \mathcal{Y}_1 \otimes \mathcal{Z}_1$ such that $\Phi_1 : X \mapsto \text{Tr}_{\mathcal{Z}_1}(A_1 X A_1^*)$.

For the case $r \geq 2$, let A_1, \dots, A_{r-1} be isometries that describe R and let $A = A_{r-1} \cdots A_1$ denote the $(r-1)$ -fold composition of these isometries, so that

$$R = \text{Tr}_{\mathcal{Z}_{r-1}}(\text{vec}(A) \text{vec}(A)^*).$$

By the induction hypothesis, the memory space \mathcal{Z}_{r-1} has $\dim(\mathcal{Z}_{r-1}) = \text{rank}(R)$. As required for item 3, we let \mathcal{Z}_r be a complex Euclidean space with dimension equal to $\text{rank}(Q)$. Let $B : \mathcal{X}_{1..r} \rightarrow \mathcal{Y}_{1..r} \otimes \mathcal{Z}_r$ be any operator satisfying

$$\text{Tr}_{\mathcal{Z}_r}(\text{vec}(B) \text{vec}(B)^*) = Q.$$

Such a choice of B must exist given that the dimension of \mathcal{Z}_r is large enough to admit a purification of Q . Note that $\text{vec}(B)$ is also a purification of $R \otimes I_{\mathcal{X}_r}$:

$$\text{Tr}_{\mathcal{Y}_r \otimes \mathcal{Z}_r}(\text{vec}(B) \text{vec}(B)^*) = \text{Tr}_{\mathcal{Y}_r}(Q) = R \otimes I_{\mathcal{X}_r}.$$

We will now identify a second purification of $R \otimes I_{\mathcal{X}_r}$. Toward that end, let \mathcal{V} be a complex Euclidean space with $\dim(\mathcal{V}) = \dim(\mathcal{X}_r)$ and let $V : \mathcal{X}_r \rightarrow \mathcal{V}$ be an arbitrary unitary operator. Then

$$\text{Tr}_{\mathcal{V}}(\text{vec}(V) \text{vec}(V)^*) = I_{\mathcal{X}_r}$$

and so

$$\mathrm{Tr}_{\mathcal{Z}_{r-1} \otimes \mathcal{V}}(\mathrm{vec}(A \otimes V) \mathrm{vec}(A \otimes V)^*) = R \otimes I_{\mathcal{X}_r}.$$

We will now use the isometric equivalence of purifications to define an isometry

$$A_r : \mathcal{X}_r \otimes \mathcal{Z}_{r-1} \rightarrow \mathcal{Y}_r \otimes \mathcal{Z}_r$$

for the r th round that will complete the proof. Because $\mathcal{Z}_{r-1} \otimes \mathcal{V}$ has the minimal dimension required to admit a purification of $R \otimes I_{\mathcal{X}_r}$, it follows that there exists an isometry $U : \mathcal{Z}_{r-1} \otimes \mathcal{V} \rightarrow \mathcal{Y}_r \otimes \mathcal{Z}_r$ such that

$$(I_{\mathcal{Y}_{1\dots r-1}} \otimes U \otimes I_{\mathcal{X}_{1\dots r}}) \mathrm{vec}(A \otimes V) = \mathrm{vec}(B).$$

This expression may equivalently be written

$$B = (I_{\mathcal{Y}_{1\dots r-1}} \otimes U) (A \otimes V).$$

We now define $A_r \stackrel{\mathrm{def}}{=} U(I_{\mathcal{Z}_{r-1}} \otimes V)$ so that

$$B = (I_{\mathcal{Y}_{1\dots r-1}} \otimes A_r)(A \otimes I_{\mathcal{X}_r}).$$

In other words, B is given by the r -fold composition $A_r \cdots A_1$. As

$$Q = \mathrm{Tr}_{\mathcal{Z}_r}(\mathrm{vec}(B) \mathrm{vec}(B)^*),$$

it follows that Q is an r -round strategy described by the isometries A_1, \dots, A_r . The proofs of items 1 and 3 are thus complete.

Let us now prove item 2. Essentially, the proof consists of several applications of item 1 with different choices of input and output spaces.

As noted in Definition 2.6 (New formalism for co-strategies), each r -round co-strategy may be viewed as an $(r+1)$ -round strategy with input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_r$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r, \mathbb{C}$ and *vice versa*. With this fact in mind, item 1 tells us

$$\begin{aligned} & Q \text{ is an } (r+1)\text{-round strategy} \\ & \quad \text{for input spaces } \mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_r \text{ and output spaces } \mathcal{X}_1, \dots, \mathcal{X}_r, \mathbb{C} \\ \iff & \mathrm{Tr}_{\mathbb{C}}(Q) = R \otimes I_{\mathcal{Y}_r} \text{ where } R \text{ is an } r\text{-round strategy} \\ & \quad \text{for the input spaces } \mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1} \text{ and output spaces } \mathcal{X}_1, \dots, \mathcal{X}_r. \end{aligned}$$

(Of course, $\mathrm{Tr}_{\mathbb{C}}(X) = X$ for every X .) To complete the proof, it suffices to show that

$$\begin{aligned} & R \text{ is an } r\text{-round strategy} \\ & \quad \text{for input spaces } \mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1} \text{ and output spaces } \mathcal{X}_1, \dots, \mathcal{X}_r \\ \iff & \mathrm{Tr}_{\mathcal{X}_r}(R) \text{ is an } r\text{-round strategy} \\ & \quad \text{for input spaces } \mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1} \text{ and output spaces } \mathcal{X}_1, \dots, \mathcal{X}_{r-1}, \mathbb{C}. \end{aligned}$$

For once this equivalence is shown, item 2 follows from the fact that the linear conditions in question are robust with respect to taking adjoints. In particular, it holds that

$$\mathrm{Tr}_{\mathcal{X}}(J(\Xi)) = J(\Lambda) \otimes I_{\mathcal{Y}} \iff \mathrm{Tr}_{\mathcal{X}}(J(\Xi^*)) = J(\Lambda^*) \otimes I_{\mathcal{Y}}$$

for all super-operators Ξ, Λ and all appropriate choices of spaces \mathcal{X}, \mathcal{Y} .

It remains only to establish the stated equivalence between R and $\mathrm{Tr}_{\mathcal{X}_r}(R)$. Toward that end, suppose first that R is an r -round strategy for the input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1}$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$. By item 1 it follows that

$$\mathrm{Tr}_{\mathcal{X}_r}(R) = S \otimes I_{\mathcal{Y}_{r-1}}$$

where S is an $(r-1)$ -round strategy for the input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-2}$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_{r-1}$. It is not hard to see that $S \otimes I_{\mathcal{Y}_{r-1}}$ denotes a valid r -round strategy for the input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1}$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_{r-1}, \mathbb{C}$. Hence, $\mathrm{Tr}_{\mathcal{X}_r}(R)$ is as claimed.

Conversely, suppose that $\mathrm{Tr}_{\mathcal{X}_r}(R)$ is an r -round strategy for the input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1}$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_{r-1}, \mathbb{C}$. By item 1 it follows that

$$\mathrm{Tr}_{\mathbb{C}}(\mathrm{Tr}_{\mathcal{X}_r}(R)) = S \otimes I_{\mathcal{Y}_{r-1}}$$

where S is an $(r-1)$ -round strategy for the input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-2}$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_{r-1}$. By item 1 again, it must be that R is an r -round strategy for the input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_{r-1}$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ as desired. \square

3.3 Maximum output probabilities

This section is devoted to a proof of Theorem 3.3 (Maximum output probability). We offer two distinct proofs of this theorem. The first is provided in Section 3.3.1 and employs the formalism of convex polarity, while the second is provided in Section 3.3.2 and employs the formalism of semidefinite optimization duality. Both proofs make use of the fact that the sets of strategies and co-strategies are compact and convex—an immediate implication of Theorem 3.2 that was noted at the start of this chapter and shall be proven at the end of this chapter.

The theorem is restated here for convenience and is followed by a simple lemma employed in both our proofs.

Theorem 3.3 (Maximum output probability). *Let $\{Q_m\}$ be a measuring strategy. For each outcome m , the maximum probability with which $\{Q_m\}$ can be forced to output m by a compatible co-strategy is given by*

$$\min \{p \in [0, 1] : Q_m \preceq pR \text{ for some strategy } R\}.$$

An analogous result holds for co-strategies.

Lemma 3.6. *Suppose that Theorem 3.3 is known to hold for strategies. Then Theorem 3.3 also holds for co-strategies. The converse is also true.*

Proof. Suppose that Theorem 3.3 is known to hold for strategies. Let $\{R_b\}$ be an r -round measuring co-strategy for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$. For each outcome b let $\Xi_b : \mathbf{L}(\mathcal{Y}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{X}_{1\dots r})$ be the super-operator with $R_b = J(\Xi_b^*)$. Then $\{J(\Xi_b)\}$ denotes an $(r+1)$ -round measuring strategy for input spaces $\mathbb{C}, \mathcal{Y}_1, \dots, \mathcal{Y}_r$ and output spaces $\mathcal{X}_1, \dots, \mathcal{X}_r, \mathbb{C}$. As such, it holds that

$$\begin{aligned} & \max \{ \langle J(\Xi_b), J(\Phi) \rangle : J(\Phi) \text{ is a compatible co-strategy} \} \\ &= \min \{ p \in [0, 1] : J(\Xi_b) \preceq pJ(\Psi) \text{ for some strategy } J(\Psi) \}. \end{aligned}$$

The result follows from the fact that $\langle J(\Xi_b), J(\Phi) \rangle = \langle J(\Xi_b^*), J(\Phi^*) \rangle$ and

$$J(\Xi_b) \preceq pJ(\Psi) \iff J(\Xi_b^*) \preceq pJ(\Psi^*)$$

for all choices of super-operators Ξ_b, Φ , and Ψ . The proof of the converse statement is identical. \square

3.3.1 Proof by convex polarity

Our first proof of Theorem 3.3 (Maximum output probability) employs a notion from convex analysis known as ‘‘polarity.’’ But before we discuss this notion in detail, let us first introduce some notation. We let

$$\mathbf{S}_r \subset \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$$

denote the set of all r -round non-measuring strategies for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$. Similarly, we let

$$\text{co-}\mathbf{S}_r \subset \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$$

denote the set of all r -round non-measuring co-strategies for these input and output spaces. In keeping with our convention, we have $\mathbf{S}_0 = \text{co-}\mathbf{S}_0 = \{1\}$. We mentioned earlier that the sets \mathbf{S}_r and $\text{co-}\mathbf{S}_r$ are compact and convex.

For any set \mathbf{C} of positive semidefinite operators, we write

$$\downarrow \mathbf{C} \stackrel{\text{def}}{=} \{X : 0 \preceq X \preceq Y \text{ for some } Y \in \mathbf{C}\}.$$

A key component of this proof of Theorem 3.3 is a characterization of the polar sets of $\downarrow \mathbf{S}_r$ and $\downarrow \text{co-}\mathbf{S}_r$.

Introduction to polarity

For any non-empty set \mathbf{C} of Hermitian operators, the *polar* \mathbf{C}° of \mathbf{C} is defined as

$$\mathbf{C}^\circ \stackrel{\text{def}}{=} \{A : \langle B, A \rangle \leq 1 \text{ for all } B \in \mathbf{C}\}$$

and the *support* and *gauge* functions for \mathbf{C} are defined as

$$\begin{aligned} s(X \mid \mathbf{C}) &\stackrel{\text{def}}{=} \sup \{\langle X, Y \rangle : Y \in \mathbf{C}\} \\ g(X \mid \mathbf{C}) &\stackrel{\text{def}}{=} \inf \{\lambda \geq 0 : X \in \lambda \mathbf{C}\}. \end{aligned}$$

Let us list some basic facts concerning these objects.

Proposition 3.7. *Let \mathbf{C}, \mathbf{D} be non-empty sets of Hermitian operators. The following hold:*

1. *If $\mathbf{C} \subseteq \mathbf{D}$ then $\mathbf{D}^\circ \subseteq \mathbf{C}^\circ$.*
2. *If $-X \in \mathbf{C}$ for each positive semidefinite operator X then every element of \mathbf{C}° is positive semidefinite.*
3. *If \mathbf{C} is closed, convex, and contains the origin then the same is true of \mathbf{C}° . In this case we have $\mathbf{C}^{\circ\circ} = \mathbf{C}$ and $s(\cdot \mid \mathbf{C}) = g(\cdot \mid \mathbf{C}^\circ)$.*

The first two items of Proposition 3.7 are elementary; a proof of the third may be found in Rockafellar [Roc70].

A characterization of polar sets of strategies

We now establish a useful characterization of polar sets derived from strategies and co-strategies.

Proposition 3.8 (Polar sets of strategies). *The following polarity relations hold for the sets $\downarrow \mathbf{S}_r$ and $\downarrow \text{co-}\mathbf{S}_r$:*

$$\begin{aligned} (\downarrow \mathbf{S}_r)^\circ &= \{X : X \preceq Q \text{ for some } Q \in \text{co-}\mathbf{S}_r\}, \\ (\downarrow \text{co-}\mathbf{S}_r)^\circ &= \{X : X \preceq Q \text{ for some } Q \in \mathbf{S}_r\}. \end{aligned}$$

We begin with a simple observation implying that the two equalities in Proposition 3.8 are equivalent.

Lemma 3.9. *Let \mathbf{A}, \mathbf{B} be non-empty, closed, and convex sets of positive semidefinite operators, and suppose*

$$(\downarrow \mathbf{A})^\circ = \{X : X \preceq Q \text{ for some } Q \in \mathbf{B}\}.$$

Then

$$(\downarrow \mathbf{B})^\circ = \{Y : Y \preceq R \text{ for some } R \in \mathbf{A}\}.$$

Proof. Let

$$\mathbf{C} = \{Y : Y \preceq R \text{ for some } R \in \mathbf{A}\}.$$

The lemma is proved by showing $\mathbf{C}^\circ = \downarrow \mathbf{B}$, from which the desired result $\mathbf{C} = (\downarrow \mathbf{B})^\circ$ immediately follows. Let us start by proving $\mathbf{C}^\circ \subseteq \downarrow \mathbf{B}$. As $-P \in \mathbf{C}$ for every positive semidefinite P , it follows that every element of the polar \mathbf{C}° is positive semidefinite. Clearly $\downarrow \mathbf{A} \subseteq \mathbf{C}$, and therefore $\mathbf{C}^\circ \subseteq (\downarrow \mathbf{A})^\circ$. By definition, $\downarrow \mathbf{B}$ consists of the positive semidefinite elements of $(\downarrow \mathbf{A})^\circ$, so we have $\mathbf{C}^\circ \subseteq \downarrow \mathbf{B}$.

On the other hand, we have that every $Q \in \downarrow \mathbf{B}$ is contained in $(\downarrow \mathbf{A})^\circ$, implying that $\langle Q, R \rangle \leq 1$ for all $R \in \mathbf{A}$. As Q is positive semidefinite, this also implies that $\langle Q, X \rangle \leq 1$ for all $X \preceq R$. Consequently, $Q \in \mathbf{C}^\circ$. Thus $\downarrow \mathbf{B} = \mathbf{C}^\circ$ as desired. \square

We also require a technical statement that simplifies computations involving $\downarrow \mathbf{S}_r$ and $\downarrow \text{co-}\mathbf{S}_r$.

Lemma 3.10. *Let $\mathbf{D} \subseteq \mathbf{H}(\mathcal{V})$ be any closed, convex set that contains the origin and let $\mathbf{C} \subseteq \mathbf{H}(\mathcal{V} \otimes \mathcal{W})$ be given by*

$$\mathbf{C} = \{X : X \preceq Y \otimes I_{\mathcal{W}} \text{ for some } Y \in \mathbf{D}\}.$$

Then \mathbf{C} is also a closed, convex set that contains the origin and

$$\mathbf{C}^\circ = \{Q : Q \succeq 0 \text{ and } \text{Tr}_{\mathcal{W}}(Q) \in \mathbf{D}^\circ\}.$$

Proof. That \mathbf{C} is closed, convex, and contains the origin follows immediately from its definition. Let us compute \mathbf{C}° . The assumption $0 \in \mathbf{D}$ implies that $-R \in \mathbf{C}$ for every positive semidefinite R , and hence every element of the polar \mathbf{C}° must be positive semidefinite. If it is the case that $Q \in \mathbf{C}^\circ$ then for all $Y \in \mathbf{D}$ we have

$$1 \geq \langle Q, Y \otimes I_{\mathcal{W}} \rangle = \langle \text{Tr}_{\mathcal{W}}(Q), Y \rangle$$

and so $\text{Tr}_{\mathcal{W}}(Q) \in \mathbf{D}^\circ$ as desired.

On the other hand, if $\text{Tr}_{\mathcal{W}}(Q) \in \mathbf{D}^\circ$ then

$$1 \geq \langle \text{Tr}_{\mathcal{W}}(Q), Y \rangle = \langle Q, Y \otimes I_{\mathcal{W}} \rangle$$

for all $Y \in \mathbf{D}$. If in addition Q is positive semidefinite then it also holds that $\langle Q, X \rangle \leq 1$ for all $X \preceq Y \otimes I_{\mathcal{W}}$ and therefore $Q \in \mathbf{C}^\circ$. \square

We are now ready to prove the desired characterization of polar sets of strategies.

Proof of Proposition 3.8 (Polar sets of strategies). The proof is by induction on the number of rounds r . For the base case $r = 0$, we have $\downarrow \mathbf{S}_0 = \downarrow \text{co-}\mathbf{S}_0 = [0, 1]$ and $(\downarrow \mathbf{S}_0)^\circ = (\downarrow \text{co-}\mathbf{S}_0)^\circ = (-\infty, 1]$ and so the lemma holds when $r = 0$.

For the general case, we note that the two items in the statement of the proposition are equivalent by Lemma 3.9, so it suffices to prove only the first. Let

$$\begin{aligned}\mathbf{A} &= \{X : X \succeq 0 \text{ and } \text{Tr}_{\mathcal{Y}_r}(X) \in \mathbf{B}\}, \\ \mathbf{B} &= \{Y : Y \preceq R \otimes I_{\mathcal{X}_r} \text{ for some } R \in \downarrow \mathbf{S}_{r-1}\}.\end{aligned}$$

Then by Theorem 3.2 (Characterization of strategies) we have

$$\downarrow \mathbf{S}_r = \mathbf{A}.$$

We apply Lemma 3.10 twice—once with $(\mathbf{C}^\circ, \mathbf{D}^\circ, \mathcal{W}) = (\mathbf{A}, \mathbf{B}, \mathcal{Y}_r)$ and once with $(\mathbf{C}, \mathbf{D}, \mathcal{W}) = (\mathbf{B}, \downarrow \mathbf{S}_{r-1}, \mathcal{X}_r)$ —to obtain

$$\begin{aligned}(\downarrow \mathbf{S}_r)^\circ &= \mathbf{A}^\circ = \{X : X \preceq Q \otimes I_{\mathcal{Y}_r} \text{ for some } Q \in \mathbf{B}^\circ\}, \\ \mathbf{B}^\circ &= \{Y : Y \succeq 0 \text{ and } \text{Tr}_{\mathcal{X}_r}(Y) \in (\downarrow \mathbf{S}_{r-1})^\circ\}.\end{aligned}$$

By the induction hypothesis we have

$$(\downarrow \mathbf{S}_{r-1})^\circ = \{X : X \preceq Q \text{ for some } Q \in \text{co-}\mathbf{S}_{r-1}\}.$$

Substituting this expression into the above expression for $(\downarrow \mathbf{S}_r)^\circ$ we find that the proposition follows from Theorem 3.2 (Characterization of strategies). \square

First proof of Theorem 3.3 (Maximum output probability)

First proof of Theorem 3.3. We prove the theorem for strategies—the result for co-strategies then follows from Lemma 3.6. Let p denote the maximum probability with which $\{Q_m\}$ can be forced to output m in an interaction with some compatible co-strategy. It follows from Theorem 3.1 (Interaction output probabilities) that $p = s(Q_m \mid \text{co-}\mathbf{S}_r)$. Using Proposition 3.8 (Polar sets of strategies), along with the fact that Q_m is positive semidefinite, we have

$$s(Q_m \mid \text{co-}\mathbf{S}_r) = s(Q_m \mid \downarrow \text{co-}\mathbf{S}_r) = g(Q_m \mid (\downarrow \text{co-}\mathbf{S}_r)^\circ) = g(Q_m \mid \downarrow \mathbf{S}_r),$$

which completes the proof. \square

3.3.2 Proof by semidefinite optimization duality

Our second proof of Theorem 3.3 (Maximum output probability) employs the powerful machinery of semidefinite optimization duality. The idea is to construct a semidefinite optimization problem that captures the maximum output probability, compute its dual problem, and then show that the primal and dual problems satisfy the conditions for so-called “strong” duality.

Overview of semidefinite optimization

The semidefinite optimization problems we consider are expressed in *super-operator form*. While the super-operator form differs superficially from the more conventional *standard form* for these problems, the two forms can be shown to be equivalent and the super-operator form is more convenient for our purpose. Watrous provides a helpful overview of this form of semidefinite optimization [Wat09]. For completeness, that overview is reproduced here.

Let \mathcal{X}, \mathcal{Y} be complex Euclidean spaces. A *semidefinite optimization problem* for these spaces is specified by a triple (Φ, A, B) where $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is a Hermitian-preserving super-operator and $A \in \mathbf{H}(\mathcal{X})$ and $B \in \mathbf{H}(\mathcal{Y})$. This triple specifies two optimization problems:

Primal problem	Dual problem
maximize $\langle A, X \rangle$	minimize $\langle B, Y \rangle$
subject to $\Phi(X) \preceq B$	subject to $\Phi^*(Y) \succeq A$
$X \in \mathbf{H}^+(\mathcal{X})$	$Y \in \mathbf{H}^+(\mathcal{Y})$

An operator X obeying the constraints of the primal problem is said to be *primal feasible*, while an operator Y obeying the constraints of the dual problem is called *dual feasible*. The functions $X \mapsto \langle A, X \rangle$ and $Y \mapsto \langle B, Y \rangle$ are called the primal and dual *objective functions*, respectively. We let

$$\alpha \stackrel{\text{def}}{=} \sup \{ \langle A, X \rangle : X \text{ is primal feasible} \}$$

$$\beta \stackrel{\text{def}}{=} \inf \{ \langle B, Y \rangle : Y \text{ is dual feasible} \}$$

denote the *optimal values* of the primal and dual problems. (If there are no primal or dual feasible operators then we adopt the convention $\alpha = -\infty$ and $\beta = \infty$, respectively.) It is not always the case that the optimal value is actually attained by a feasible operator for these problems.

Semidefinite optimization problems derive great utility from the notions of *weak* and *strong duality*. Essentially, weak duality asserts that $\alpha \leq \beta$ for all triples (Φ, A, B) , whereas strong duality provides conditions on (Φ, A, B) under which $\alpha = \beta$. Two such conditions are stated explicitly as follows.

Fact 3.11 (Strong duality conditions). *Let (Φ, A, B) be a semidefinite optimization problem. The following hold:*

1. (*Strict primal feasibility.*) *Suppose β is finite and there exists a primal feasible operator X such that X is positive definite and $\Phi(X) \prec B$. Then $\alpha = \beta$ and β is achieved by some dual feasible operator.*
2. (*Strict dual feasibility.*) *Suppose α is finite and there exists a dual feasible operator Y such that Y is positive definite and $\Phi^*(Y) \succ A$. Then $\alpha = \beta$ and α is achieved by some primal feasible operator.*

A semidefinite optimization problem for maximum output probabilities

We now construct a triple (Φ, A, B) whose primal problem captures the maximum probability with which a given measuring co-strategy $\{R_b\}$ can be forced to output a given outcome b by some compatible strategy. The dual problem will capture the minimization stated in Theorem 3.3. Equality between these two values will follow when we establish strong duality for (Φ, A, B) .

It is convenient to express the components of our triple (Φ, A, B) in block form. The operators A, B are given by

$$A = \begin{pmatrix} 0 & & 0 \\ & \ddots & \\ 0 & & 0 \\ & & & R_b \end{pmatrix} \quad B = \begin{pmatrix} I_{\mathcal{X}_1} & & 0 \\ & 0 & \\ & & \ddots & \\ 0 & & & 0 \end{pmatrix}$$

and the super-operator Φ is given by

$$\Phi : \begin{pmatrix} S_1 & & \\ & \ddots & \\ & & S_r \end{pmatrix} \mapsto \begin{pmatrix} \text{Tr}_{\mathcal{Y}_1}(S_1) & & & 0 \\ & \text{Tr}_{\mathcal{Y}_2}(S_2) - S_1 \otimes I_{\mathcal{X}_2} & & \\ & & \ddots & \\ 0 & & & \text{Tr}_{\mathcal{Y}_r}(S_r) - S_{r-1} \otimes I_{\mathcal{X}_r} \end{pmatrix}.$$

To be clear, Φ depends only upon the diagonal blocks S_1, \dots, S_r of the input matrix.

Correctness of the primal problem

Let us verify that the primal problem expresses the desired maximum output probability for the measurement outcome b . For any operator S with diagonal blocks S_1, \dots, S_r , the primal objective value at S is given by $\langle A, S \rangle = \langle R_b, S_r \rangle$. Hence, the desired quantity is precisely the supremum of $\langle A, S \rangle$ over all S for which S_r is a strategy. It remains only to verify that the constraint $\Phi(S) \preceq B$ enforces this property of S_r . The following lemma serves that purpose.

Lemma 3.12 (Correctness of the primal problem). *For every primal feasible solution S there exists another primal feasible solution S' whose objective value meets or exceeds that of S and whose diagonal blocks S'_1, \dots, S'_r have the property that S'_i is an i -round non-measuring strategy for each $i = 1, \dots, r$.*

Proof. Let S_1, \dots, S_r denote the diagonal blocks of S . As S is primal feasible, we

know that

$$\begin{aligned}\mathrm{Tr}_{\mathcal{Y}_1}(S_1) &\preceq I_{\mathcal{X}_1} \\ \mathrm{Tr}_{\mathcal{Y}_2}(S_2) &\preceq S_1 \otimes I_{\mathcal{X}_2} \\ &\vdots \\ \mathrm{Tr}_{\mathcal{Y}_r}(S_r) &\preceq S_{r-1} \otimes I_{\mathcal{X}_r}\end{aligned}$$

It is clear that there is a $S'_1 \succeq S_1$ such that $\mathrm{Tr}_{\mathcal{Y}_1}(S'_1) = I_{\mathcal{X}_1}$. Similarly, for each $i = 2, \dots, r$ there is a $S'_i \succeq S_i$ such that $\mathrm{Tr}_{\mathcal{Y}_i}(S'_i) = S'_{i-1} \otimes I_{\mathcal{X}_i}$. The desired operator S' is then obtained from S by replacing its diagonal blocks with S'_1, \dots, S'_r . That $\langle A, S' \rangle \geq \langle A, S \rangle$ follows from the fact that $A \succeq 0$ and $S' \succeq S$. That S'_1, \dots, S'_r are non-measuring strategies follows from Theorem 3.2 (Characterization of strategies). \square

Correctness of the dual problem

Before we can show correctness of the dual problem, we must compute the adjoint super-operator Φ^* . Let us verify that $\Phi^* = \Psi$ where Ψ is given by

$$\begin{aligned}\Psi &: \begin{pmatrix} T_1 & & \\ & \ddots & \\ & & T_r \end{pmatrix} \\ &\mapsto \begin{pmatrix} T_1 \otimes I_{\mathcal{Y}_1} - \mathrm{Tr}_{\mathcal{X}_2}(T_2) & & & 0 \\ & \ddots & & \\ & & T_{r-1} \otimes I_{\mathcal{Y}_{r-1}} - \mathrm{Tr}_{\mathcal{X}_r}(T_r) & \\ 0 & & & T_r \otimes I_{\mathcal{Y}_r} \end{pmatrix}.\end{aligned}$$

Let S, T be any operators in the domains of Φ, Ψ , respectively, and let S_1, \dots, S_r and T_1, \dots, T_r denote the diagonal blocks of S and T , respectively, so that

$$\begin{aligned}\langle \Phi(S), T \rangle &= \langle \mathrm{Tr}_{\mathcal{Y}_1}(S_1), T_1 \rangle + \sum_{i=2}^r \langle \mathrm{Tr}_{\mathcal{Y}_i}(S_i) - S_{i-1} \otimes I_{\mathcal{X}_i}, T_i \rangle \\ &= \langle S_1, T_1 \otimes I_{\mathcal{Y}_1} \rangle + \sum_{i=2}^r \langle S_i, T_i \otimes I_{\mathcal{Y}_i} \rangle - \langle S_{i-1}, \mathrm{Tr}_{\mathcal{X}_i}(T_i) \rangle \\ &= \langle S_r, T_r \otimes I_{\mathcal{Y}_r} \rangle + \sum_{i=1}^{r-1} \langle S_i, T_i \otimes I_{\mathcal{Y}_i} \rangle - \langle S_i, \mathrm{Tr}_{\mathcal{X}_{i+1}}(T_{i+1}) \rangle \\ &= \langle S_r, T_r \otimes I_{\mathcal{Y}_r} \rangle + \sum_{i=1}^{r-1} \langle S_i, T_i \otimes I_{\mathcal{Y}_i} - \mathrm{Tr}_{\mathcal{X}_{i+1}}(T_{i+1}) \rangle \\ &= \langle S, \Psi(T) \rangle.\end{aligned}$$

As this equality holds for all S, T , it follows from the definition of the adjoint that $\Phi^* = \Psi$ as claimed.

We now prove a lemma that establishes the link between the dual problem and the minimization condition of Theorem 3.3 (Maximum output probability).

Lemma 3.13 (Correctness of the dual problem). *For every dual feasible solution T there exists another dual feasible solution T' whose objective value p equals that of T and whose diagonal blocks T'_1, \dots, T'_r have the property that $T'_i \otimes I_{\mathcal{Y}_i}$ is a non-measuring co-strategy multiplied by p for each $i = 1, \dots, r$.*

Proof. Let T_1, \dots, T_r denote the diagonal blocks of T . As T is feasible, we know that

$$\begin{aligned} T_1 \otimes I_{\mathcal{Y}_1} &\succeq \text{Tr}_{\mathcal{X}_2}(T_2) \\ &\vdots \\ T_{r-1} \otimes I_{\mathcal{Y}_{r-1}} &\succeq \text{Tr}_{\mathcal{X}_r}(T_r) \\ T_r \otimes I_{\mathcal{Y}_r} &\succeq R_b \end{aligned}$$

The objective value for T is then given by $p = \langle B, T \rangle = \text{Tr}(T_1)$. If $p = 0$ then it must hold that $T_1 = \dots = T_r = 0$, so the lemma holds trivially in this case. For the remainder of the proof we shall assume $p > 0$.

It is clear that there is a $T'_2 \succeq T_2$ such that $T_1 \otimes I_{\mathcal{Y}_1} = \text{Tr}_{\mathcal{X}_2}(T'_2)$. Similarly, for each $i = 3, \dots, r$ there is a $T'_i \succeq T_i$ such that $T'_{i-1} \otimes I_{\mathcal{Y}_{i-1}} = \text{Tr}_{\mathcal{X}_i}(T'_i)$. As $T'_r \succeq T_r$, it must also be the case that $T'_r \otimes I_{\mathcal{Y}_r} \succeq R_b$. The desired operator T' is then obtained from T by replacing its diagonal blocks with T_1, T'_2, \dots, T'_r . That $\langle B, T' \rangle = \langle B, T \rangle$ follows from the equality $\langle B, T' \rangle = \text{Tr}(T_1) = p$. Finally, it follows from Theorem 3.2 (Characterization of strategies) that each $\frac{1}{p}T'_i \otimes I_{\mathcal{Y}_i}$ is a non-measuring co-strategy. \square

As any feasible solution T has $\Phi^*(T) \succeq A$, it must be that the r th diagonal block T_r of T has the property that $T_r \otimes I_{\mathcal{Y}_r} \succeq R_b$. By Lemma 3.13 we may also assume that $T_r \otimes I_{\mathcal{Y}_r}$ is a scalar multiple of a co-strategy. It is clear then that the optimal value of the dual problem equals the infimum over all such scalar multiples, subject to $T_r \otimes I_{\mathcal{Y}_r} \succeq R_b$, as required by Theorem 3.3.

Strong duality of the semidefinite optimization problem

We already argued that the two quantities appearing in Theorem 3.3 are captured by the primal and dual semidefinite optimization problems associated with the triple (Φ, A, B) . To prove Theorem 3.3, it remains only to show that these two quantities are equal. This equality is established by showing that (Φ, A, B) satisfies the conditions for strong duality.

Lemma 3.14 (Strong duality of (Φ, A, B)). *There exists a primal feasible operator S and a dual feasible operator T such that $\langle A, S \rangle = \langle B, T \rangle$.*

Proof. We prove strong duality via item 1 of Fact 3.11 (Strong duality conditions). That is, we show that β is finite and the primal problem is strictly feasible. Then by Fact 3.11 it follows that $\alpha = \beta$ and that β is achieved for some dual feasible operator. We complete the proof by noting that the optimal value α is also achieved by a primal feasible operator.

First, let us argue that β is finite. As $B \succeq 0$, any dual feasible solution has nonnegative objective value. Thus, to show that β is finite it suffices to exhibit a single dual feasible solution. Toward that end, let R be a non-measuring co-strategy with $R \succeq R_b$. Let T_r be such that $T_r \otimes I_{\mathcal{Y}_r} = R$ and for each $i = r-1, \dots, 1$ choose T_i so that $T_i \otimes I_{\mathcal{Y}_i} = \text{Tr}_{\mathcal{X}_{i+1}}(T_{i+1})$. (That each of T_1, \dots, T_r exists follows from Theorem 3.2 (Characterization of strategies).) Finally, let T be the block-diagonal operator whose diagonal blocks are T_1, \dots, T_r . It is clear that T is dual feasible.

Next, we show that the primal is strictly feasible. Choose $\delta \in (0, 1/r)$ and let S be the block-diagonal operator whose i th diagonal block S_i is given by

$$S_i = \frac{1 - i\delta}{\dim(\mathcal{Y}_{1\dots i})} I_{\mathcal{Y}_{1\dots i} \otimes \mathcal{X}_{1\dots i}}.$$

It is clear that $S \succ 0$ and it is tedious but straightforward to verify that $\Phi(S) \succ B$. In particular, we have

$$\begin{aligned} \text{Tr}_{\mathcal{Y}_1}(S_1) &= (1 - \delta) I_{\mathcal{X}_1} \prec I_{\mathcal{X}_1} \\ \text{Tr}_{\mathcal{Y}_2}(S_2) &= \frac{1 - 2\delta}{\dim(\mathcal{Y}_1)} I_{\mathcal{Y}_1 \otimes \mathcal{X}_{1\dots 2}} \prec \frac{1 - \delta}{\dim(\mathcal{Y}_1)} I_{\mathcal{Y}_1 \otimes \mathcal{X}_{1\dots 2}} = S_1 \otimes I_{\mathcal{X}_2} \\ &\quad \vdots \\ \text{Tr}_{\mathcal{Y}_r}(S_r) &= \frac{1 - r\delta}{\dim(\mathcal{Y}_{1\dots r-1})} I_{\mathcal{Y}_{1\dots r-1} \otimes \mathcal{X}_{1\dots r}} \prec \frac{1 - (r-1)\delta}{\dim(\mathcal{Y}_{1\dots r-1})} I_{\mathcal{Y}_{1\dots r-1} \otimes \mathcal{X}_{1\dots r}} = S_{r-1} \otimes I_{\mathcal{X}_r} \end{aligned}$$

It now follows from item 1 of Fact 3.11 that $\alpha = \beta$ and that β is achieved by some dual feasible operator.

It remains only to show that α is also achieved by some primal feasible operator. By Lemma 3.12 it suffices to consider only those primal feasible S whose diagonal blocks are strategies. As the set of strategies is compact (Proposition 3.15), it follows that the optimal α is finite and is achieved by a primal feasible solution. \square

Second proof of Theorem 3.3 (Maximum output probability)

Second proof of Theorem 3.3. We prove the theorem for co-strategies—the result for strategies then follows from Lemma 3.6. Let $p_m \in [0, 1]$ denote the maximum probability with which $\{Q_m\}$ can be forced to output m in an interaction with some compatible strategy. By Lemma 3.12, the optimal value of the primal problem associated with (Φ, A, B) equals p_m . By Lemma 3.13, the optimal value of the dual problem associated with (Φ, A, B) equals the minimum p such that $Q_m \preceq pR$ for some non-measuring co-strategy R . That these two values are equal follows from Lemma 3.14. \square

3.4 Other properties of strategies

We conclude this chapter with three simple applications of the properties of strategies established in Theorems 3.1 and 3.2. In particular, we point out that Theorem 3.2 immediately implies that the set of all strategies is compact and convex. We also show that representations of strategies are unique and that they satisfy a convenient distributive property. The results in this section are proven only for strategies, but it is trivial to repeat each proof for co-strategies.

The set of strategies is compact and convex

Proposition 3.15 (Convexity of strategies). *The set of all r -round strategies is compact and convex, as is the set of all r -round co-strategies.*

Proof. That the set of strategies is bounded follows from Proposition 1.3 and the fact that every strategy is the Choi-Jamiołkowski representation of some quantum operation.

That the set of strategies is closed and convex follows from Theorem 3.2, which characterizes this set as an intersection between two closed and convex sets—the positive semidefinite operators and those operators satisfying the linear constraints appearing in the theorem. \square

Equivalence and uniqueness of strategies

In Section 2.1.2 we noted that operational descriptions of strategies are not unique in the sense that two distinct descriptions could specify equivalent strategies. For the new representation of strategies, this is not so.

Proposition 3.16 (Uniqueness of strategies). *Two strategies Q, Q' may be distinguished with nonzero bias by a compatible measuring co-strategy if and only if $Q \neq Q'$. A similar statement holds for co-strategies.*

Proof. Let $\{R_b\}$ be any measuring co-strategy. By Theorem 3.1, the probability with which $\{R_b\}$ outputs a given outcome b after an interaction with Q or with Q' is given by $\langle R_b, Q \rangle$ or $\langle R_b, Q' \rangle$, respectively. If $Q = Q'$ then it holds that $\langle X, Q \rangle = \langle X, Q' \rangle$ for every operator X and so $\{R_b\}$ cannot distinguish Q from Q' with nonzero bias.

Conversely, if $Q \neq Q'$ then for each spanning set \mathbf{B} of operators there must exist some $B \in \mathbf{B}$ with $\langle B, Q \rangle \neq \langle B, Q' \rangle$. We claim that the set

$$\{R_b : R_b \text{ is an element of some measuring co-strategy}\}$$

is a spanning set, from which the proposition follows.

One way to verify this claim is to note that the identity operator I (suitably normalized) denotes a valid non-measuring co-strategy. It follows from Proposition 2.7 that for each positive semidefinite operator $R \preceq I$ (again, suitably normalized) there exists a measuring co-strategy of which R is an element. The claim then follows from the fact that the set of all operators $0 \preceq R \preceq I$ is a spanning set. \square

Distributive property for probabilistic combinations of strategies

In Example 2.3 it was shown that the operational description of strategies lacks a convenient distributive property for probabilistic combinations of strategies. This problem is rectified under the new representation for strategies.

Proposition 3.17 (Distributive property for strategies). *Let Q, Q' be strategies and let $p \in [0, 1]$. The strategy that plays according to Q with probability p and according to Q' otherwise is given by the convex combination $pQ + (1 - p)Q'$. A similar statement holds for co-strategies.*

Proof. Let Q'' denote the probabilistic combination of Q and Q' described in the statement of the proposition. For any measuring co-strategy $\{R_b\}$ and any outcome b we let r_b denote the probability with which the interaction between Q and $\{R_b\}$ yields the outcome b . We also define r'_b and r''_b similarly in terms of Q' and Q'' , respectively. By definition, $r''_b = pr_b + (1 - p)r'_b$.

By Theorem 3.1 and the linearity of the inner product, we have

$$\langle Q'', R_b \rangle = r''_b = pr_b + (1 - p)r'_b = p \langle Q, R_b \rangle + (1 - p) \langle Q', R_b \rangle = \langle pQ + (1 - p)Q', R_b \rangle$$

from which it follows that $\{R_b\}$ cannot distinguish Q'' from $pQ + (1 - p)Q'$ with nonzero bias. That these two operators are equal then follows from Proposition 3.16. \square

Chapter 4

Applications

In this chapter we present several applications of the formalism for quantum strategies developed in Chapter 3:

- In Section 4.1 we develop a general formalism for zero-sum quantum games. A quantum analogue of von Neumann's Min-Max Theorem for two-player zero-sum games is established, and an efficient algorithm to compute the value of a zero-sum quantum game is presented.
- Section 4.2 contains complexity theoretic applications of quantum strategies, including the collapse of the complexity classes QRG and EXP and parallel repetition results for single-prover quantum interactive proofs and for quantum interactive proofs with two competing provers.
- In Section 4.3 we provide a simplified proof of Kitaev's bound for strong quantum coin-flipping.

With the exception of the parallel repetition results, each of these applications first appeared in Ref. [GW07]. The parallel repetition result for single-prover quantum interactive proofs is due to Watrous and was not published prior to the present thesis.

4.1 Theory of zero-sum quantum games

In this section we develop a general formalism for zero-sum quantum games. Specifically, we prove a quantum analogue of von Neumann's Min-Max Theorem for two-player zero-sum games in Section 4.1.1 and we provide an efficient algorithm to compute the value of such a game in Section 4.1.2.

4.1.1 Min-max theorem for zero-sum quantum games

This section begins with a brief introduction to games and min-max theorems, after which we provide our formalism for quantum games and the quantum min-max theorem.

While it is natural to expect a min-max theorem to hold for two-player zero-sum quantum games, the absence of a mathematically convenient formalism for these games has precluded the appearance of such a theorem in the literature. (Lee and Johnson established a min-max theorem for the special case of one-round quantum games, wherein only a single round of messages is exchanged between the players and a referee [LJ03].)

Two-player games and min-max theorems

A *game* consists of an interaction between two or more players, followed by a *payout* that is awarded to each player at the end of the interaction. It is the goal of each player to maximize his or her own payout. A two-player game is *zero-sum* if the sum of the payouts awarded to the players is always zero. Two-player zero-sum games are always competitive, as the players never have incentive to co-operate. Many popular examples of games fall into this category, including Poker, Checkers, and Go. (Win-lose games such as Checkers and Go can be represented as a zero-sum game wherein the only possible payouts are ± 1 .)

Let us call the two players in a zero-sum game Alice and Bob. In the context of these games, a *min-max theorem* is an assertion that every game has a *value* v with the following properties:

- (i) There exists a strategy for Alice that ensures a payout of at least v regardless of Bob's strategy.
- (ii) There exists a strategy for Bob that ensures a payout of at most v to Alice regardless of her strategy.

In other words, there always exist strategies for the players that are *optimal* in the sense that the players never have incentive to deviate from their optimal strategies. The original Min-Max Theorem for classical games was established by von Neumann in 1928 [vN28].

In an analytical context, min-max theorems are statements about sets A , B and functions $f : A \times B \rightarrow \mathbb{R}$. While it must always hold that

$$\sup_{a \in A} \inf_{b \in B} f(a, b) \leq \inf_{b \in B} \sup_{a \in A} f(a, b),$$

a *min-max theorem* provides conditions upon A , B , f under which these two quantities are equal. For example, the following well-known min-max theorem will prove useful for our purpose.

Fact 4.1 (Convex-bilinear min-max theorem). *If A, B are compact convex subsets of finite-dimensional real vector spaces and $f : A \times B \rightarrow \mathbb{R}$ is bilinear then*

$$\max_{a \in A} \min_{b \in B} f(a, b) = \min_{b \in B} \max_{a \in A} f(a, b).$$

While Fact 4.1 does not follow immediately from von Neumann’s original Min-Max Theorem, it does follow from an early generalization due to Ville [Vil38]. Additional citations and an English-language proof can be found in Rockafellar [Roc70].

Analytical min-max theorems such as Fact 4.1 can be used to establish a min-max theorem for a given game or class of games. To do so, it suffices that A, B represent the sets of all possible strategies for Alice and Bob and that $f(a, b)$ denotes the payout to Alice when she employs strategy a and Bob employs strategy b .

Formal definition of a zero-sum quantum game

Classical two-player games have several distinct mathematical formalizations, each with its own advantages. One of these formalisms—the *refereed game*—lends itself particularly well to generalization to quantum games. This formalism encapsulates the rules of a particular game into a *referee*, who exchanges messages with each of the two players. The referee enforces the rules of the game and decides when to terminate the interaction and award a payout to the players. In a *quantum* refereed game the players and referee may exchange and process quantum information and the payout is determined by a measurement made by the referee at the end of the interaction.

Definition 4.2 (Quantum game). *An r -round referee is an r -round measuring co-strategy $\{R_m\}_{m \in \Sigma}$ whose input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ take the form*

$$\mathcal{X}_i = \mathcal{A}_i \otimes \mathcal{B}_i \quad \text{and} \quad \mathcal{Y}_i = \mathcal{C}_i \otimes \mathcal{D}_i$$

for complex Euclidean spaces $\mathcal{A}_i, \mathcal{B}_i, \mathcal{C}_i,$ and \mathcal{D}_i for $1 \leq i \leq r$. An *r -round quantum game* consists of an r -round referee along with *payout functions*

$$V_A, V_B : \Sigma \rightarrow \mathbb{R}$$

defined on the referee’s set Σ of measurement outcomes. For each such outcome $m \in \Sigma$, Alice’s payout is $V_A(m)$ and Bob’s payout is $V_B(m)$. Such a game is *zero-sum* if $V_A(m) + V_B(m) = 0$ for all $m \in \Sigma$.

During each round, the referee simultaneously sends a message to Alice and a message to Bob, and a response is expected from each player. The spaces \mathcal{A}_i and \mathcal{B}_i correspond to the messages sent by the referee during the i th round, while \mathcal{C}_i and \mathcal{D}_i correspond to their responses. After r rounds, the referee produces an output $m \in \Sigma$ and awards the payouts $V_A(m)$ to Alice and $V_B(m)$ to Bob. \square

Notice that Definition 4.2 places no restrictions on the strategies available to the players. For example, the players might employ a strategy that allows them to exchange quantum information directly, as opposed to an indirect exchange via messages to the referee. Alternatively, they might share entanglement or randomness but be forbidden from direct communication, or they might even be forbidden from sharing entanglement or randomness altogether.

Min-max theorem for zero-sum quantum games

Let

$$\mathbf{A} \subset \mathbf{H}^+(\mathcal{C}_{1\dots r} \otimes \mathcal{A}_{1\dots r})$$

denote the set of all r -round non-measuring strategies for Alice's input spaces $\mathcal{A}_1, \dots, \mathcal{A}_r$ and output spaces $\mathcal{C}_1, \dots, \mathcal{C}_r$. Similarly, let

$$\mathbf{B} \subset \mathbf{H}^+(\mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r})$$

denote the set of all r -round non-measuring strategies for Bob's input spaces $\mathcal{B}_1, \dots, \mathcal{B}_r$ and output spaces $\mathcal{D}_1, \dots, \mathcal{D}_r$.

In a zero-sum quantum game it cannot simultaneously be to both players' advantage to communicate directly with each other or to share a source of randomness or entanglement. Thus, we may assume that Alice and Bob play independent strategies represented by $A \in \mathbf{A}$ and $B \in \mathbf{B}$, respectively. In particular, their combined r -round strategy for the referee's input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ is described by the tensor product $A \otimes B$.

For any zero-sum quantum game with referee $\{R_m\}$ and payout functions V_A, V_B we write

$$V(m) \stackrel{\text{def}}{=} V_A(m) = -V_B(m)$$

and define the Hermitian operator

$$R \stackrel{\text{def}}{=} \sum_{m \in \Sigma} V(m) R_m.$$

By Theorem 3.1 (Interaction output probabilities) Alice's expected payout for this game is given by

$$\sum_{m \in \Sigma} V(m) \langle A \otimes B, R_m \rangle = \langle A \otimes B, R \rangle$$

while Bob's expected payout is $-\langle A \otimes B, R \rangle$. Because the inner product $\langle A \otimes B, R \rangle$ is a bilinear function of A and B and because the sets \mathbf{A}, \mathbf{B} are compact and convex (as noted in Proposition 3.15), we may employ Fact 4.1 (Convex-bilinear min-max theorem) to obtain

$$\max_{A \in \mathbf{A}} \min_{B \in \mathbf{B}} \langle A \otimes B, R \rangle = \min_{B \in \mathbf{B}} \max_{A \in \mathbf{A}} \langle A \otimes B, R \rangle.$$

The real number represented by the two sides of this equation is called the *value* of the game. Any strategy $A \in \mathbf{A}$ achieving the maximum of the left side of this equality is an optimal strategy for Alice, while any strategy $B \in \mathbf{B}$ achieving the minimum of the right side of this equality is an optimal strategy for Bob. The following theorem is now proved.

Theorem 4.3 (Min-max theorem for zero-sum quantum games). *Every two-player zero-sum r -round quantum game has a value v with the following properties:*

- (i) *There exists a strategy for Alice that ensures a payout of at least v regardless of Bob's strategy.*
- (ii) *There exists a strategy for Bob that ensures a payout of at most v to Alice regardless of her strategy.*

4.1.2 Efficient algorithm to compute the value

Theorem 4.3 (Min-max theorem for zero-sum quantum games) asserts that each zero-sum quantum game has a value. But can this value be easily computed? In this section we answer this question in the affirmative by exhibiting an efficient deterministic algorithm that takes as input a description of a referee $\{R_m\}$ and produces as output the value of the zero-sum quantum game defined by $\{R_m\}$.

This goal is achieved by expressing the value of the game as a semidefinite optimization problem and then employing the existence of efficient algorithms for semidefinite optimization. As these algorithms apply only to certain semidefinite optimization problems whose feasible sets are “well-bounded,” it is necessary to establish the well-boundedness of our semidefinite optimization problem for zero-sum quantum games.

An optimization problem for the value of a game

Let $\{R_m\}_{m \in \Sigma}$ be an r -round referee and consider the following super-operator defined for each measurement outcome $m \in \Sigma$:

$$\begin{aligned} \Omega_m &: \mathbf{L}(\mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{C}_{1\dots r} \otimes \mathcal{A}_{1\dots r}) \\ &: B \mapsto \text{Tr}_{\mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r}} ((B \otimes I_{\mathcal{C}_{1\dots r} \otimes \mathcal{A}_{1\dots r}}) R_m). \end{aligned}$$

These super-operators have the property that $\langle A \otimes B, R_m \rangle = \langle A, \Omega_m(B) \rangle$ for each measurement outcome $m \in \Sigma$ and each choice of strategies A for Alice and B for Bob. Indeed, the set $\{\Omega_m(B)\}_{m \in \Sigma}$ is the r -round measuring co-strategy for Alice's input spaces $\mathcal{A}_1, \dots, \mathcal{A}_r$ and output spaces $\mathcal{C}_1, \dots, \mathcal{C}_r$ obtained by “hard-wiring” Bob's strategy B into the referee.

Let $V(m)$ denote the payout to Alice in a zero-sum quantum game with referee $\{R_m\}$. For the moment, it is convenient to restrict our attention to payout functions

with $0 \leq V(m) \leq 1$ for all m . But we shall soon see that the ensuing discussion is easily generalized to arbitrary payout functions.

Borrowing from the previous subsection, we define the super-operator

$$\Omega_R \stackrel{\text{def}}{=} \sum_{m \in \Sigma} V(m) \Omega_m$$

so that Alice's expected payout is given by $\langle A, \Omega_R(B) \rangle$ when Alice and Bob play according to the strategies A and B , respectively. Naturally, Alice's maximum expected payout when Bob plays according to B is

$$\max_{A \in \mathbf{A}} \langle A, \Omega_R(B) \rangle.$$

Let $\text{co-}\mathbf{A}$ denote the set of co-strategies for Alice's input and output spaces. Because $0 \leq V(m) \leq 1$ for all m , it follows that, regardless of the choice of $B \in \mathbf{B}$, the operator $\Omega_R(B)$ is always an element of some measuring co-strategy. Hence, we may apply Theorem 3.3 (Maximum output probability) to obtain the following alternate expression for Alice's maximum expected payout:

$$\min \{ \lambda \geq 0 : \Omega_R(B) \preceq \lambda Q \text{ for some } Q \in \text{co-}\mathbf{A} \}.$$

Thus, the value of the game is given by

$$\min_{B \in \mathbf{B}} \max_{A \in \mathbf{A}} \langle A, \Omega_R(B) \rangle = \min_{B \in \mathbf{B}} \min \{ \lambda \geq 0 : \Omega_R(B) \preceq \lambda Q \text{ for some } Q \in \text{co-}\mathbf{A} \},$$

which can equivalently be written as an optimization problem:

$$\begin{aligned} & \text{minimize} && \lambda \\ & \text{subject to} && \Omega_R(B) \preceq \lambda Q \\ & && B \in \mathbf{B} \\ & && Q \in \text{co-}\mathbf{A}. \end{aligned} \tag{4.1}$$

Thus, an efficient solution for this optimization problem yields an efficient algorithm to compute the value of a zero-sum quantum game, provided that the payouts all fall within the interval $[0, 1]$. To compute the value of a game with an *arbitrary* payout function V we simply shift and scale V to a payout function V' with payouts in $[0, 1]$ by defining

$$V'(m) \stackrel{\text{def}}{=} \frac{V(m) + |\min_n \{V(n)\}|}{|\max_n \{V(n)\}| + |\min_n \{V(n)\}|}$$

and then solve the optimization problem associated with V' , from which the original value is easily recovered.

A semidefinite optimization problem for the value of a game

Let us argue that the optimization problem (4.1) can be expressed as a semidefinite optimization problem in the super-operator form described in Section 3.3.2. Specifically, we construct a triple (Φ, E, F) with the property that the optimal value of (4.1) equals the optimal value of the dual problem

$$\begin{aligned} & \text{minimize} && \langle F, X \rangle \\ & \text{subject to} && \Phi(X) \succeq E \\ & && X \succeq 0. \end{aligned}$$

To this end, we choose the super-operator Φ so that the semidefinite variable X can be assumed without loss of generality to be block-diagonal with the form

$$X = \begin{pmatrix} B & \\ & Q \end{pmatrix}, \quad B = \begin{pmatrix} B_1 & & \\ & \ddots & \\ & & B_r \end{pmatrix}, \quad Q = \begin{pmatrix} Q_1 & & \\ & \ddots & \\ & & Q_r \end{pmatrix}$$

It is convenient to decompose Φ , E , and F into a hierarchy of block-diagonal operators as follows:

$$E = \begin{pmatrix} E_{\text{Bob}} & & \\ & E_{\text{co-Alice}} & \\ & & 0 \end{pmatrix}, \quad F = \begin{pmatrix} F_{\text{Bob}} & & \\ & F_{\text{co-Alice}} & \end{pmatrix},$$

and

$$\Phi : \begin{pmatrix} B & \\ & Q \end{pmatrix} \mapsto \begin{pmatrix} \Phi_{\text{Bob}}(B) & & 0 \\ & \Phi_{\text{co-Alice}}(Q) & \\ 0 & & Q_r \otimes I_{C_r} - \Omega_R(B_r) \end{pmatrix}.$$

It is clear that the condition $\Phi(X) \succeq E$ is equivalent to

$$\begin{aligned} \Phi_{\text{Bob}}(B) &\succeq E_{\text{Bob}}, \\ \Phi_{\text{co-Alice}}(Q) &\succeq E_{\text{co-Alice}}, \\ Q_r \otimes I_{C_r} &\succeq \Omega_R(B_r). \end{aligned}$$

The idea is that we will choose Φ_{Bob} , $\Phi_{\text{co-Alice}}$ so that the first two constraints ensure that B_r is a valid r -round strategy for Bob and that $Q_r \otimes I_{C_r}$ is (a scalar multiple of) a valid r -round co-strategy for Alice's input and output spaces. The third constraint clearly captures the semidefinite inequality condition of (4.1).

Let us fill in the details for Φ , E , and F . Define

$$\Phi_{\text{Bob}} : B \mapsto \begin{pmatrix} \text{Tr}_{\mathcal{D}_1}(B_1) & & & 0 \\ & \text{Tr}_{\mathcal{D}_2}(B_2) - B_1 \otimes I_{\mathcal{B}_2} & & \\ & & \ddots & \\ 0 & & & \text{Tr}_{\mathcal{D}_r}(B_r) - B_{r-1} \otimes I_{\mathcal{B}_r} \end{pmatrix},$$

$$\Phi_{\text{co-Alice}} : Q \mapsto \begin{pmatrix} Q_1 \otimes I_{\mathcal{C}_1} - \text{Tr}_{\mathcal{A}_2}(Q_2) & & 0 \\ & \ddots & \\ 0 & & Q_{r-1} \otimes I_{\mathcal{C}_{r-1}} - \text{Tr}_{\mathcal{A}_r}(Q_r) \end{pmatrix},$$

and

$$E_{\text{Bob}} = \begin{pmatrix} I_{\mathcal{B}_1} & & 0 \\ & 0 & \\ & & \ddots \\ 0 & & & 0 \end{pmatrix}, \quad E_{\text{co-Alice}} = 0.$$

For the objective function, we select

$$F_{\text{Bob}} = 0, \quad F_{\text{co-Alice}} = \begin{pmatrix} I_{\mathcal{A}_1} & & 0 \\ & 0 & \\ & & \ddots \\ 0 & & & 0 \end{pmatrix}$$

so that

$$\langle F, X \rangle = \langle F_{\text{co-Alice}}, Q \rangle = \langle I_{\mathcal{A}_1}, Q_1 \rangle = \text{Tr}(Q_1).$$

Correctness of the semidefinite optimization problem

Let us verify that the optimal value of the dual problem (Φ, E, F) equals the optimal value of (4.1). To the extent that (Φ, E, F) resembles the semidefinite optimization problem appearing in Section 3.3.2, the material in that section—particularly the proofs of Lemmas 3.12 and 3.13—can be reused in the present setting.

Lemma 4.4 (Correctness of (Φ, E, F)). *For each dual feasible X for (Φ, E, F) there is another dual feasible X' whose objective value p equals that of X and whose diagonal blocks $B'_1, \dots, B'_r, Q'_1, \dots, Q'_r$ have the property that*

1. B'_i is an i -round non-measuring strategy, and
2. $Q'_i \otimes I_{\mathcal{C}_i}$ is an i -round non-measuring co-strategy multiplied by p

for each $i = 1, \dots, r$.

Proof. Let $B_1, \dots, B_r, Q_1, \dots, Q_r$ denote the diagonal blocks of X . First, we choose B'_1, \dots, B'_r and show that they satisfy item 1 in the statement of the lemma. The proof is nearly identical to that of Lemma 3.12. As X is feasible, we know that $\Phi_{\text{Bob}}(B) \succeq E_{\text{Bob}}$ and therefore

$$\begin{aligned} \text{Tr}_{\mathcal{D}_1}(B_1) &\succeq I_{\mathcal{B}_1} \\ \text{Tr}_{\mathcal{D}_2}(B_2) &\succeq B_1 \otimes I_{\mathcal{B}_2} \\ &\vdots \\ \text{Tr}_{\mathcal{D}_r}(B_r) &\succeq B_{r-1} \otimes I_{\mathcal{B}_r}. \end{aligned}$$

It is clear that there is a $B'_1 \preceq B_1$ such that $\text{Tr}_{\mathcal{D}_1}(B'_1) = I_{\mathcal{B}_1}$. Similarly, for each $i = 2, \dots, r$ there is a $B'_i \preceq B_i$ such that $\text{Tr}_{\mathcal{D}_i}(B'_i) = B'_{i-1} \otimes I_{\mathcal{B}_i}$. That B'_1, \dots, B'_r are non-measuring strategies follows from Theorem 3.2 (Characterization of strategies).

Next, we choose Q'_1, \dots, Q'_r and show that they satisfy item 2 in the statement of the proposition. The proof is nearly identical to that of Lemma 3.13. As X is feasible, we know that $\Phi_{\text{co-Alice}}(Q) \succeq E_{\text{co-Alice}}$ and therefore

$$\begin{aligned} Q_1 \otimes I_{\mathcal{C}_1} &\succeq \text{Tr}_{\mathcal{A}_2}(Q_2) \\ &\vdots \\ Q_{r-1} \otimes I_{\mathcal{C}_{r-1}} &\succeq \text{Tr}_{\mathcal{A}_r}(Q_r). \end{aligned}$$

As X has objective value p , it must be that $p = \text{Tr}(Q_1)$. If $p = 0$ then it must hold that $Q_1 = \dots = Q_r = 0$, so item 2 holds trivially for the choice $Q'_1 = \dots = Q'_r = 0$.

Assume then that $p > 0$. It is clear that there is a $Q'_2 \succeq Q_2$ such that $Q_1 \otimes I_{\mathcal{C}_1} = \text{Tr}_{\mathcal{A}_2}(Q'_2)$. Similarly, for each $i = 3, \dots, r$ there is a $Q'_i \succeq Q_i$ such that $Q'_{i-1} \otimes I_{\mathcal{C}_{i-1}} = \text{Tr}_{\mathcal{A}_i}(Q'_i)$. Taking $Q'_1 = Q_1$ and using the fact that $p = \text{Tr}(Q_1)$, it follows from Theorem 3.2 (Characterization of strategies) that each $\frac{1}{p}Q'_i \otimes I_{\mathcal{C}_i}$ is a non-measuring co-strategy.

At this point, we have chosen the diagonal blocks $B'_1, \dots, B'_r, Q'_1, \dots, Q'_r$ of X' . That X and X' have the same objective value follows immediately from the choice $Q'_1 = Q_1$. It remains only to verify that $Q'_r \otimes I_{\mathcal{C}_r} \succeq \Omega_R(B'_r)$. But this inequality follows immediately from the facts that $Q'_r \succeq Q_r$, that $B'_r \preceq B_r$, and that the super-operator Ω_R is completely positive. \square

Computational efficiency and well-boundedness

Now that we have expressed the value of a zero-sum quantum game as a semidefinite optimization problem, it remains only to argue that an existing algorithm for semidefinite optimization can be applied to our specific problem (Φ, E, F) .

The algorithm we employ is the ellipsoid method [Kha79, GLS88], which approximates the optimal value of a semidefinite optimization problem to arbitrary

precision in polynomial-time, provided that the set of feasible solutions is *well-bounded*—that is, it contains a ball of radius δ and is, in turn, contained in a ball of radius Δ such that the ratio Δ/δ is not too large. More precisely, we have the following.

Fact 4.5 (Efficient algorithm for semidefinite optimization). *The following promise problem admits a deterministic polynomial-time solution:*

Input. *A semidefinite optimization problem (Φ, A, B) and positive real numbers $\varepsilon, \delta, \Delta$. The numbers ε, δ , and Δ are given explicitly in binary, as are the real and complex parts of each entry of $J(\Phi)$, A , and B .*

Promise. *There exists a primal [dual] feasible solution X_0 such that for all Hermitian operators H with $\|H\|_F \leq \delta$ it holds that $X_0 + H$ is primal [dual] feasible. Moreover, for all primal [dual] feasible solutions X it holds that $\|X\|_F \leq \Delta$.*

Output. *A real number γ such that $|\gamma - \alpha| < \varepsilon$, where α is the optimal value of the primal [dual] problem associated with (Φ, A, B) .*

In addition to the ellipsoid method, interior point methods are also used for semidefinite optimization [dK02, BV04]. The reader is referred to Watrous [Wat09] and the references therein for more detailed discussion of algorithms for semidefinite optimization problems written in super-operator form.

Unfortunately, the set of dual feasible solutions associated with our problem (Φ, E, F) is unbounded—if X is dual feasible then so is λX for all $\lambda \geq 1$. To remedy this problem, it suffices to augment the original problem (Φ, E, F) with the additional constraint $\text{Tr}(X) \leq t$ on dual feasible solutions X for some appropriately large choice of t . Such a constraint can be incorporated into the super-operator form by defining

$$\begin{aligned} \Phi' : X &\mapsto \begin{pmatrix} \Phi(X) & \\ & -\text{Tr}(X) \end{pmatrix}, \\ E' &= \begin{pmatrix} E & \\ & -t \end{pmatrix}. \end{aligned}$$

If t is large enough so that an optimal solution X^* for the dual problem (Φ, E, F) has $\text{Tr}(X^*) \leq t$ then X^* is also optimal for (Φ', E', F) . In this case, we can obtain the optimal value for (Φ, E, F) by solving (Φ', E', F) . Thus, it suffices to prove well-boundedness for (Φ', E', F) .

Lemma 4.6 (Well-boundedness of (Φ', E', F)). *There exists a dual feasible X_0 for (Φ, E, F) and positive real numbers δ, t such that:*

1. *For all Hermitian operators H with $\|H\|_F \leq \delta$ it holds that $X_0 + H$ is dual feasible with $\text{Tr}(X_0 + H) \leq t$.*

2. The optimal value of the dual problem (Φ, E, F) is achieved by a dual feasible X^* with $\text{Tr}(X^*) \leq t$.

In particular, we may select

$$\begin{aligned}\delta &= \Omega (\dim(\mathcal{D}_{1\dots r})^{-1}), \\ t &= O(r^3 \dim(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r} \otimes \mathcal{C}_{1\dots r})).\end{aligned}$$

Proof. First, we select X_0 and δ ; the quantity t will be chosen at the end of the proof. We begin by choosing the first collection B_1, \dots, B_r of diagonal blocks for X_0 . For each $i = 1, \dots, r$ let

$$B_i = \frac{i+1}{\dim(\mathcal{D}_{1\dots i})} I_{\mathcal{D}_{1\dots i} \otimes \mathcal{B}_{1\dots i}}$$

and let Y_i be any Hermitian operator with

$$\|Y_i\| \leq \frac{1}{3 \dim(\mathcal{D}_{1\dots i})}$$

so that

$$\frac{i + \frac{2}{3}}{\dim(\mathcal{D}_{1\dots i})} I_{\mathcal{D}_{1\dots i} \otimes \mathcal{B}_{1\dots i}} \preceq B_i + Y_i \preceq \frac{i + \frac{4}{3}}{\dim(\mathcal{D}_{1\dots i})} I_{\mathcal{D}_{1\dots i} \otimes \mathcal{B}_{1\dots i}}.$$

As in the proof of Lemma 3.14, it is tedious but straightforward to verify that

$$\begin{aligned}\text{Tr}_{\mathcal{D}_1}(B_1 + Y_1) &\succ I_{\mathcal{B}_1} \\ \text{Tr}_{\mathcal{D}_2}(B_2 + Y_2) &\succ (B_1 + Y_1) \otimes I_{\mathcal{B}_2} \\ &\vdots \\ \text{Tr}_{\mathcal{D}_r}(B_r + Y_r) &\succ (B_{r-1} + Y_{r-1}) \otimes I_{\mathcal{B}_r}.\end{aligned}$$

Next, we choose the remaining diagonal blocks Q_1, \dots, Q_r for X_0 . To this end, let γ be a real number large enough to guarantee that

$$\Omega_R(B_r + Y_r) \preceq \gamma I_{\mathcal{C}_{1\dots r} \otimes \mathcal{A}_{1\dots r}},$$

regardless of the choice of Y_r . (The precise value of γ will be chosen later.) For each $i = 1, \dots, r$ let

$$Q_i = (r - i + 2) \dim(\mathcal{A}_{i+1\dots r}) \gamma I_{\mathcal{C}_{1\dots i-1} \otimes \mathcal{A}_{1\dots i}}$$

and let Z_i be any Hermitian operator with

$$\|Z_i\| \leq \frac{\dim(\mathcal{A}_{i+1\dots r}) \gamma}{3}$$

so that

$$\begin{aligned}
(Q_r + Z_r) \otimes I_{\mathcal{C}_r} &\succ \Omega_R(B_r + Y_r) \\
(Q_{r-1} + Z_{r-1}) \otimes I_{\mathcal{C}_{r-1}} &\succ \text{Tr}_{\mathcal{A}_r}(Q_r + Z_r) \\
&\vdots \\
(Q_1 + Z_1) \otimes I_{\mathcal{C}_1} &\succ \text{Tr}_{\mathcal{A}_2}(Q_2 + Z_2).
\end{aligned}$$

It follows from these semidefinite inequalities that the block-diagonal operator X_0 with diagonal blocks $B_1, \dots, B_r, Q_1, \dots, Q_r$ is a strictly dual feasible solution to (Φ, E, F) . Moreover, for any Hermitian operator H with diagonal blocks $Y_1, \dots, Y_r, Z_1, \dots, Z_r$ and $\|H\|_F \leq \delta$ it must be that $X_0 + H$ is also a strictly feasible solution, provided

$$\delta \leq \frac{1}{3 \dim(\mathcal{D}_{1\dots r})}.$$

Next, let us choose an appropriate value for γ , which will enable us to establish the desired upper bound t . For notational convenience, write

$$\nu = \frac{r + \frac{4}{3}}{\dim(\mathcal{D}_{1\dots r})}$$

so that

$$B_r + Y_r \preceq \nu I_{\mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r}}.$$

As Ω_R is completely positive, it follows that

$$\Omega_R(B_r + Y_r) \preceq \nu \Omega_R(I_{\mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r}}) = \nu \text{Tr}_{\mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r}}(R)$$

where R is an operator with $0 \preceq R \preceq J(\Psi^*)$ for some completely positive and trace-preserving super-operator $\Psi : \mathbf{L}(\mathcal{Y}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{X}_{1\dots r})$. Proposition 1.3 tells us that $\|J(\Psi^*)\|_{\text{Tr}} = \dim(\mathcal{Y}_{1\dots r})$, from which we obtain $\|R\| \leq \dim(\mathcal{Y}_{1\dots r})$ and hence

$$R \preceq \dim(\mathcal{Y}_{1\dots r}) I_{\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}}.$$

It follows that

$$\Omega_R(B_r + Y_r) \preceq \nu \dim(\mathcal{Y}_{1\dots r} \otimes \mathcal{D}_{1\dots r} \otimes \mathcal{B}_{1\dots r}) I_{\mathcal{C}_{1\dots r} \otimes \mathcal{A}_{1\dots r}}$$

and hence we may select

$$\gamma = \left(r + \frac{4}{3}\right) \dim(\mathcal{Y}_{1\dots r} \otimes \mathcal{B}_{1\dots r}).$$

It is straightforward but tedious to verify that for every Hermitian H with $\|H\|_F \leq \delta$ it holds that

$$\text{Tr}(X_0 + H) < r \left(r + \frac{4}{3}\right) (\dim(\mathcal{B}_{1\dots r}) + \gamma \dim(\mathcal{A}_{1\dots r} \otimes \mathcal{C}_{1\dots r})).$$

Substituting our choice of γ , we find that it suffices to select

$$t = 2(r + 2)^3 \dim(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r} \otimes \mathcal{C}_{1\dots r}).$$

To see that this choice of t also bounds the trace of an optimal solution, we note that by Lemma 4.4 the diagonal blocks $B_1^*, \dots, B_r^*, Q_1^*, \dots, Q_r^*$ of any optimal solution X^* can be assumed to be strategies and co-strategies, respectively, from which it is easy to see that

$$\text{Tr}(X^*) < r(\dim(\mathcal{B}_{1\dots r}) + \dim(\mathcal{C}_{1\dots r})) < t.$$

□

Lemma 4.6 (Well-boundedness of (Φ', E', F)) provides all that we need in order to apply Fact 4.5 (Efficient algorithm for semidefinite optimization). While the bound t on the set of feasible solutions to (Φ', E', F) from Lemma 4.6 is stated in terms of the trace norm, it is straightforward to convert this quantity into a bound Δ in terms of the Frobenius norm via the norm inequalities listed in Section 1.2.1. The following theorem is now proved.

Theorem 4.7 (Efficient algorithm to compute the value of a quantum game). *The following problem admits a deterministic polynomial-time solution:*

Input. *A two-player zero-sum quantum game specified by an r -round quantum referee $\{R_m\}$, a payout function $V(m)$, and an accuracy parameter $\varepsilon > 0$. The real numbers $V(m)$ and ε are each given explicitly in binary, as are the real and complex parts of each entry of the matrices R_m .*

Output. *A real number v such that the value of the game specified by $\{R_m\}$ and $V(m)$ lies in the open interval $(v - \varepsilon, v + \varepsilon)$.*

4.2 Quantum interactive proofs

In this section we provide an application of the theory of zero-sum quantum games developed in Section 4.1 to quantum interactive proofs. In particular, we observe that the existence of an efficient algorithm that computes the value of a zero-sum quantum game implies the equivalence of the complexity classes QRG and EXP. Here QRG is the class of problems that admit quantum interactive proofs with two competing provers, whereas EXP is the fundamental class of problems that admit deterministic (classical) exponential-time solutions.

We also extend an existing parallel repetition result for single-prover quantum interactive proofs so that it applies to interactions with an arbitrary number of messages, as opposed to interactions with only three messages. The technique employed toward this end is then applied to yield a similar parallel repetition result for quantum interactive proofs with competing provers.

This section begins with a primer on interactive proofs, followed by a brief historical survey before the results are presented.

Interactive proofs

An *interactive proof* consists of an interaction between a *verifier* and a *prover* regarding some common input x , which is viewed as an instance of a yes-no decision problem P . Throughout the interaction the verifier is restricted to randomized polynomial-time computations, while the prover’s computational power is unlimited. After the interaction the verifier produces a binary outcome indicating acceptance or rejection of x . The verifier’s goal is to accept those inputs that are yes-instances of P and reject those inputs that are no-instances of P . Typically, the verifier does not have the computational power to make this determination himself and so he must look to the prover and his unlimited computational power for help. However, the goal of the prover is always to convince the verifier to accept, regardless of whether the input x is actually a yes-instance. Hence, the verifier must be careful to distinguish truthful proofs from false proofs.

A decision problem P is said to *admit an interactive proof* if there exists a verifier with the property that a prover can convince him to accept every yes-instance of P with high probability, yet no prover can convince him to accept any no-instance of P except with small probability. Such an interactive proof is said to *solve* the problem P . The complexity class of decision problems that admit interactive proofs is denoted IP. In a *quantum* interactive proof, the verifier and prover may process and exchange quantum information. Whereas a classical verifier is restricted to randomized polynomial-time computations, a quantum verifier is instead restricted to quantum circuits that can be generated uniformly in deterministic polynomial time. The corresponding complexity class is denoted QIP.

An interactive proof *with competing provers* has two provers—one, as before, whose goal is always to convince the verifier to accept, and another, whose goal is to convince the verifier to reject. The class of problems that admit interactive proofs with competing provers is denoted RG, for “refereed games”, owing to the similarity between these two models of interaction. As above, we may speak of *quantum* interactive proofs with competing provers and the associated complexity class QRG.

Background

Interactive proofs were introduced [Bab85, GMR89] as a generalization of efficiently verifiable proofs—a fruitful concept that characterizes the ubiquitous complexity class NP. The generalization lies in the interaction: whereas any problem in NP may be verified with a single message from the prover to the verifier, interactive proofs allow the verifier to ask a series of questions of the prover, possibly basing future questions upon previous answers. In order to make this generalization from NP to IP nontrivial, the verifier is permitted a source of randomness. (Hence the allowance in interactive proofs for a small probability of error.)

Naturally, every problem in NP admits an interactive proof. Moreover, it is conjectured that interactive proofs with only a constant number of messages cannot

solve problems outside NP [MV99]. On the other hand, interactive proofs with an unbounded number of messages are surprisingly powerful: every problem in PSPACE admits such an interactive proof [LFKN92, Sha92]. By contrast, in any *quantum* interactive proof it suffices that the verifier and prover exchange only three messages [KW00]. It is not difficult to see that PSPACE contains IP, from which the characterization $IP = PSPACE$ then follows. In a recent breakthrough, Jain, Ji, Upadhyay, and Watrous proved that PSPACE also contains QIP, from which one obtains $QIP = PSPACE$ [JJUW09].

Feige and Kilian showed that every problem that can be solved in deterministic exponential-time can also be solved by an interactive proof with competing provers [FK97]. That these proofs can be simulated in deterministic exponential time follows from Ref. [KM92], and hence we have the complexity theoretic equality $RG = EXP$. In this section we employ the algorithm from Section 4.1.2 for computing the value of a zero-sum quantum game to show that $QRG = EXP$.

Containment of QRG inside EXP

Theorem 4.8 ($QRG = EXP$). *Every decision problem that admits a quantum interactive proof with competing provers also admits a deterministic exponential-time solution. It follows that $QRG = EXP$.*

Proof. Let P be any problem that admits a quantum interactive proof with competing provers. We will exhibit a deterministic exponential-time algorithm that, for each input x , correctly decides whether x is a yes-instance of P or a no-instance.

For each fixed input x , the actions of the verifier may be represented by a two-outcome measuring co-strategy $\{R_{\text{accept}}, R_{\text{reject}}\}$ where the two provers combine to implement the corresponding strategy. The operators in this co-strategy may be computed explicitly in deterministic exponential time by running the polynomial-time algorithm that generates a description of the quantum circuits corresponding to the actions of the verifier on input x and then converting that description into an exponential-sized measuring co-strategy in the standard way.

Consider the two-player zero-sum quantum refereed game defined by the referee $\{R_{\text{accept}}, R_{\text{reject}}\}$ and the payout function

$$V(\text{accept}) = 1, \quad V(\text{reject}) = 0.$$

The value of this game equals the probability with which the verifier is convinced to accept the input when the two-provers act according to optimal strategies. By Theorem 4.7 (Efficient algorithm to compute the value of a quantum game), there is a deterministic algorithm that approximates this value to arbitrary precision in time polynomial in the bit length of $\{R_{\text{accept}}, R_{\text{reject}}\}$, which we know to be exponential in the input x . By running this algorithm and computing this probability, it is possible to determine whether or not x is a yes-instance of P . Thus, $P \in EXP$. As P was selected arbitrarily from QRG, it follows that $QRG \subseteq EXP$. As QRG is already known to contain EXP, it follows that $QRG = EXP$. \square

Parallel repetition of many-message quantum interactive proofs

Thus far, our discussion of interactive proofs has been restricted to protocols that succeed with high probability. But this condition is quite vague—what is a “high” probability? Can interactive proofs be somehow transformed so as to amplify their probability of success?

To address these questions, we need a more specific definition of interactive proof. A decision problem P is said to admit an interactive proof with *completeness error* c and *soundness error* s if the following hold:

- (i) If x is a yes-instance of P then there is a prover who can convince the verifier to accept with probability at least $1 - c$.
- (ii) If x is a no-instance of P then no prover can convince the verifier to accept with probability larger than s .

It is not difficult to see that *any decision problem whatsoever* admits an interactive proof with $1 - c \leq s$. Clearly then, interactive proofs are only interesting when $1 - c > s$. When this condition is met, the interactive proof may be transformed so as to achieve $c, s < \varepsilon$ for any desired $\varepsilon > 0$. This reduction in error is achieved by a method called *sequential repetition*, whereby the verifier simply repeats the protocol many times in succession and then bases his final decision upon a weighted vote of the outcomes of the individual repetitions. Under sequential repetition, the probability of error decreases exponentially in the number of repetitions.

A side-effect of sequential repetition is that the number of messages exchanged between the verifier and prover increases with each repetition. Sometimes, it is desirable to achieve error reduction without increasing the number of messages in the interaction. The standard transformation meeting this criterion is called *parallel repetition*, whereby many copies of the protocol are executed simultaneously instead of sequentially. The danger of parallel repetition is that the prover need not treat each repetition independently. Instead, he might somehow attempt to correlate the repetitions so as to thwart the exponential reduction in error. Therefore, before parallel repetition may be relied upon to reduce error, it must first be established that the prover cannot significantly affect the probability of acceptance by deviating from a strategy that treats the repetitions independently.

For classical single-prover interactive proofs, it can be shown that parallel repetition works as desired to achieve exponential error reduction. For quantum single-prover interactive proofs, parallel repetition is known to work only when

- (i) the verifier and prover exchange at most *three messages*, and
- (ii) the verifier’s final decision depends upon a *unanimous vote* of the repetitions (as opposed to some other weighting of the repetitions, such as a majority vote).

(See Kitaev and Watrous for proofs of this and other properties of quantum interactive proofs [KW00].)

In this section, we show that condition (i) is unnecessary. Specifically, we prove the following.

Theorem 4.9 (Parallel repetition of many-message quantum interactive proofs). *Consider the k -fold parallel repetition of an arbitrary r -round quantum interactive proof with soundness error s . If the verifier in the repeated protocol accepts only when all k repetitions accept then the repeated protocol has soundness error s^k .*

Proof. Suppose that the interactive proof in the statement of the theorem solves the decision problem P . The theorem makes no claim about the case where the input is a yes-instance of P , so we need only consider those inputs which are no-instances of P . We must show that the verifier in the repeated protocol can be made to accept such an input with probability not exceeding s^k .

For any such input, the actions of the verifier may be represented by a two-outcome r -round measuring strategy $\{R_{\text{accept}}, R_{\text{reject}}\}$ where the prover implements the corresponding r -round co-strategy. As the verifier has soundness error s , it follows from Theorem 3.3 (Maximum output probability) that there exists an r -round non-measuring strategy R for the referee's input and output spaces with the property that

$$R_{\text{accept}} \preceq sR.$$

It follows from a semidefinite inequality proven in Ref. [CSUU06, Proposition 5] that

$$R_{\text{accept}}^{\otimes k} \preceq s^k R^{\otimes k}.$$

(The cited inequality is not as explicit in the final journal version [CSUU08] of Ref. [CSUU06].)

For the repeated protocol with unanimous vote, Theorem 3.3 implies that the maximum probability with which the verifier can be made to accept is given by

$$\min \{ \lambda \geq 0 : R_{\text{accept}}^{\otimes k} \preceq \lambda Q \text{ for some } r\text{-round non-measuring strategy } Q \}.$$

Taking $\lambda = s^k$ and $Q = R^{\otimes k}$ completes the proof. □

Theorem 4.9 makes no claim about the effect of a unanimous vote on the completeness error. But for single-prover interactive proofs the question is moot, since these proofs can be assumed to have zero completeness error [KW00]. In particular, for yes-instances of P the prover in the k -fold repeated protocol can achieve zero completeness error by playing an independent copy of the zero-error strategy for each of the k repetitions.

Parallel repetition of quantum interactive proofs with competing provers

Fortunately, Theorem 4.9 can be adapted with little difficulty to say something meaningful about the parallel repetition of quantum interactive proofs with competing provers.

Theorem 4.10 (Parallel repetition of quantum interactive proofs with competing provers). *Consider the k -fold parallel repetition of an arbitrary r -round quantum interactive proof with competing provers that has completeness error c and soundness error s .*

If the verifier in the repeated protocol accepts only when all k repetitions accept then the repeated protocol has completeness error kc and soundness error s^k . Similarly, if the verifier in the repeated protocol rejects only when all k repetitions reject then the repeated protocol has completeness error c^k and soundness error ks .

Proof. We prove only the first claim in the statement of the theorem, as the second claim follows by symmetry.

Suppose that the interactive proof in the statement of the theorem solves the decision problem P and suppose first that the input is a yes-instance of P . We must show that there is a “yes-prover” who convinces the verifier in the repeated protocol to accept with probability at least $1 - kc$, regardless of the strategy employed by the other prover, whom we call the “no-prover”.

The yes-prover we seek merely plays an independent copy of the optimal strategy for each of the k repetitions. Of course, no no-prover can win any one of the k repetitions with probability greater than c . It then follows from the union bound that the repeated game with unanimous vote has completeness error at most kc .

Now suppose that the input is a no-instance of P . For this case, we must find a no-prover who convinces the verifier in the repeated protocol to reject with probability at least $1 - s^k$, regardless of the strategy employed by the yes-prover.

Toward that end, consider an optimal no-prover for the original protocol. Borrowing from the proof of Theorem 4.9, the combined actions of the verifier and no-prover in this protocol may be represented by a two-outcome r -round measuring strategy $\{N_{\text{accept}}, N_{\text{reject}}\}$ where the yes-prover implements the corresponding r -round co-strategy. That the verifier-no-prover combination for the repeated protocol specified by $\{N'_{\text{accept}}, N'_{\text{reject}}\}$ with

$$N'_{\text{accept}} = N_{\text{accept}}^{\otimes k}$$

has soundness error s^k now follows exactly as in the proof of Theorem 4.9.

Moreover, we see from this proof that the desired optimal no-prover for the repeated protocol simply plays an independent copy of the optimal strategy for each of the k repetitions. \square

Theorem 4.10 tells us that we may achieve an exponential reduction in completeness (or soundness) at the cost of a linear increase in soundness (or completeness). Such a transformation is useful, for example, for protocols for which either completeness or soundness is already small—these proofs can be transformed so that *both* completeness *and* soundness are small.

More specifically, suppose we have a family of protocols with the property that, for some fixed constants $\varepsilon, s \in (0, 1)$ and for any desired m , there is a protocol in the family with completeness error ε^m and soundness error s . Then for any desired k , we may choose m large enough and employ k -fold parallel repetition with unanimous vote to obtain a protocol in which both completeness error *and* soundness error are no larger than s^k . Indeed, such a transformation was employed in Ref. [GW05] to reduce error for so-called “short quantum games”.

4.3 Kitaev’s bound for strong coin-flipping

In this section we provide a simplified proof of Kitaev’s bound for strong quantum coin-flipping, which states that a cheating party can always force any desired outcome upon an honest party with probability at least $1/\sqrt{2}$. We begin with a definition of the coin-flipping problem, followed by a brief survey, before ending the section with our contribution.

Strong and weak coin-flipping

Suppose that two parties—Alice and Bob—wish to agree on a random bit. (That is, they wish to flip a fair coin.) The parties are physically separated, so the agreement must be reached by an exchange of messages via remote communication. Moreover, the two parties do not trust each other, meaning that each party suspects that the other might attempt to force a particular result of the coin flip rather than settle for a uniformly random result. The goal is to devise a protocol that produces the fairest possible coin flip, even in the case where one cheating party deliberately attempts to bias the result of the flip.

For example, Alice could simply flip a fair coin for herself and announce the result to Bob, who meekly agrees to whatever Alice dictates. If both parties adhere to these honest strategies then it is clear that they will produce a perfectly correlated random bit as desired. Even if Bob cheats, there’s nothing he can do to prevent Alice from producing a perfectly fair coin toss. (Indeed, Alice’s honest strategy is to completely ignore Bob.) But if Alice decides to cheat then Bob is out of luck; his mindless strategy of taking Alice at her word allows Alice to force Bob to produce any outcome she desires. While this simple protocol is robust against a cheating Bob, it is completely vulnerable to a cheating Alice and is therefore a poor solution to the problem at hand.

Let us be more specific about this problem. A *strong coin-flipping protocol with bias ε* consists of an honest strategy for Alice and an honest strategy for Bob such that:

1. If both parties follow their honest strategies then both parties produce the same outcome, and that outcome is chosen uniformly at random from $\{0, 1\}$.
2. If only one party follows his or her honest strategy then the maximum probability with which the cheating party can force the honest party to produce a given outcome is at most $1/2 + \varepsilon$.

The adjective *strong* refers to the fact that a cheater cannot bias an honest party's outcome toward either result 0 or 1. By contrast, *weak* protocols assume that one player desires outcome 0 and the other desires outcome 1. The only requirement for weak protocols is that a cheater cannot bias the result toward his or her desired outcome. In a *quantum* coin-flipping protocol, the parties may process and exchange quantum information.

Background

This problem was introduced in the classical setting by Blum [Blu81], who called it *coin-flipping by telephone*. The problem is of interest to cryptographers because it is a primitive—a basic tool—for secure two-party computations. In this setting there is little need to distinguish between strong and weak coin-flipping. Indeed, Blum showed that even strong coin-flipping with zero bias is possible under certain computational assumptions. Conversely, without any such assumptions it is not difficult to see that even weak coin-flipping is impossible. In particular, for each outcome $b \in \{0, 1\}$ either (i) a computationally unrestricted cheating Alice can force outcome b upon honest-Bob with certainty, or (ii) a computationally unrestricted cheating Bob can force the opposite outcome upon honest-Alice with certainty.

The quantum version of this problem admits some surprising contrasts to its classical counterpart. Whereas unconditional classical coin-flipping is impossible, there is a weak quantum coin-flipping protocol that achieves *arbitrarily small bias* in the limit of the number of messages exchanged between the parties [Moc07]. Moreover, the existence of such a protocol for weak quantum coin-flipping implies the existence of a *strong* quantum coin-flipping protocol that achieves bias arbitrarily close to $1/\sqrt{2} - 1/2 \approx 0.207$ [CK09]. This protocol for strong quantum coin-flipping is the best possible, as it matches a lower bound due to Kitaev [Kit02]. (Kitaev did not publish this proof, but it appears in Refs. [ABDR04, Röh04].) More comprehensive histories of quantum coin-flipping can be found in the breakthrough works of Mochon [Moc07], and Chailloux and Kerenidis [CK09].

Our contribution

We now provide an alternate and simplified proof of Kitaev's lower bound for strong quantum coin-flipping.

Theorem 4.11 (Kitaev's bound for strong quantum coin-flipping). *In any strong quantum coin-flipping protocol, a computationally unrestricted cheating party can always force a given outcome on an honest party with probability at least $1/\sqrt{2}$.*

Our new proof. In any quantum coin-flipping protocol the actions of honest-Alice are represented by a two-outcome measuring strategy $\{A_0, A_1\}$ and the actions of honest-Bob are represented by a two-outcome measuring co-strategy $\{B_0, B_1\}$. By Theorem 3.1 (Interaction output probabilities), the definition of a quantum coin-flipping protocol requires

$$1/2 = \langle A_0, B_0 \rangle = \langle A_1, B_1 \rangle.$$

Choose any outcome $b \in \{0, 1\}$ and let p denote the maximum probability with which a cheating Bob could force honest-Alice to output b . Obviously we have $p \geq 1/2$. Theorem 3.3 (Maximum output probability) implies that there must exist a strategy Q for Alice such that $A_b \preceq pQ$. If a cheating Alice plays this strategy Q then honest-Bob outputs b with probability

$$\langle Q, B_b \rangle \geq \frac{1}{p} \langle A_b, B_b \rangle = \frac{1}{2p}.$$

As $\max\{p, \frac{1}{2p}\} \geq 1/\sqrt{2}$ for all $p > 0$, it follows that either honest-Alice or honest-Bob can be convinced to output b with probability at least $1/\sqrt{2}$. \square

This proof makes clear the limitations of strong coin-flipping protocols: the inability of Bob to force Alice to output b directly implies that Alice can herself bias the outcome toward b . By definition, weak coin-flipping protocols are not subject to this limitation.

Chapter 5

Distance Measures

In this chapter we introduce a new norm for super-operators that generalizes the diamond norm and we argue that this norm quantifies the observable difference between quantum strategies. Indeed, we establish an extension of a result from Ref. [GW05] stating that each pair of convex sets $\mathbf{S}_0, \mathbf{S}_1$ of r -round strategies has a fixed r -round measuring co-strategy $\{T_a\}$ with the property that any pair $S_0 \in \mathbf{S}_0, S_1 \in \mathbf{S}_1$ can be distinguished by $\{T_a\}$ with probability determined by the minimal distance between \mathbf{S}_0 and \mathbf{S}_1 as measured by our new norm.

In order to prove this distinguishability result we also establish several properties of the new norm, including characterizations of its unit ball and dual norm. We end the chapter with a discussion of the dual of the diamond norm.

When not explicitly stated otherwise, the results of this chapter are the sole work of the author and were not published prior to the present thesis.

5.1 A new norm for strategies

In order to best argue that our new norm captures the distinguishability of quantum strategies, we begin with a discussion of the trace norm and explain its use in distinguishing quantum states. We then proceed with a discussion of the diamond norm for super-operators and explain its use in distinguishing quantum operations. Only then do we introduce our new norm and note the similarities with the diamond norm and trace norm for the purpose of distinguishing quantum strategies. We conclude the section with several immediate observations concerning this new norm, including the fact that it agrees with the diamond norm wherever it is defined.

The trace norm as a distance measure for quantum states

Recall from Section 1.2.1 that the *trace norm* $\|X\|_{\text{Tr}}$ of an arbitrary operator X is defined as the sum of the singular values of X . If X is Hermitian then it is a simple

exercise to verify that its trace norm is given by

$$\begin{aligned}\|X\|_{\text{Tr}} &= \max \{ \langle P_0 - P_1, X \rangle : P_0, P_1 \succeq 0, P_0 + P_1 = I \} \\ &= \max \{ \langle P_0 - P_1, X \rangle : \{P_0, P_1\} \text{ is a two-outcome measurement} \}.\end{aligned}$$

The trace norm provides a physically meaningful distance measure for quantum states in the sense that it captures the maximum likelihood with which two states can be correctly distinguished. Let us illustrate this fact with a simple example involving two parties called Alice and Bob and a fixed pair of quantum states ρ_0, ρ_1 . Suppose Bob selects a state $\rho \in \{\rho_0, \rho_1\}$ uniformly at random and gives Alice a quantum system prepared in state ρ . Alice has a complete description of both ρ_0 and ρ_1 , but she does not know which of the two was selected by Bob. Her goal is to correctly guess which of $\{\rho_0, \rho_1\}$ was selected based upon the outcome of a measurement she conducts on ρ .

Since Alice's guess is binary-valued and completely determined by her measurement, that measurement can be assumed to be a two-outcome measurement $\{P_0, P_1\}$ wherein outcome $a \in \{0, 1\}$ indicates a guess that Bob prepared $\rho = \rho_a$. Letting C denote the event that Alice's guess is correct, basic quantum formalism tells us that

$$\begin{aligned}\Pr[C] &= \frac{1}{2} \langle P_0, \rho_0 \rangle + \frac{1}{2} \langle P_1, \rho_1 \rangle \\ \Pr[\neg C] &= \frac{1}{2} \langle P_1, \rho_0 \rangle + \frac{1}{2} \langle P_0, \rho_1 \rangle,\end{aligned}$$

implying that

$$\Pr[C] - \Pr[\neg C] = \frac{1}{2} \langle P_0 - P_1, \rho_0 - \rho_1 \rangle.$$

As $\Pr[C] + \Pr[\neg C] = 1$, we obtain the following alternate expression for the probability with which Alice's guess is correct:

$$\Pr[C] = \frac{1}{2} + \frac{1}{4} \langle P_0 - P_1, \rho_0 - \rho_1 \rangle \leq \frac{1}{2} + \frac{1}{4} \|\rho_0 - \rho_1\|_{\text{Tr}}$$

with equality achieved at the optimal measurement $\{P_0, P_1\}$ for Alice. This fundamental observation was originally made by Helstrom [Hel69].

The diamond norm as a distance measure for quantum operations

The trace norm extends naturally to super-operators, but this extension does not lead to an overly useful distance measure for quantum operations. To achieve such a measure, the trace norm must be modified as follows.

Definition 5.1 (Super-operator trace norm, diamond norm). For an arbitrary super-operator Φ the *super-operator trace norm* $\|\Phi\|_{\text{Tr}}$ of Φ is defined as

$$\|\Phi\|_{\text{Tr}} \stackrel{\text{def}}{=} \max_{\|X\|_{\text{Tr}}=1} \|\Phi(X)\|_{\text{Tr}}.$$

The *diamond norm* $\|\Phi\|_\diamond$ of Φ is defined as

$$\|\Phi\|_\diamond \stackrel{\text{def}}{=} \sup_{\mathcal{W}} \|\Phi \otimes \mathbb{1}_{\mathcal{W}}\|_{\text{Tr}}$$

where the supremum is taken over all finite-dimensional complex Euclidean spaces \mathcal{W} . \square

Much is known of the diamond norm. For example, if Φ has the form $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ then the supremum in the definition of $\|\Phi\|_\diamond$ is always achieved by some space \mathcal{W} whose dimension does not exceed that of the input space \mathcal{X} [Kit97, AKN98]. As a consequence, the supremum in the definition of the diamond norm can be replaced by a maximum. Moreover, if Φ is Hermitian-preserving and $\dim(\mathcal{W}) \geq \dim(\mathcal{X})$ then the maximum in the definition of $\|\Phi \otimes \mathbb{1}_{\mathcal{W}}\|_{\text{Tr}}$ is always achieved by some positive semidefinite operator X [RW05]. Thus, if Φ is Hermitian-preserving then its diamond norm is given by

$$\begin{aligned} \|\Phi\|_\diamond &= \max \|(\Phi \otimes \mathbb{1}_{\mathcal{W}})(\rho)\|_{\text{Tr}} \\ &= \max \langle P_0 - P_1, (\Phi \otimes \mathbb{1}_{\mathcal{W}})(\rho) \rangle \end{aligned}$$

where the maxima in these two expressions are taken over all spaces \mathcal{W} with dimension at most \mathcal{X} , all states $\rho \in \mathbf{H}^+(\mathcal{X} \otimes \mathcal{W})$, and all two-outcome measurements $\{P_0, P_1\} \subset \mathbf{H}^+(\mathcal{Y} \otimes \mathcal{W})$.

The diamond norm is to quantum operations as the trace norm is to quantum states: it provides a physically meaningful distance measure for quantum operations in the sense that the value $\|\Phi_0 - \Phi_1\|_\diamond$ quantifies the observable difference between two quantum operations Φ_0, Φ_1 . As before, this fact may be illustrated with a simple example. Suppose Bob selects an operation from $\Phi \in \{\Phi_0, \Phi_1\}$ uniformly at random. Alice is granted “one-shot, black-box” access to Φ and her goal is to correctly guess which of Φ_0, Φ_1 was applied. Specifically, Alice may prepare a quantum system in state ρ and send a portion of that system to Bob, who applies Φ to that portion and then returns it to Alice. Finally, Alice performs a two-outcome measurement $\{P_0, P_1\}$ on the resulting system $(\Phi \otimes \mathbb{1})(\rho)$ where outcome $a \in \{0, 1\}$ indicates a guess that $\Phi = \Phi_a$.

Letting C denote the event that Alice’s guess is correct, we may repeat our previous derivation of $\Pr[C]$ to obtain

$$\begin{aligned} \Pr[C] &= \frac{1}{2} + \frac{1}{4} \langle P_0 - P_1, (\Phi_0 \otimes \mathbb{1})(\rho) - (\Phi_1 \otimes \mathbb{1})(\rho) \rangle \\ &\leq \frac{1}{2} + \frac{1}{4} \|\Phi_0 - \Phi_1\|_\diamond \end{aligned}$$

with equality achieved at the optimal input state ρ and measurement $\{P_0, P_1\}$ for Alice.

It is interesting to note that the ability to send only *part* of the input state ρ to Bob and keep the rest for herself can enhance Alice’s ability to distinguish some pairs of quantum operations, as compared to a simpler test that involves

sending the *entire* input state to Bob. Indeed, there exist pairs Φ_0, Φ_1 of quantum operations that are perfectly distinguishable when applied to half of a maximally entangled input state—that is, $\|\Phi_0 - \Phi_1\|_\diamond = 2$ —yet they appear nearly identical when an auxiliary system is not used—that is, $\|\Phi_0 - \Phi_1\|_{\text{Tr}} \approx 0$. An example of such a pair of super-operators can be found in Watrous [Wat08], along with much of the discussion that has occurred thus far in this section. It is this phenomenon that renders the super-operator trace norm less useful than the diamond norm for the study of quantum information.

Additional properties of the diamond norm are established in this thesis by Lemmas 5.12, 5.14, and 5.17 of Section 5.3.

The strategy r -norm as a distance measure for quantum strategies

The simple guessing game between Alice and Bob extends naturally from quantum operations to quantum strategies. Let S_0, S_1 be arbitrary r -round quantum strategies and suppose Bob selects $S \in \{S_0, S_1\}$ uniformly at random. Alice’s task is to interact with Bob and then decide after the interaction whether Bob selected $S = S_0$ or $S = S_1$.

Theorem 3.1 (Interaction output probabilities) establishes an inner product relationship between measuring strategies and co-strategies that is analogous to the relationship between states and measurements. As such, much of our previous discussion concerning the task of distinguishing pairs of *states* can be re-applied to the task of distinguishing pairs of *strategies*. In particular, Alice can be assumed to act according to some two-outcome r -round measuring co-strategy $\{T_0, T_1\}$ for Bob’s input and output spaces, with outcome $a \in \{0, 1\}$ indicating a guess that Bob acted according to strategy S_a . Moreover, letting C denote the event that Alice’s guess is correct, we have

$$\Pr[C] = \frac{1}{2} + \frac{1}{4}\langle T_0 - T_1, S_0 - S_1 \rangle.$$

Naturally, Alice maximizes the probability of a correct guess by maximizing this expression over all r -round measuring co-strategies $\{T_0, T_1\}$.

Of course, this guessing game is symmetric with respect to strategies and co-strategies. In particular, if Bob’s actions S_0, S_1 are *co-strategies* instead of strategies then Alice’s actions $\{T_0, T_1\}$ must be a measuring *strategy* instead of a measuring co-strategy. The probability with which Alice correctly guesses Bob’s strategy is still given by

$$\Pr[C] = \frac{1}{2} + \frac{1}{4}\langle T_0 - T_1, S_0 - S_1 \rangle$$

except that Alice now maximizes this probability over all r -round measuring *strategies* $\{T_0, T_1\}$.

With this example in mind, we propose two new norms—one that captures the distinguishability of strategies and one that captures the distinguishability of co-strategies.

Definition 5.2 (Strategy r -norm). For any Hermitian-preserving super-operator $\Phi : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$ we define

$$\begin{aligned} \|\Phi\|_{\diamond r} &\stackrel{\text{def}}{=} \max \{ \langle T_0 - T_1, J(\Phi) \rangle : \{T_0, T_1\} \text{ is a } r\text{-round measuring co-strategy} \}, \\ \|\Phi\|_{\diamond r}^* &\stackrel{\text{def}}{=} \max \{ \langle S_0 - S_1, J(\Phi) \rangle : \{S_0, S_1\} \text{ is a } r\text{-round measuring strategy} \}. \end{aligned}$$

Sometimes it is convenient to write $\|J(\Phi)\|_{\diamond r}$ instead of $\|\Phi\|_{\diamond r}$, particularly when $J(\Phi)$ is an operator derived from quantum strategies. Similarly, we may sometimes write $\|J(\Phi)\|_{\diamond r}^*$ instead of $\|\Phi\|_{\diamond r}^*$. \square

If S_0, S_1 are strategies for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$ then it follows immediately that the maximum probability with which Alice can correctly distinguish S_0 from S_1 is

$$\frac{1}{2} + \frac{1}{4} \|S_0 - S_1\|_{\diamond r}$$

Likewise, if S_0, S_1 are co-strategies rather than strategies then the maximum probability with which Alice can correctly distinguish S_0 from S_1 is

$$\frac{1}{2} + \frac{1}{4} \|S_0 - S_1\|_{\diamond r}^*$$

It may seem superfluous to allow both strategies and co-strategies as descriptions for Bob's actions in this simple example, as every co-strategy may be written as a strategy via suitable relabelling of input and output spaces. But there is something to be gained by considering both the norms $\|\cdot\|_{\diamond r}$ and $\|\cdot\|_{\diamond r}^*$. Later, we will show that these norms are dual to each other and we will use this duality to generalize the simple guessing game of this section.

Both a trace norm and an operator norm

We just argued that the norms $\|\cdot\|_{\diamond r}$ and $\|\cdot\|_{\diamond r}^*$ are to strategies and co-strategies as the trace norm is to states—a distance measure.

But these new norms also admit a different interpretation. It is easy to see that the maximum probability over all states ρ with which a standard quantum measurement $\{P_a\}$ produces a given outcome a is given by $\|P_a\|$. (This maximum is achieved for $\rho = vv^*$ where v is an eigenvector associated with the largest eigenvalue of P_a .) This observation extends unhindered to measuring strategies and co-strategies.

Proposition 5.3 (Maximum output probability from the strategy r -norm). *Let $\{S_a\}$ be an r -round measuring strategy. The maximum probability with which a compatible r -round co-strategy could force $\{S_a\}$ to produce a given measurement outcome a is given by $\|S_a\|_{\diamond r}$.*

Similarly, if $\{T_b\}$ is an r -round measuring co-strategy then the maximum probability with which a compatible r -round strategy could force $\{T_b\}$ to produce a given measurement outcome b is given by $\|T_b\|_{\diamond r}^$.*

Proof. The proofs of the two statements are completely symmetric, so we prove only the first. As each S_a is positive semidefinite, the maximum in the definition of $\|S_a\|_{\diamond r}$ is achieved at an r -round measuring co-strategy $\{T_0, T_1\}$ with $T_1 = 0$. That is,

$$\begin{aligned} \|S_a\|_{\diamond r} &= \max \{ \langle T_0 - T_1, S_a \rangle : \{T_0, T_1\} \text{ is an } r\text{-round measuring co-strategy} \} \\ &= \max \{ \langle T_0, S_a \rangle : T_0 \text{ is an } r\text{-round non-measuring co-strategy} \} \end{aligned}$$

The proposition follows from Theorem 3.1 (Interaction output probabilities). \square

Thus, the strategy r -norm and its dual are to measuring strategies and co-strategies as the operator norm is to measurements—a measure of maximum output probability. Juxtaposing these two interpretations, we see that both the norms $\|\cdot\|_{\diamond r}$ and $\|\cdot\|_{\diamond r}^*$ act as *both* a trace norm *and* an operator norm in the appropriate contexts.

The strategy 1-norm agrees with the diamond norm

As suggested by the notation $\|\cdot\|_{\diamond r}$, this new norm agrees with the diamond norm on Hermitian-preserving super-operators for the case $r = 1$. The proof of this fact is a simple exercise that serves as a convenient introduction to some of the techniques that will be employed later in this chapter.

Proposition 5.4 (Agreement with the diamond norm). *For every Hermitian-preserving super-operator Φ it holds that $\|\Phi\|_{\diamond} = \|\Phi\|_{\diamond 1}$.*

Proof. The proposition is trivially true for $\Phi = 0$, so we assume $\Phi \neq 0$ throughout. We also assume that Φ has the form $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$.

The proposition is a simple consequence of Theorem 3.1 that is somewhat complicated by the fact that it must be proven for *all* Hermitian-preserving super-operators—not just those which represent strategies. Strictly speaking, Theorem 3.1 applies only to measuring strategies and co-strategies and therefore cannot be directly applied to an arbitrary Hermitian-preserving super-operator Φ . In order to use Theorem 3.1, we first decompose Φ into a linear combination of measuring strategy elements.

Toward that end, we note that the Hermitian-preserving property of Φ implies that $J(\Phi)$ is a Hermitian operator. In particular, it has a Jordan decomposition

$$J(\Phi) = R^+ - R^-$$

where the operators $R^+, R^- \in \mathbf{H}^+(\mathcal{Y} \otimes \mathcal{X})$ are positive semidefinite and act on orthogonal subspaces. Let ε be a positive real number with the property that εR^+ and εR^- are both elements of some three-outcome one-round measuring strategy

$\{S_0, S_1, S_2\}$ for the input space \mathcal{X} and output space \mathcal{Y} . An easy way to accomplish this is to take

$$\varepsilon = \frac{1}{\dim(\mathcal{Y}) \max\{\|R^+\|, \|R^-\|\}},$$

$$(S_0, S_1, S_2) = \left(\varepsilon R^+, \varepsilon R^-, \frac{1}{\dim(\mathcal{Y})} I_{\mathcal{Y} \otimes \mathcal{X}} - \varepsilon R^+ - \varepsilon R^- \right).$$

For any one-round measuring co-strategy $\{T_0, T_1\}$ for the input space \mathcal{X} and output space \mathcal{Y} , Theorem 3.1 tells us that the probability with which an interaction between $\{S_a\}$ and $\{T_b\}$ yields measurement outcomes (a, b) is given by

$$\Pr[(a, b)] = \langle T_b, S_a \rangle.$$

In particular,

$$\Pr[(0, b)] - \Pr[(1, b)] = \langle T_b, S_0 - S_1 \rangle = \varepsilon \langle T_b, J(\Phi) \rangle.$$

Similarly, let $(\Psi, \{Q_a\})$ be an operational description of the strategy $\{S_a\}$ and let $(\rho, \{P_b\})$ be an operational description of the co-strategy $\{T_b\}$. These objects take the form

$$\begin{aligned} \Psi : \mathbf{L}(\mathcal{X}) &\rightarrow \mathbf{L}(\mathcal{Y} \otimes \mathcal{Z}) & \rho &\in \mathbf{H}^+(\mathcal{X} \otimes \mathcal{W}) \\ \{Q_a\} &\subset \mathbf{H}^+(\mathcal{Z}) & \{P_b\} &\subset \mathbf{H}^+(\mathcal{W}) \end{aligned}$$

for some choice of spaces \mathcal{W} and \mathcal{Z} . Basic quantum formalism tells us that the probability with which an interaction between $(\Psi, \{Q_a\})$ and $(\rho, \{P_b\})$ yields measurement outcomes (a, b) is given by

$$\Pr[(a, b)] = \langle Q_a \otimes P_b, (\Psi \otimes \mathbb{1}_{\mathcal{W}})(\rho) \rangle.$$

Let Φ^\pm be the super-operators with $J(\Phi^\pm) = R^\pm$, so that $\Phi = \Phi^+ - \Phi^-$. According to Definition 2.5 (New formalism for measuring strategies), the actions of Φ^\pm on any $X \in \mathbf{L}(\mathcal{X})$ are given by

$$\begin{aligned} \Phi^+ : X &\mapsto \frac{1}{\varepsilon} \text{Tr}_{\mathcal{Z}}((Q_0 \otimes I_{\mathcal{Y}}) \Psi(X)), \\ \Phi^- : X &\mapsto \frac{1}{\varepsilon} \text{Tr}_{\mathcal{Z}}((Q_1 \otimes I_{\mathcal{Y}}) \Psi(X)). \end{aligned}$$

In particular,

$$\begin{aligned} \Pr[(0, b)] &= \varepsilon \langle P_b, (\Phi^+ \otimes \mathbb{1}_{\mathcal{W}})(\rho) \rangle, \\ \Pr[(1, b)] &= \varepsilon \langle P_b, (\Phi^- \otimes \mathbb{1}_{\mathcal{W}})(\rho) \rangle. \end{aligned}$$

and hence

$$\Pr[(0, b)] - \Pr[(1, b)] = \varepsilon \langle P_b, (\Phi \otimes \mathbb{1}_{\mathcal{W}})(\rho) \rangle.$$

Then by linearity we obtain

$$\begin{aligned}
& \langle P_0 - P_1, (\Phi \otimes \mathbb{1}_{\mathcal{W}})(\rho) \rangle \\
&= \frac{1}{\varepsilon} ((\Pr[(0, 0)] - \Pr[(1, 0)]) - (\Pr[(0, 1)] - \Pr[(1, 1)])) \\
&= \langle T_0 - T_1, J(\Phi) \rangle.
\end{aligned}$$

To prove the proposition, first select the co-strategy $\{T_0, T_1\}$ so that its operational description $(\rho, \{P_a\})$ achieves the maximum in the definition of the diamond norm. In this case, the above expression implies $\|\Phi\|_{\diamond} \leq \|\Phi\|_{\diamond 1}$. For the reverse inequality, select a co-strategy $\{T_0, T_1\}$ that achieves the maximum in the definition of $\|\Phi\|_{\diamond 1}$. In this case, the above expression implies $\|\Phi\|_{\diamond} \geq \|\Phi\|_{\diamond 1}$. \square

Extension to non-Hermitian-preserving super-operators?

According to Definition 5.2 (Strategy r -norm), the norms $\|\Phi\|_{\diamond r}$ and $\|\Phi\|_{\diamond r}^*$ are defined only when Φ is a Hermitian-preserving super-operator. Is there a natural extension of these norms to *all* super-operators? Ideally, such an extension would agree with Definition 5.2 on Hermitian-preserving super-operators and with the diamond norm when $r = 1$.

What would such an extension look like? Presumably, it would require a generalization of quantum strategies as they appear in Definitions 2.1 and 2.4. For example, an r -tuple (Φ_1, \dots, Φ_r) might denote an operational representation of a generalized strategy if and only if each of the super-operators Φ_i preserves trace norm.

Indeed, a generalization of Theorem 3.1 (Interaction output probabilities) would probably be required in order to reason about these generalized strategies and define a generalized strategy r -norm. Moreover, in the next section we will see that several basic properties of the strategy r -norm are established via Proposition 3.8 (Polar sets of strategies) and ultimately Theorem 3.2 (Characterization of strategies). Presumably, a proper generalization of the strategy r -norm would also require corresponding generalizations of these two results. Essentially, much of the formalism developed so far in this thesis for quantum strategies would need to be re-derived in a more general setting.

What use could such a generalization find in quantum information theory? In looking to the diamond norm for inspiration, we find few examples of its use on non-Hermitian-preserving super-operators in a quantum information context. Perhaps the most notable such use is the “maximum output fidelity” characterization of the diamond norm:

Fact 5.5 (Maximum output fidelity characterization of the diamond norm). *Let $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ be an arbitrary super-operator. Let \mathcal{Z} be a space and let $A, B : \mathcal{X} \rightarrow \mathcal{Y} \otimes \mathcal{Z}$ be operators with $\Phi : X \mapsto \text{Tr}_{\mathcal{Z}}(AXB^*)$. Then*

$$\|\Phi\|_{\diamond} = \max_{\rho, \sigma} F(\Psi_A(\rho), \Psi_B(\sigma))$$

where the super-operators $\Psi_A, \Psi_B : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Z})$ are given by

$$\begin{aligned}\Psi_A &: X \mapsto \text{Tr}_{\mathcal{Y}}(AXA^*), \\ \Psi_B &: X \mapsto \text{Tr}_{\mathcal{Y}}(BXB^*)\end{aligned}$$

and $F(\rho, \sigma)$ denotes the fidelity function on pairs of quantum states and the maximum is taken over all quantum states $\rho, \sigma \in \mathbf{H}^+(\mathcal{X})$.

This characterization has been employed in the study of quantum interactive proofs [RW05] and a proof may be found in Ref. [KSV02]. Could this characterization be generalized so as to yield an interesting quantum information theoretic application of the strategy r -norm to non-Hermitian-preserving super-operators?

While an endeavor to generalize the strategy r -norm would no doubt be an interesting mathematical exercise, its applicability to quantum information is not readily apparent. In this thesis, our foray into the realm of mathematical curiosity of this flavor is limited to the study of the dual of the diamond norm in Section 5.3.

5.2 Distinguishing convex sets of strategies

In this section we generalize the guessing game example from the previous section from a problem of distinguishing *individual* states, operations, or strategies to distinguishing *convex sets* of states, operations, or strategies.

Specifically, suppose two convex sets $\mathbf{A}_0, \mathbf{A}_1$ of quantum states are fixed. Suppose that Bob arbitrarily selects $\rho_0 \in \mathbf{A}_0$ and $\rho_1 \in \mathbf{A}_1$ and then selects $\rho \in \{\rho_0, \rho_1\}$ uniformly at random and gives Alice a quantum system prepared in state ρ . Alice's goal is to correctly guess whether $\rho \in \mathbf{A}_0$ or $\rho \in \mathbf{A}_1$ based upon the outcome of a measurement she conducts on ρ . It is clear that this problem is a generalization of that from the previous section, as the original problem is recovered by considering the singleton sets $\mathbf{A}_0 = \{\rho_0\}$ and $\mathbf{A}_1 = \{\rho_1\}$.

This problem of distinguishing convex sets of states was solved in Ref. [GW05], wherein it was shown that there exists a *single* measurement $\{P_0, P_1\}$ that depends only upon the sets $\mathbf{A}_0, \mathbf{A}_1$ with the property that *any* pair $\rho_0 \in \mathbf{A}_0, \rho_1 \in \mathbf{A}_1$ may be correctly distinguished with probability at least

$$\frac{1}{2} + \frac{1}{4} \min_{\sigma_a \in \mathbf{A}_a} \|\sigma_0 - \sigma_1\|_{\text{Tr}}.$$

In particular, even if two distinct pairs ρ_0, ρ_1 and ρ'_0, ρ'_1 both minimize the trace distance between \mathbf{A}_0 and \mathbf{A}_1 then *both* pairs may be optimally distinguished by the *same* measurement $\{P_0, P_1\}$.

What about distinguishing convex sets of quantum operations or strategies? Nothing was known of either problem prior to the work of the present thesis. In this

section, we prove that the distinguishability result for convex sets of states extends unhindered to operations and strategies. In particular, we prove that two convex sets $\mathbf{S}_0, \mathbf{S}_1$ of r -round strategies can be correctly distinguished with probability at least

$$\frac{1}{2} + \frac{1}{4} \min_{S_a \in \mathbf{S}_a} \|S_0 - S_1\|_{\diamond r}.$$

It then follows trivially that two convex sets $\mathbf{T}_0, \mathbf{T}_1$ of r -round co-strategies can be correctly distinguished with probability at least

$$\frac{1}{2} + \frac{1}{4} \min_{T_a \in \mathbf{T}_a} \|T_0 - T_1\|_{\diamond r}^*.$$

As a special case, it holds that two convex sets Φ_0, Φ_1 of quantum operations can be distinguished with probability at least

$$\frac{1}{2} + \frac{1}{4} \min_{\Phi_a \in \Phi_a} \|\Phi_0 - \Phi_1\|_{\diamond}.$$

Properties of the strategy r -norm

Our proof of the distinguishability of convex sets of strategies is essentially a copy of the proof appearing in Ref. [GW05] with states and measurements replaced by strategies and co-strategies and the trace and operator norms replaced with the strategy r -norm and its dual.

To ensure correctness of the new proof, we must identify the relevant properties of the trace and operator norms employed in Ref. [GW05] and then establish suitable analogues of those properties for the new norms $\|\cdot\|_{\diamond r}$ and $\|\cdot\|_{\diamond r}^*$. The relevant properties of the trace and operator norms are:

1. The trace norm and operator norm are *dual* to each other, meaning that

$$\begin{aligned} \|X\|_{\text{Tr}} &= \max_{\|Y\| \leq 1} |\langle Y, X \rangle|, \\ \|X\| &= \max_{\|Y\|_{\text{Tr}} \leq 1} |\langle Y, X \rangle| \end{aligned}$$

for all operators X .

2. If X is Hermitian then

$$\|X\| \leq 1 \iff |X| \preceq I.$$

Here $|X| \stackrel{\text{def}}{=} X^+ + X^-$ where $X = X^+ - X^-$ is a Jordan decomposition of X .

Property 2 is really just a circuitous way of saying that a Hermitian operator has operator norm at most 1 if and only if all its eigenvalues are no larger than 1 in absolute value. Compared to the duality of the trace and operator norms (property

1), this observation is not deep or significant in and of itself. It is only phrased as such so as to highlight the forthcoming generalization to the strategy r -norm and its dual.

Our generalization of property 1 states that the norms $\|\cdot\|_{\diamond_r}$ and $\|\cdot\|_{\diamond_r}^*$ are dual to each other. Because our proof of this fact relies upon the generalization of property 2, we first establish the latter.

Unit ball of the strategy r -norms

Proposition 5.6 (Unit ball of the strategy r -norms). *For any Hermitian-preserving super-operator $\Phi : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$ it holds that*

$$\|\Phi\|_{\diamond_r} \leq 1 \iff |J(\Phi)| \preceq S$$

for some r -round non-measuring strategy S for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$. Similarly,

$$\|\Phi\|_{\diamond_r}^* \leq 1 \iff |J(\Phi)| \preceq T$$

for some r -round non-measuring co-strategy T for input spaces $\mathcal{X}_1, \dots, \mathcal{X}_r$ and output spaces $\mathcal{Y}_1, \dots, \mathcal{Y}_r$.

Proposition 5.6 follows immediately from the following two lemmas.

Lemma 5.7. *For any Hermitian-preserving super-operator $\Phi : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$ it holds that*

$$\begin{aligned} \|\Phi\|_{\diamond_r} \leq 1 &\iff \langle |J(\Phi)|, T \rangle \leq 1 \text{ for every } r\text{-round non-measuring co-strategy } T, \\ \|\Phi\|_{\diamond_r}^* \leq 1 &\iff \langle |J(\Phi)|, S \rangle \leq 1 \text{ for every } r\text{-round non-measuring strategy } S. \end{aligned}$$

Lemma 5.8. *For each positive semidefinite operator $X \in \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$ it holds that*

$$\begin{aligned} \langle X, T \rangle \leq 1 &\text{ for every } r\text{-round non-measuring co-strategy } T \\ \iff X \preceq S &\text{ for some } r\text{-round non-measuring strategy } S, \\ \langle X, S \rangle \leq 1 &\text{ for every } r\text{-round non-measuring strategy } S \\ \iff X \preceq T &\text{ for some } r\text{-round non-measuring co-strategy } T. \end{aligned}$$

Lemma 5.8 is an immediate corollary of Proposition 3.8 (Polar sets of strategies), which was proven in Section 3.3.1. Hence, we do not prove Lemma 5.8 here—it is restated only for ease of reference without the polarity notation. It remains only to prove Lemma 5.7.

Proof of Lemma 5.7 and hence also of Proposition 5.6. The proofs of the two desired equivalences are completely symmetric, so we prove only the first. As Φ is Hermitian-preserving, it holds that $J(\Phi)$ is a Hermitian operator. As such, it has a Jordan decomposition $J(\Phi) = R^+ - R^-$ with $|J(\Phi)| = R^+ + R^-$.

We first prove the easier implication. Suppose $\langle |J(\Phi)|, T \rangle \leq 1$ for all co-strategies T and choose any measuring co-strategy $\{T_0, T_1\}$. We must show that $\langle J(\Phi), T_0 - T_1 \rangle \leq 1$. As $\{T_0, T_1\}$ is a measuring co-strategy, it holds that $T_0 + T_1$ is a nonmeasuring co-strategy. Taking $T = T_0 + T_1$, we get

$$1 \geq \langle R^+ + R^-, T_0 + T_1 \rangle \geq \langle R^+ - R^-, T_0 - T_1 \rangle$$

as desired.

For the other direction, suppose $\|\Phi\|_{\diamond r} \leq 1$ and choose any co-strategy T . We must show that $\langle |J(\Phi)|, T \rangle \leq 1$. Our goal is to decompose $T = T_0 + T_1$ in such a way that $\{T_0, T_1\}$ is a measuring co-strategy with $\langle R^+, T_1 \rangle = \langle R^-, T_0 \rangle = 0$. Such a choice of T_0, T_1 ensures that

$$\langle R^+ - R^-, T_0 - T_1 \rangle = \langle R^+ + R^-, T_0 + T_1 \rangle.$$

The lemma is then established by noting that the left side of this expression is at most 1 and the right side of this expression equals $\langle |J(\Phi)|, T \rangle$.

It remains to choose the appropriate T_0, T_1 . Letting Π^+ denote the projection onto the support of R^+ , we choose

$$\begin{aligned} T_0 &= \Pi^+ T \Pi^+, \\ T_1 &= T - T_0. \end{aligned}$$

That $\langle R^+, T_1 \rangle = \langle R^-, T_0 \rangle = 0$ follows from the fact that R^+ and R^- have orthogonal support. Finally, it is easy to verify that $\{T_0, T_1\}$ denotes a valid measuring co-strategy: as T is a co-strategy, so too is the sum $T_0 + T_1 = T$. As T is positive semidefinite, so too are T_0, T_1 . \square

Duality of the strategy r -norms

We are now ready to prove that the norms $\|\cdot\|_{\diamond r}$ and $\|\cdot\|_{\diamond r}^*$ are dual to each other.

Proposition 5.9 (Duality of the strategy r -norms). *The norms $\|\cdot\|_{\diamond r}$ and $\|\cdot\|_{\diamond r}^*$ are dual to each other. In other words, for any Hermitian-preserving super-operator $\Phi : \mathbf{L}(\mathcal{X}_{1\dots r}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots r})$ it holds that*

$$\begin{aligned} \|\Phi\|_{\diamond r} &= \max_{\|\Psi\|_{\diamond r}^* \leq 1} \langle \Psi, \Phi \rangle, \\ \|\Phi\|_{\diamond r}^* &= \max_{\|\Psi\|_{\diamond r} \leq 1} \langle \Psi, \Phi \rangle. \end{aligned}$$

Proof. The proofs of the two desired equalities are completely symmetric, so we prove only the first. We begin by proving

$$\|\Phi\|_{\diamond r} \leq \max_{\|\Psi\|_{\diamond r}^* \leq 1} \langle \Psi, \Phi \rangle.$$

Let $\{T_0, T_1\}$ be an r -round measuring co-strategy attaining the maximum in the definition of $\|\Phi\|_{\diamond r}$, so that

$$\|\Phi\|_{\diamond r} = \langle T_0 - T_1, J(\Phi) \rangle.$$

The desired inequality follows from the claim that $\|T_0 - T_1\|_{\diamond r}^* \leq 1$. To verify this claim, we note that any r -round measuring strategy $\{S_0, S_1\}$ attaining the maximum in the definition of $\|T_0 - T_1\|_{\diamond r}^*$ has

$$\|T_0 - T_1\|_{\diamond r}^* = \langle S_0 - S_1, T_0 - T_1 \rangle \leq \langle S_0 + S_1, T_0 + T_1 \rangle = 1.$$

That the final inner product equals one follows immediately from Theorem 3.1 (Interaction output probabilities) and the observation that $S_0 + S_1$ may be viewed as a one-outcome measuring strategy and $T_0 + T_1$ may be viewed as a one-outcome measuring co-strategy.

For the reverse inequality, choose a Hermitian-preserving super-operator Ψ with $\|\Psi\|_{\diamond r}^* \leq 1$ that maximizes the inner product $\langle \Psi, \Phi \rangle$. Proposition 5.6 tells us that $|J(\Psi)| \preceq T$ for some r -round non-measuring co-strategy T . Let $J(\Psi) = R^+ - R^-$ be the Jordan decomposition of $J(\Psi)$ and let $\{T_0, T_1\}$ be the r -round measuring co-strategy given by

$$\begin{aligned} T_0 &= R^+ + \frac{1}{2}(T - |J(\Psi)|), \\ T_1 &= R^- + \frac{1}{2}(T - |J(\Psi)|). \end{aligned}$$

Then

$$\langle \Psi, \Phi \rangle = \langle T_0 - T_1, J(\Phi) \rangle \leq \|\Phi\|_{\diamond r}$$

as desired. □

Distinguishability of convex sets of strategies

As mentioned earlier, our proof of the distinguishability of convex sets of strategies closely resembles the proof of the distinguishability of convex sets of states appearing in Ref. [GW05]. Now that we have identified and established the relevant properties of the norm $\|\cdot\|_{\diamond r}$ and its dual $\|\cdot\|_{\diamond r}^*$, we are ready to translate the proof of Ref. [GW05] onto the domain of quantum strategies.

Theorem 5.10 (Distinguishability of convex sets of strategies). *Let $\mathbf{S}_0, \mathbf{S}_1 \subset \mathbf{H}^+(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r})$ be nonempty convex sets of r -round strategies. There exists an r -round measuring co-strategy $\{T_0, T_1\}$ with the property that*

$$\langle T_0 - T_1, S_0 - S_1 \rangle \geq \min_{R_a \in \mathbf{S}_a} \|R_0 - R_1\|_{\diamond r}$$

for all choices of $S_0 \in \mathbf{S}_0$ and $S_1 \in \mathbf{S}_1$. A similar statement holds in terms of the dual norm $\|\cdot\|_{\diamond r}^*$ for convex sets of co-strategies.

Proof. The proof for co-strategies is completely symmetric to the proof for strategies, so we only address strategies here. Let d denote the minimum distance between \mathbf{S}_0 and \mathbf{S}_1 as stated in the theorem. If $d = 0$ then the theorem is satisfied by the trivial r -round measuring co-strategy corresponding to a random coin flip. (For this trivial co-strategy, both T_0 and T_1 are equal to the identity divided by $2 \dim(\mathcal{X}_{1\dots r})$.) For the remainder of this proof, we shall restrict our attention to the case $d > 0$.

Define

$$\mathbf{S} \stackrel{\text{def}}{=} \mathbf{S}_0 - \mathbf{S}_1 = \{S_0 - S_1 : S_0 \in \mathbf{S}_0, S_1 \in \mathbf{S}_1\}$$

and let

$$\mathbf{B} \stackrel{\text{def}}{=} \{B \in \mathbf{H}(\mathcal{Y}_{1\dots r} \otimes \mathcal{X}_{1\dots r}) : \|B\|_{\diamond r} < d\}$$

denote the open ball of radius d with respect to the $\|\cdot\|_{\diamond r}$ norm. The sets \mathbf{S} and \mathbf{B} are nonempty disjoint sets of Hermitian operators, both are convex, and \mathbf{B} is open. By the Separation Theorem (Fact 1.6), there exists a Hermitian operator H and a scalar α such that

$$\langle H, S \rangle \geq \alpha > \langle H, B \rangle$$

for all $S \in \mathbf{S}$ and $B \in \mathbf{B}$.

For every choice of $B \in \mathbf{B}$ we have $-B \in \mathbf{B}$ as well, from which it follows that $|\langle H, B \rangle| < \alpha$ for all $B \in \mathbf{B}$ and hence $\alpha > 0$. Moreover, as \mathbf{B} is the open ball of radius d in the norm $\|\cdot\|_{\diamond r}$, it follows from Proposition 5.9 (Duality of the strategy r -norms) that

$$\|H\|_{\diamond r}^* \leq \alpha/d.$$

Now let $\hat{H} = \frac{d}{\alpha}H$ be the normalization of H with $\|\hat{H}\|_{\diamond r}^* \leq 1$ and let $\hat{H} = \hat{H}^+ - \hat{H}^-$ be the Jordan decomposition of \hat{H} . By Proposition 5.6 (Unit ball of the strategy r -norms) we have $|\hat{H}| \preceq T$ for some co-strategy T . Let $\{T_0, T_1\}$ be the measuring co-strategy given by

$$\begin{aligned} T_0 &= \hat{H}^+ + \frac{1}{2} \left(T - |\hat{H}| \right), \\ T_1 &= \hat{H}^- + \frac{1}{2} \left(T - |\hat{H}| \right). \end{aligned}$$

It remains only to verify that $\{T_0, T_1\}$ has the desired properties: for every choice of $S_0 \in \mathbf{S}_0$ and $S_1 \in \mathbf{S}_1$ we have

$$\langle T_0 - T_1, S_0 - S_1 \rangle = \langle \hat{H}, S_0 - S_1 \rangle = \frac{d}{\alpha} \langle H, S_0 - S_1 \rangle \geq d$$

as desired. □

The claimed result regarding the distinguishability of convex sets of strategies now follows immediately. To recap, let $\mathbf{S}_0, \mathbf{S}_1$ be convex sets of strategies and let

$\{T_0, T_1\}$ denote the measuring co-strategy from Theorem 5.10 that distinguishes elements in \mathbf{S}_0 from elements in \mathbf{S}_1 . Suppose Bob selects $S_0 \in \mathbf{S}_0$ and $S_1 \in \mathbf{S}_1$ arbitrarily and then selects $S \in \{S_0, S_1\}$ uniformly at random. As derived in Section 5.1, if Alice acts according to $\{T_0, T_1\}$ then the probability with which she correctly guesses whether $S \in \mathbf{S}_0$ or $S \in \mathbf{S}_1$ is given by

$$\frac{1}{2} + \frac{1}{4} \langle T_0 - T_1, S_0 - S_1 \rangle \geq \frac{1}{2} + \frac{1}{4} \min_{R_a \in \mathbf{S}_a} \|R_0 - R_1\|_{\diamond r}$$

as desired.

5.3 The dual of the diamond norm

In Proposition 5.4 we established that the strategy 1-norm $\|\Phi\|_{\diamond 1}$ agrees with the diamond norm $\|\Phi\|_{\diamond}$ on Hermitian-preserving super-operators Φ . Can the same be said of the duals of these norms? In this section we answer that question in the affirmative. In order to do so, we also establish several basic facts concerning the dual of the diamond norm. While none of these facts are surprising, the only proofs we can offer are nontrivial and possibly interesting in their own right. Thus, in this section we take a “detour” from the strategy r -norm in order to study the dual of the diamond norm, returning only at the very end to answer this section’s opening question.

Definition 5.11 (Dual of the diamond norm). As with any norm, the *dual* $\|\cdot\|_{\diamond}^*$ of the diamond norm $\|\cdot\|_{\diamond}$ is defined for every super-operator Φ as

$$\|\Phi\|_{\diamond}^* \stackrel{\text{def}}{=} \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle|.$$

□

That the dual of the dual of the diamond norm equals the diamond norm follows from the Duality Theorem, a proof of which can be found in Horn and Johnson [HJ85]. In other words, it holds that

$$\|\Phi\|_{\diamond}^{**} = \|\Phi\|_{\diamond} = \max_{\|\Psi\|_{\diamond}^*=1} |\langle \Psi, \Phi \rangle|.$$

While much is known of the diamond norm, its dual has never been studied. In this section, we establish the following basic facts about this norm:

1. If Φ is completely positive then the maximum in the definition of $\|\Phi\|_{\diamond}^*$ is achieved by a completely positive and trace-preserving super-operator.
2. If Φ is Hermitian-preserving then the maximum in the definition of $\|\Phi\|_{\diamond}^*$ is achieved by a Hermitian-preserving super-operator.

3. For each of the two maxima

$$\|\Phi\|_{\diamond}^* = \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle|, \quad \|\Phi\|_{\diamond} = \max_{\|\Psi\|_{\diamond}^*=1} |\langle \Psi, \Phi \rangle|$$

there exist super-operators Φ such that the maximum is not attained by any Hermitian-preserving super-operator Ψ . Moreover, there exist Hermitian-preserving super-operators Φ such that the maximum is not attained by any completely positive super-operator Ψ .

4. If Φ is Hermitian-preserving then $\|\Phi\|_{\diamond}^* = \|\Phi\|_{\diamond 1}^*$.

Useful lemmas involving the diamond norm

We begin with three technical lemmas, two of which establish facts about the diamond norm. The first lemma follows immediately from Refs. [RW05, AHW00]. We provide an alternate proof for completeness.

Lemma 5.12. *For any completely positive super-operator Φ it holds that $\|\Phi\|_{\text{Tr}} = \|\Phi\|_{\diamond}$. Moreover, the maximum in the definition of $\|\Phi\|_{\text{Tr}}$ is achieved by a positive semidefinite operator.*

Proof. It is clear that $\|\Phi\|_{\diamond} \geq \|\Phi\|_{\text{Tr}}$, so let us concentrate only on the reverse inequality. Suppose Φ has the form $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$. As Φ is completely positive, there exists a space \mathcal{W} and a density operator $\rho \in \mathbf{H}^+(\mathcal{X} \otimes \mathcal{W})$ with the property that

$$\|\Phi\|_{\diamond} = \|(\Phi \otimes \mathbb{1}_{\mathcal{W}})(\rho)\|_{\text{Tr}} = \text{Tr}((\Phi \otimes \mathbb{1}_{\mathcal{W}})(\rho)).$$

By convexity, we may assume that ρ is a pure state—that is, $\rho = uu^*$ for some unit vector $u \in \mathcal{X} \otimes \mathcal{W}$. For any orthonormal basis $\{e_1, \dots, e_{\dim(\mathcal{W})}\}$ of \mathcal{W} we may write

$$u = \sum_{i=1}^{\dim(\mathcal{W})} \sqrt{q_i} x_i \otimes e_i$$

for some choice of unit vectors $x_1, \dots, x_{\dim(\mathcal{W})} \in \mathcal{X}$ (not necessarily orthogonal) and nonnegative real numbers $q_1, \dots, q_{\dim(\mathcal{W})}$ that sum to one. Then

$$\|\Phi\|_{\diamond} = \sum_{i,j=1}^{\dim(\mathcal{W})} \sqrt{q_i q_j} \text{Tr}(\Phi(x_i x_j^*) \otimes e_i e_j^*) = \sum_{i=1}^{\dim(\mathcal{W})} q_i \text{Tr}(\Phi(x_i x_i^*)) = \text{Tr}(\Phi(\sigma))$$

for $\sigma = \sum_{i=1}^{\dim(\mathcal{W})} q_i x_i x_i^*$. □

For any operator $A : \mathcal{X} \rightarrow \mathcal{Y}$ of rank r it follows from the Singular Value Theorem (Section 1.2.1) that it is possible to choose vectors $u_1, \dots, u_r \in \mathcal{Y}$, $v_1, \dots, v_r \in \mathcal{X}$ such that

$$A = \sum_{i=1}^r u_i v_i^*.$$

Given such a choice of vectors, we define the positive semidefinite operators $A_L \in \mathbf{H}^+(\mathcal{Y})$, $A_R \in \mathbf{H}^+(\mathcal{X})$ by

$$A_L = \sum_{i=1}^r u_i u_i^*, \quad A_R = \sum_{i=1}^r v_i v_i^*.$$

If A is Hermitian then it is clear that we may take $A_L = A_R = |A|$. As a consequence, if A is positive semidefinite then we may take $A_L = A_R = A$. The following simple lemma is a special case of Ref. [Wat05, Lemma 2].

Lemma 5.13. *For any operators $A, B : \mathcal{X} \rightarrow \mathcal{Y}$ and any decompositions A_L, A_R and B_L, B_R of those operators it holds that*

$$|\langle A, B \rangle|^2 \leq \langle A_L, B_L \rangle \cdot \langle A_R, B_R \rangle.$$

Proof. Let

$$A = \sum_{i=1}^r u_i v_i^*, \quad B = \sum_{j=1}^q w_j x_j^*$$

be decompositions of A, B that yield A_L, A_R and B_L, B_R . Then

$$\begin{aligned} |\langle A, B \rangle| &= \left| \sum_{i=1}^r \sum_{j=1}^q \langle u_i v_i^*, w_j x_j^* \rangle \right| = \left| \sum_{i=1}^r \sum_{j=1}^q \langle u_i, w_j \rangle \cdot \langle x_j, v_i \rangle \right| \\ &\leq \sqrt{\sum_{i=1}^r \sum_{j=1}^q |\langle u_i, w_j \rangle|^2} \cdot \sqrt{\sum_{i=1}^r \sum_{j=1}^q |\langle x_j, v_i \rangle|^2} \\ &= \sqrt{\sum_{i=1}^r \sum_{j=1}^q \langle u_i u_i^*, w_j w_j^* \rangle} \cdot \sqrt{\sum_{i=1}^r \sum_{j=1}^q \langle v_i v_i^*, x_j x_j^* \rangle} \\ &= \sqrt{\langle A_L, B_L \rangle} \cdot \sqrt{\langle A_R, B_R \rangle} \end{aligned}$$

where the inequality follows from Cauchy-Schwarz. \square

For any super-operator Φ and any decomposition $J(\Phi)_L, J(\Phi)_R$ of $J(\Phi)$ we let Φ_L, Φ_R denote the super-operators with $J(\Phi_L) = J(\Phi)_L$ and $J(\Phi_R) = J(\Phi)_R$. As $J(\Phi_L), J(\Phi_R)$ are positive semidefinite, it holds that Φ_L, Φ_R are completely positive.

Lemma 5.14. *For any super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ there exists a decomposition Φ_L, Φ_R of Φ with*

$$\|\Phi\|_\diamond = \|\Phi_L\|_\diamond = \|\Phi_R\|_\diamond.$$

Proof. It was noted in the conclusion of Ref. [Wat05] that

$$\|\Phi\|_\diamond^2 = \inf \{ \|\Phi_L\|_{\text{Tr}} \cdot \|\Phi_R\|_{\text{Tr}} \}$$

where the infimum is taken over all decompositions Φ_L, Φ_R of Φ . The existence of a fixed pair (Φ_L, Φ_R) that achieve this infimum may be argued as in Ref. [KSV02]. By rescaling, we may assume that $\|\Phi_L\|_{\text{Tr}} = \|\Phi_R\|_{\text{Tr}}$, from which it follows that

$$\|\Phi\|_{\diamond} = \|\Phi_L\|_{\text{Tr}} = \|\Phi_R\|_{\text{Tr}}.$$

The lemma then follows from Lemma 5.12. \square

Achieving the maximum

We are now ready to exhibit two theorems that establish two of the claims from the beginning of this section.

For any Hermitian-preserving super-operator Φ , we let $|\Phi|$ denote the completely positive super-operator with $J(|\Phi|) = |J(\Phi)|$. Just as with operators, there is a decomposition Φ_L, Φ_R of Φ with $\Phi_L = \Phi_R = |\Phi|$. Moreover, if Φ is completely positive then we may take $\Phi_L = \Phi_R = \Phi$.

Theorem 5.15 (Achieving the maximum for completely positive super-operators). *For any Hermitian-preserving super-operator Φ it holds that*

$$\|\Phi\|_{\diamond}^* \leq \||\Phi|\|_{\diamond}^*.$$

Moreover, the maximum in the definition of $\||\Phi|\|_{\diamond}^$ is achieved by a completely positive and trace-preserving super-operator.*

As a consequence, if Φ is completely positive then the maximum in the definition of $\|\Phi\|_{\diamond}^$ is achieved by a completely positive and trace-preserving super-operator.*

Proof. First we show that the maximum is attained by a completely positive super-operator—the trace-preserving property will be established later.

Let Ψ be any super-operator with $\|\Psi\|_{\diamond} = 1$ achieving the maximum in the definition of $\|\Phi\|_{\diamond}^*$, so that $\|\Phi\|_{\diamond}^* = |\langle \Psi, \Phi \rangle|$. By Lemma 5.14 there exists a decomposition Ψ_L, Ψ_R of Ψ with

$$\|\Psi\|_{\diamond} = \|\Psi_L\|_{\diamond} = \|\Psi_R\|_{\diamond} = 1.$$

Moreover, Ψ_L and Ψ_R are completely positive.

As Φ is Hermitian-preserving, there is a decomposition Φ_L, Φ_R of Φ with $\Phi_L = \Phi_R = |\Phi|$. By Lemma 5.13 we have

$$\begin{aligned} \|\Phi\|_{\diamond}^* &= |\langle \Psi, \Phi \rangle| \leq \sqrt{\langle \Psi_L, |\Phi| \rangle} \cdot \sqrt{\langle \Psi_R, |\Phi| \rangle} \\ &\leq \max \{ \langle \Xi, |\Phi| \rangle : \|\Xi\|_{\diamond} = 1 \text{ and } \Xi \text{ is completely positive} \}. \end{aligned}$$

Thus, $\|\Phi\|_{\diamond}^* \leq \||\Phi|\|_{\diamond}^*$ and the maximum in the definition of $\||\Phi|\|_{\diamond}^*$ is attained by a completely positive super-operator Ξ .

For the trace-preserving property, we note that $1 = \|\Xi\|_{\diamond} = \|\Xi\|_{\diamond 1}$. Then by Proposition 5.6 and the complete positivity of Ξ it holds that $J(\Xi) \preceq J(\Xi')$ for some completely positive and trace-preserving super-operator Ξ' with $\|\Xi'\|_{\diamond} = 1$. The desired result follows from the observation that $\langle \Xi', |\Phi| \rangle \geq \langle \Xi, |\Phi| \rangle$. \square

Theorem 5.16 (Achieving the maximum for Hermitian-preserving super-operators). *For any Hermitian-preserving super-operator Φ the maximum in the definition of $\|\Phi\|_{\diamond}^*$ is achieved by a Hermitian-preserving super-operator.*

Proof. Let $J(\Phi) = T^+ - T^-$ be a Jordan decomposition of $J(\Phi)$, so that $J(|\Phi|) = T^+ + T^-$. By Theorem 5.15 it holds that $\|\Phi\|_{\diamond}^* \leq \||\Phi|\|_{\diamond}^*$ and the maximum in the definition of $\||\Phi|\|_{\diamond}^*$ is achieved by a completely positive and trace-preserving super-operator Ψ .

Let Π^{\pm} denote the projections onto the support of T^{\pm} and let Ξ denote the Hermitian-preserving super-operator with

$$J(\Xi) = \Pi^+ J(\Psi) \Pi^+ - \Pi^- J(\Psi) \Pi^-.$$

It is easily verified that

$$\langle \Psi, |\Phi| \rangle = \langle \Xi, \Phi \rangle.$$

The left side of this equality is $\||\Phi|\|_{\diamond}^*$, which we know to be at least as large as $\|\Phi\|_{\diamond}^*$. Hence, the desired result will follow once we establish that the Hermitian-preserving super-operator Ξ has $\|\Xi\|_{\diamond} \leq 1$.

Toward that end, we note that

$$|J(\Xi)| = \Pi^+ J(\Psi) \Pi^+ + \Pi^- J(\Psi) \Pi^- \preceq J(\Psi).$$

As Ψ is completely positive and trace-preserving, it holds that $J(\Psi)$ denotes a one-round non-measuring strategy. Thus, by Proposition 5.6 it holds that $\|\Xi\|_{\diamond 1} \leq 1$. The theorem follows from Proposition 5.4, which tells us that $\|\Xi\|_{\diamond} = \|\Xi\|_{\diamond 1}$. \square

Not achieving the maximum

We now show that Theorems 5.15 and 5.16 cannot be extended beyond completely positive and Hermitian-preserving super-operators, respectively. Our counterexamples rely upon the following lemma.

Lemma 5.17 (Diamond norm for 1-dimensional spaces). *The following hold for any super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$:*

1. *Suppose $\dim(\mathcal{X}) = 1$ and let $A \in \mathbf{L}(\mathcal{Y})$ be the operator with $\Phi : \alpha \mapsto \alpha A$. It holds that $\|\Phi\|_{\diamond} = \|\Phi\|_{\text{Tr}} = \|A\|_{\text{Tr}}$.*
2. *Suppose $\dim(\mathcal{Y}) = 1$ and let $A \in \mathbf{L}(\mathcal{X})$ be the operator with $\Phi : X \mapsto \langle A, X \rangle$. It holds that $\|\Phi\|_{\diamond} = \|\Phi\|_{\text{Tr}} = \|A\|$.*

Proof. Item 1 is a simple consequence of the fact that the auxiliary space \mathcal{W} in the definition of the diamond norm can be assumed to have dimension no larger than $\dim(\mathcal{X})$. This fact immediately implies $\|\Phi\|_{\diamond} = \|\Phi\|_{\text{Tr}}$. It follows immediately from the definition of the super-operator trace norm that $\|\Phi\|_{\text{Tr}} = \|A\|_{\text{Tr}}$.

For item 2, choose a space \mathcal{W} and vectors $u, v \in \mathcal{X} \otimes \mathcal{W}$ with

$$\|\Phi\|_{\diamond} = \|(\Phi \otimes \mathbb{1}_{\mathcal{W}})(uv^*)\|_{\text{Tr}}.$$

Let $U \in \mathbf{L}(\mathcal{W})$ be a unitary such that

$$\|(\Phi \otimes \mathbb{1}_{\mathcal{W}})(uv^*)\|_{\text{Tr}} = \text{Tr}(U(\Phi \otimes \mathbb{1}_{\mathcal{W}})(uv^*))$$

and let $w \in \mathcal{X} \otimes \mathcal{W}$ be the vector given by $w = (I_{\mathcal{X}} \otimes U)u$, so that

$$\|(\Phi \otimes \mathbb{1}_{\mathcal{W}})(uv^*)\|_{\text{Tr}} = \text{Tr}((\Phi \otimes \mathbb{1}_{\mathcal{W}})(wv^*)).$$

As in the proof of Lemma 5.12, choose an orthonormal basis $\{e_i\}$ of \mathcal{W} and write

$$w = \sum_{i=1}^{\dim(\mathcal{W})} \sqrt{q_i} w_i \otimes e_i, \quad v = \sum_{i=1}^{\dim(\mathcal{W})} \sqrt{p_i} v_i \otimes e_i$$

for some choice of unit vectors $\{w_i\}, \{v_i\} \subset \mathcal{X}$ and probability distributions $\{q_i\}, \{p_i\}$. Then

$$\text{Tr}((\Phi \otimes \mathbb{1}_{\mathcal{W}})(wv^*)) = \sum_{i=1}^{\dim(\mathcal{W})} \sqrt{p_i} \sqrt{q_i} \langle A, w_i v_i \rangle \leq \|A\| \sum_{i=1}^{\dim(\mathcal{W})} \sqrt{p_i} \sqrt{q_i} \leq \|A\|$$

where the final inequality is Cauchy-Schwarz. Equality is obtained for real numbers $p_1 = q_1 = 1$ and unit vectors w_1, v_1 maximizing the inner product $\langle A, w_1 v_1 \rangle$. We have thus established $\|\Phi\|_{\diamond} = \|A\|$. As $\|\Phi(w_1 v_1^*)\|_{\text{Tr}} = \|A\|$, we also have $\|\Phi\|_{\diamond} = \|\Phi\|_{\text{Tr}}$. \square

Proposition 5.18 (Theorems 5.15 and 5.16 do not extend). *For each of the two maxima*

$$\|\Phi\|_{\diamond}^* = \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle|, \quad \|\Phi\|_{\diamond} = \max_{\|\Psi\|_{\diamond}^*=1} |\langle \Psi, \Phi \rangle|$$

there exist super-operators Φ such that the maximum is not attained by any Hermitian-preserving super-operator Ψ .

Moreover, there exist Hermitian-preserving super-operators Φ such that the maximum is not attained by any completely positive super-operator Ψ .

Proof. The counterexamples presented here are all achieved via reduction from the diamond norm and its dual to the trace norm and its dual (the operator norm).

Let \mathcal{X} be a space of dimension one, so that for each super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ there is an operator $A \in \mathbf{L}(\mathcal{Y})$ with $\Phi : \alpha \mapsto \alpha A$. For any $\Psi : \alpha \mapsto \alpha B$ it is easily verified that

$$|\langle \Psi, \Phi \rangle| = |\langle B, A \rangle|.$$

Moreover, it is clear that Ψ is Hermitian-preserving if and only if B is Hermitian and that Ψ is completely positive if and only if B is positive semidefinite. By Lemma 5.17 it holds that

$$\|\Phi\|_{\diamond}^* = \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle| = \max_{\|B\|_{\text{Tr}}=1} |\langle B, A \rangle| = \|A\|. \quad (5.1)$$

We now exhibit an operator A such that the maximum is not achieved by any Hermitian operator B . Consider the operator

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

The quantity

$$\max_{\substack{\|B\|_{\text{Tr}}=1, \\ B \text{ Hermitian}}} |\langle B, A \rangle| = \max_{\|v\|=1} |v^* A v|$$

is known as the *numerical radius* $\varrho(A)$ of A . It is easy to compute $\varrho(A) = \frac{3}{2}$, which is strictly smaller than $\|A\| = \frac{1+\sqrt{5}}{2}$.

Next, let \mathcal{Y} be a space of dimension one, so that for each super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ there is an operator $A \in \mathbf{L}(\mathcal{Y})$ with $\Phi : X \mapsto \langle A, X \rangle$. For any $\Psi : X \mapsto \langle B, X \rangle$ it is easily verified that

$$|\langle \Psi, \Phi \rangle| = |\langle B, A \rangle|.$$

Moreover, it is clear that Ψ is Hermitian-preserving if and only if B is Hermitian and that Ψ is completely positive if and only if B is positive semidefinite. By Lemma 5.17 it holds that

$$\|\Phi\|_{\diamond}^* = \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle| = \max_{\|B\|=1} |\langle B, A \rangle| = \|A\|_{\text{Tr}}. \quad (5.2)$$

Consider the operator

$$A = \Pi_0 - \Pi_1$$

where $\Pi_0, \Pi_1 \in \mathbf{H}^+(\mathcal{Y})$ are nonzero orthogonal projections. It is easy to verify that the maximum is not achieved by any positive semidefinite B .

The proposition is now proved for the maximum $\|\Phi\|_{\diamond}^* = \max_{\|\Psi\|_{\diamond}=1} |\langle \Psi, \Phi \rangle|$. The proof for the other maximum follows along similar lines. \square

The dual of the strategy 1-norm agrees with the dual of the diamond norm

Finally, we have what we need to establish agreement between $\|\cdot\|_{\diamond 1}^*$ and $\|\cdot\|_{\diamond}^*$ for Hermitian-preserving super-operators.

Proposition 5.19 (Agreement with the dual of the diamond norm). *For every Hermitian-preserving super-operator Φ it holds that $\|\Phi\|_{\diamond}^* = \|\Phi\|_{\diamond 1}^*$.*

Proof. We have

$$\begin{aligned} \|\Phi\|_{\diamond}^* &= \max \{ |\langle \Psi, \Phi \rangle| : \|\Psi\|_{\diamond} = 1 \text{ and } \Psi \text{ is Hermitian-preserving} \} \\ &= \max \{ |\langle \Psi, \Phi \rangle| : \|\Psi\|_{\diamond 1} = 1 \text{ and } \Psi \text{ is Hermitian-preserving} \} \\ &= \|\Phi\|_{\diamond 1}^*. \end{aligned}$$

The first equality is Theorem 5.16, the second Proposition 5.4, and the third Proposition 5.9. \square

Part II

Local Operations with Shared Entanglement

Chapter 6

Introduction to Local Operations

In Part II of this thesis we are interested in various classes of “local” quantum operations, which are operations that can be jointly implemented by two or more parties who act on distinct portions of the input system and who do not communicate once they receive their portions of the input.

In this introductory chapter for Part II we provide formal definitions and immediate observations for three such classes: local operations, local operations with shared randomness, and local operations with shared entanglement. We also introduce convenient shorthand notations for separable operators and for the vector space spanned by the local operations—objects which are of fundamental importance to the work in this part of the thesis.

Local operations and the space that they span

The simplest examples of quantum operations that can be implemented jointly by multiple parties without communication are the product operations, which consist of several completely independent quantum operations juxtaposed and viewed as a larger single operation.

Definition 6.1 (Local operation). A quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ is *m-party local* with respect to the input partition $\mathcal{X}_1, \dots, \mathcal{X}_m$ and output partition $\mathcal{Y}_1, \dots, \mathcal{Y}_m$ if it can be written as a product operation of the form

$$\Lambda = \Psi_1 \otimes \dots \otimes \Psi_m$$

where each $\Psi_i : \mathbf{L}(\mathcal{X}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$ is a quantum operation. Typically, the partition of the input and output spaces is implicit and clear from the context. \square

The vector space of super-operators spanned by the local operations is of paramount interest in Part II of this thesis. As such, we introduce a convenient shorthand notation for this space.

Toward that end, recall that a super-operator $\Psi_i : \mathbf{L}(\mathcal{X}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$ denotes a quantum operation if and only if Ψ_i is completely positive and trace-preserving. In terms of Choi-Jamiołkowski representations, the trace-preserving property of Ψ_i is characterized by the condition $\text{Tr}_{\mathcal{Y}_i}(J(\Psi_i)) = I_{\mathcal{X}_i}$. The set of all operators X obeying the inhomogeneous linear condition $\text{Tr}_{\mathcal{Y}_i}(X) = I_{\mathcal{X}_i}$ is not a vector space, but this set is easily extended to a unique smallest vector space by including its closure under multiplication by real scalars.

Definition 6.2 (Shorthand notation for the space spanned by local operations). For complex Euclidean spaces $\mathcal{X}_i, \mathcal{Y}_i$, let

$$\mathbf{Q}_i \stackrel{\text{def}}{=} \{X \in \mathbf{H}(\mathcal{Y}_i \otimes \mathcal{X}_i) : \text{Tr}_{\mathcal{Y}_i}(X) = \lambda I_{\mathcal{X}_i} \text{ for some } \lambda \in \mathbb{R}\}$$

denote the subspace of the real vector space $\mathbf{H}(\mathcal{Y}_i \otimes \mathcal{X}_i)$ of Hermitian operators X of the form $X = J(\Psi)$ for which $\Psi : \mathbf{L}(\mathcal{X}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$ is a trace-preserving super-operator, or a scalar multiple thereof. (Throughout this thesis, the spaces $\mathcal{X}_i, \mathcal{Y}_i$ are implicit whenever the notation \mathbf{Q}_i is used.)

Let $\mathcal{X}_1, \dots, \mathcal{X}_m$ and $\mathcal{Y}_1, \dots, \mathcal{Y}_m$ be complex Euclidean spaces. Employing our shorthand notation for Kronecker products, the subspace $\mathbf{Q}_{1\dots m}$ of $\mathbf{H}(\mathcal{Y}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$ spanned by the m -party local operations is given by

$$\mathbf{Q}_{1\dots m} = \text{span} \{X_1 \otimes \dots \otimes X_m : X_i \in \mathbf{Q}_i\}.$$

□

By the end of Part II, we will see that the space $\mathbf{Q}_{1\dots m}$ contains not only all the local operations, but also the local operations with shared randomness and with shared entanglement, as well as an even broader class of local quantum operations called “no-signaling” operations. Moreover, in Chapter 10 we will also see a converse result—that any quantum operation whose Choi-Jamiołkowski representation lies inside $\mathbf{Q}_{1\dots m}$ is necessarily a no-signaling operation.

Local operations with shared randomness, separable operators

Local (product) operations are not the only quantum operations admitted by our model. In particular, there is nothing to stop the parties from meeting ahead of time so as to prepare shared resources that might allow them to correlate their separate quantum operations.

Consider, for example, a convex combination $\sum_j p_j \Lambda_j$ of m -party local operations Λ_j . Such a quantum operation is legal in our model because it can be implemented by m parties who share prior knowledge of an integer j sampled according to the probability distribution p_j . In this case, the shared resource is randomness.

Definition 6.3 (Local operation with shared randomness (LOSR)). A quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ is a m -party LOSR operation if it can be written as a convex combination of m -party local operations. □

As a convex combination of products of positive semidefinite operators, the Choi-Jamiołkowski representation of a LOSR operation belongs to a broader class of operators called “separable” operators.

Definition 6.4 (Separable operator). Let $\mathbf{S}_1, \dots, \mathbf{S}_m$ be arbitrary spaces of Hermitian operators. An element X of the product space $\mathbf{S}_{1\dots m}$ is said to be $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable if X can be written as a convex combination of product operators of the form $P_1 \otimes \dots \otimes P_m$ where each $P_i \in \mathbf{S}_i^+$ is positive semidefinite. The set of $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operators forms a cone inside $(\mathbf{S}_{1\dots m})^+$. \square

Using this terminology, Definition 6.3 states that a quantum operation Λ is a LOSR operation if and only if $J(\Lambda)$ is a $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operator.

Another important example of separable operators are the separable quantum states. A state $\rho \in \mathbf{H}^+(\mathcal{X}_{1\dots m})$ is called *separable* if ρ can be written as a convex combination of product states of the form $\sigma_1 \otimes \dots \otimes \sigma_m$ where each $\sigma_i \in \mathbf{H}^+(\mathcal{X}_i)$ is a density operator. In other words, ρ denotes a separable state if and only if ρ is a $(\mathbf{H}(\mathcal{X}_1); \dots; \mathbf{H}(\mathcal{X}_m))$ -separable operator with trace equal to one. (Quantum states that are not separable are called *entangled*.)

Local operations with shared entanglement

LOSR operations are not the only quantum operations that can be implemented locally without communication. In the most general case, the parties could each hold a portion of some distinguished quantum state σ ; each party produces his output by applying some quantum operation to his portions of the input system and σ . In this case, the shared resource is σ .

Definition 6.5 (Local operation with shared entanglement (LOSE)). A quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ is a *m-party LOSE operation with finite entanglement* if there exist spaces $\mathcal{E}_1, \dots, \mathcal{E}_m$, a quantum state $\sigma \in \mathbf{H}^+(\mathcal{E}_{1\dots m})$, and quantum operations $\Psi_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$ for each $i = 1, \dots, m$ such that

$$\Lambda : X \mapsto (\Psi_{1\dots m})(X \otimes \sigma).$$

This arrangement is depicted for the two-party case in Figure 6.1.

The operation Λ is a *finitely approximable m-party LOSE operation* if it lies in the closure of the set of *m-party LOSE operations with finite entanglement*. The term “LOSE operation” is used to refer to any finitely approximable LOSE operation; the restriction to finite entanglement is made explicit whenever it is required. \square

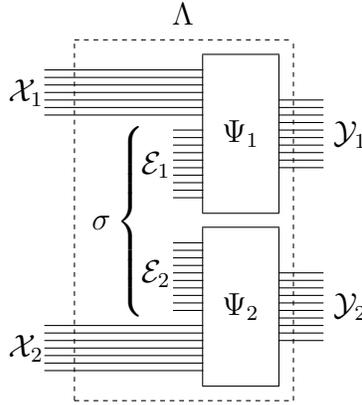


Figure 6.1: A two-party local quantum operation Λ with shared entanglement. The local operations are represented by Ψ_1, Ψ_2 ; the shared entanglement by σ .

Finite and infinite shared entanglement

The need to distinguish between LOSE operations with finite entanglement and finitely approximable LOSE operations arises from the fascinating fact that there exist LOSE operations that cannot be implemented with any finite amount of shared entanglement, yet can be approximated to arbitrary precision by LOSE operations with finite entanglement [LTW08].

Specifically, there exists an infinite sequence $\Lambda^{(1)}, \Lambda^{(2)}, \dots$ of LOSE operations with the property that the dimension of the space $\mathcal{E}_{1\dots m}^{(i)}$ associated with the shared state of $\Lambda^{(i)}$ is ever increasing with i . Moreover, this sequence is known to converge to a quantum operation that is not a LOSE operation with finite entanglement.

An analytic consequence of this fact is that the set of LOSE operations with finite entanglement is not a closed set. Fortunately, the work of this thesis is unhindered by a more encompassing notion of LOSE operation that includes the closure of that set.

Equivalence of shared randomness and shared separable states

The following straightforward but tedious proposition asserts that a LOSR operation is merely a LOSE operation for which the shared state is m -partite separable. It establishes that, as shared resources, a quantum state differs from randomness only when the state in question is entangled among the different parties.

Also established by this proposition is an explicit bound on the dimension of the shared state as a function of the dimensions of the input and output spaces. Thus, while there exist LOSE operations that cannot be implemented with a finite shared state, every LOSR operation can be implemented with a finite shared state whose size scales favorably in the size of the input and output states.

Proposition 6.6 (Equivalence of shared randomness and shared separable states).
A quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ is an m -party LOSR operation if and only if there exist

- (i) spaces $\mathcal{E}_1, \dots, \mathcal{E}_m$ of dimension $d = \dim(\mathbf{Q}_{1\dots m})$,
- (ii) a $(\mathbf{H}(\mathcal{E}_1); \dots; \mathbf{H}(\mathcal{E}_m))$ -separable state $\sigma \in \mathbf{H}^+(\mathcal{E}_{1\dots m})$, and
- (iii) quantum operations $\Psi_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$

such that $\Lambda : X \mapsto (\Psi_{1\dots m})(X \otimes \sigma)$.

Proof. Let σ be a $(\mathbf{H}(\mathcal{E}_1); \dots; \mathbf{H}(\mathcal{E}_m))$ -separable state and let $\Psi_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$ be quantum operations such that $\Lambda : X \mapsto (\Psi_{1\dots m})(X \otimes \sigma)$. Let

$$\sigma = \sum_j p_j \sigma_{1,j} \otimes \cdots \otimes \sigma_{m,j}$$

be a decomposition of σ into a convex combination of product states, where each $\sigma_{i,j} \in \mathbf{H}^+(\mathcal{E}_i)$. For each i and j define a quantum operation

$$\Phi_{i,j} : \mathbf{L}(\mathcal{X}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i) : X \mapsto \Psi_i(X \otimes \sigma_{i,j})$$

and observe that

$$\Lambda = \sum_j p_j \Phi_{1,j} \otimes \cdots \otimes \Phi_{m,j}$$

as desired.

Conversely, suppose that Λ may be decomposed into a convex combination of product quantum operations as above. By Carathéodory's Theorem (Fact 1.4), this sum may be assumed to have no more than d terms. Let $\mathcal{E}_1, \dots, \mathcal{E}_m$ be complex Euclidean spaces of dimension d and let $\rho_{i,1}, \dots, \rho_{i,d} \in \mathbf{H}^+(\mathcal{E}_i)$ be mutually orthogonal pure states for each i . Let

$$\sigma = \sum_{j=1}^d p_j \rho_{1,j} \otimes \cdots \otimes \rho_{m,j}$$

be a $(\mathbf{H}(\mathcal{E}_1); \dots; \mathbf{H}(\mathcal{E}_m))$ -separable state, let

$$\Psi_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{A}_i) : X \otimes E \mapsto \sum_{j=1}^d \langle E, \rho_{i,j} \rangle \cdot \Phi_{i,j}(X)$$

be quantum operations, and observe that $\Lambda : X \mapsto (\Psi_{1\dots m})(X \otimes \sigma)$. □

Convexity of local operations with shared entanglement

It is not difficult to see that the set of local operations is not a convex set. By definition, the set of LOSR operations is convex. That the set of LOSE operations with finite entanglement is also convex follows from a simple argument that is similar in principle to that of Proposition 6.6.

Proposition 6.7 (Convexity of LOSE operations). *The set of m -party LOSE operations with finite entanglement is convex. As a consequence, the set of finitely approximable m -party LOSE operations is also convex.*

Proof. Let $\Lambda, \Lambda' : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ be LOSE operations with finite entanglement. Choose spaces $\mathcal{E}_1, \dots, \mathcal{E}_m$ of large enough dimension so that there exist states $\sigma, \sigma' \in \mathbf{H}^+(\mathcal{E}_{1\dots m})$ and quantum operations $\Psi_i, \Psi'_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i)$ such that

$$\begin{aligned}\Lambda &: X \mapsto (\Psi_{1\dots m})(X \otimes \sigma), \\ \Lambda' &: X \mapsto (\Psi'_{1\dots m})(X \otimes \sigma').\end{aligned}$$

Let $\mathcal{F}_1, \dots, \mathcal{F}_m$ be two-dimensional spaces, let $\rho_i, \rho'_i \in \mathcal{F}_i$ be orthogonal pure states for each i , and let

$$\begin{aligned}\Phi_i &: \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i \otimes \mathcal{F}_i) \rightarrow \mathbf{L}(\mathcal{Y}_i) \\ &: X \otimes E \otimes F \mapsto \langle F, \rho_i \rangle \cdot \Psi_i(X \otimes E) + \langle F, \rho'_i \rangle \cdot \Psi'_i(X \otimes E)\end{aligned}$$

be quantum operations. For positive real numbers α, α' that sum to one, let

$$\xi = \alpha\sigma \otimes \rho_{1\dots m} + \alpha'\sigma' \otimes \rho'_{1\dots m}$$

be a state in $\mathbf{H}^+(\mathcal{E}_{1\dots m} \otimes \mathcal{F}_{1\dots m})$. It is easy to verify that

$$\alpha\Lambda + \alpha'\Lambda' : X \mapsto (\Phi_{1\dots m})(X \otimes \xi),$$

implying that the convex combination $\alpha\Lambda + \alpha'\Lambda'$ is also a LOSE operation with finite entanglement.

That the set of finitely approximable LOSE operations is convex follows immediately from the fact that this set is the closure of the set of LOSE operations with finite entanglement, which we just showed to be convex. \square

Chapter 7

Ball Around the Completely Noisy Channel

In this chapter we show that any quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ for which $J(\Lambda)$ lies in the product space $\mathbf{Q}_{1\dots m}$ and close enough to a distinguished “completely noisy” quantum operation $\tilde{\Delta}$ must necessarily be a LOSR operation. In other words, there is a ball of LOSR operations surrounding $\tilde{\Delta}$.

For each input space \mathcal{X} and output space \mathcal{Y} , the *completely noisy channel* $\tilde{\Delta} : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is the unique quantum operation defined by

$$\tilde{\Delta} : X \mapsto \frac{\text{Tr}(X)}{\dim(\mathcal{Y})} I_{\mathcal{Y}}.$$

The quantum state $\frac{1}{\dim(\mathcal{Y})} I_{\mathcal{Y}}$ is known as the *completely mixed state* of the system associated with \mathcal{Y} . Intuitively, this state denotes complete noise of the underlying system—it represents a uniform classical distribution over each of the $\dim(\mathcal{Y})$ levels in the system. Thus, the completely noisy channel is the quantum operation that always produces complete noise as output, regardless of the input.

The reason for interest in the completely noisy channel is that it is the unique super-operator whose Choi-Jamiołkowski representation equals the identity. More specifically, the completely noisy channel $\tilde{\Delta} : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ has

$$J(\tilde{\Delta}) = \frac{1}{\dim(\mathcal{Y})} I_{\mathcal{Y} \otimes \mathcal{X}}.$$

It is convenient for us to temporarily ignore the scalar multiple and deal directly with the unnormalized version Δ of $\tilde{\Delta}$ with

$$J(\Delta) = I_{\mathcal{Y} \otimes \mathcal{X}}.$$

The existence of a ball of LOSR operations around the completely noisy channel is established by proving that every operator in $\mathbf{Q}_{1\dots m}$ and close enough to the

identity must be a $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operator. Indeed, this fact is shown to hold not only for the specific choice $\mathbf{Q}_1, \dots, \mathbf{Q}_m$ of subspaces, but for *every* choice $\mathbf{S}_1, \dots, \mathbf{S}_m$ of subspaces that contain the identity.

Choosing $\mathbf{S}_1, \dots, \mathbf{S}_m$ to be full spaces of Hermitian operators yields an alternate (and simpler) proof of the existence of a ball of separable quantum states surrounding the completely mixed state and, consequently, of the NP-completeness of separability testing for quantum states. However, the ball of separable states implied by the present work is not as large as that exhibited by Gurvits and Barnum [GB02, GB03].

The technical results we require are proven in Section 7.1, as is a new bound for norms of super-operators. The application of these results to establish the existence of the ball of LOSR operations is provided in Section 7.2 along with some discussion of the details of this ball.

7.1 General results on separable operators

Due to the general nature of the results in this chapter, discussion in the present section is abstract—applications to quantum information are deferred until Section 7.2. The results presented herein were inspired by Chapter 2 of Bhatia [Bha07].

Hermitian subspaces generated by separable cones

Let \mathbf{S} be any subspace of Hermitian operators that contains the identity. The cone \mathbf{S}^+ always *generates* \mathbf{S} , meaning that each element of \mathbf{S} may be written as a difference of two elements of \mathbf{S}^+ . As proof, choose any $X \in \mathbf{S}$ and let

$$X^\pm = \frac{\|X\|I \pm X}{2}.$$

It is clear that $X = X^+ - X^-$ and that $X^\pm \in \mathbf{S}^+$ (using the fact that $I \in \mathbf{S}$). Moreover, it holds that $\|X^\pm\| \leq \|X\|$ for this particular choice of X^\pm .

In light of this observation, one might wonder whether it could be extended in product spaces to separable operators. In particular, do the $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operators generate the product space $\mathbf{S}_{1\dots m}$? If so, can elements of $\mathbf{S}_{1\dots m}$ be generated by $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operators with bounded norm? The following theorem answers these two questions in the affirmative.

Theorem 7.1 (Generation via bounded separable operators). *Let $\mathbf{S}_1, \dots, \mathbf{S}_m$ be subspaces of Hermitian operators—all of which contain the identity—and let $n = \dim(\mathbf{S}_{1\dots m})$. Then every element $X \in \mathbf{S}_{1\dots m}$ may be written $X = X^+ - X^-$ where X^\pm are $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable with*

$$\|X^\pm\| \leq 2^{m-1} \sqrt{n} \|X\|_{\mathbb{F}}.$$

Proof. First, it is proven that there is an orthonormal basis \mathbf{B} of $\mathbf{S}_{1\dots m}$ with the property that every element $E \in \mathbf{B}$ may be written $E = E^+ - E^-$ where E^\pm are $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable with $\|E^\pm\| \leq 2^{m-1}$. The proof is by straightforward induction on m . The base case $m = 1$ follows immediately from the earlier observation that every element $X \in \mathbf{S}_1$ is generated by some $X^\pm \in \mathbf{S}_1^\pm$ with $\|X^\pm\| \leq \|X\|$. In particular, any element $E = E^+ - E^-$ of any orthonormal basis of \mathbf{S}_1 has

$$\|E^\pm\| \leq \|E\| \leq \|E\|_F = 1 = 2^0.$$

In the general case, the induction hypothesis states that there is an orthonormal basis \mathbf{B}' of $\mathbf{S}_{1\dots m}$ with the desired property. Let \mathbf{B}_{m+1} be any orthonormal basis of \mathbf{S}_{m+1} . As in the base case, each $F \in \mathbf{B}_{m+1}$ is generated by some $F^\pm \in \mathbf{S}_{m+1}^\pm$ with $\|F^\pm\| \leq \|F\| \leq 1$. Define the orthonormal basis \mathbf{B} of $\mathbf{S}_{1\dots m+1}$ to consist of all product operators of the form $E \otimes F$ for $E \in \mathbf{B}'$ and $F \in \mathbf{B}_{m+1}$. Define

$$\begin{aligned} K^+ &\stackrel{\text{def}}{=} (E^+ \otimes F^+) + (E^- \otimes F^-), \\ K^- &\stackrel{\text{def}}{=} (E^+ \otimes F^-) + (E^- \otimes F^+). \end{aligned}$$

It is clear that $E \otimes F = K^+ - K^-$ and K^\pm are $(\mathbf{S}_1; \dots; \mathbf{S}_{m+1})$ -separable. Moreover,

$$\|K^+\| \leq \|E^+ \otimes F^+\| + \|E^- \otimes F^-\| \leq (2^{m-1} \times 1) + (2^{m-1} \times 1) = 2^m.$$

A similar computation yields $\|K^-\| \leq 2^m$, which establishes the induction.

Now, let $X \in \mathbf{S}_{1\dots m}$ and let $x_j \in \mathbb{R}$ be the unique coefficients of X in the aforementioned orthonormal basis $\mathbf{B} = \{E_1, \dots, E_n\}$. Define

$$\begin{aligned} X^+ &\stackrel{\text{def}}{=} \sum_{j: x_j > 0} x_j E_j^+ - \sum_{j: x_j < 0} x_j E_j^-, \\ X^- &\stackrel{\text{def}}{=} \sum_{j: x_j > 0} x_j E_j^- - \sum_{j: x_j < 0} x_j E_j^+. \end{aligned}$$

It is clear that $X = X^+ - X^-$ and X^\pm are $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable. Employing the triangle inequality and the vector norm inequality $\|x\|_1 \leq \sqrt{n}\|x\|_2$, it follows that

$$\|X^+\| \leq 2^{m-1} \sum_{j=1}^n |x_j| \leq 2^{m-1} \sqrt{n} \sqrt{\sum_{j=1}^n |x_j|^2} = 2^{m-1} \sqrt{n} \|X\|_F.$$

A similar computation yields $\|X^-\| \leq 2^{m-1} \sqrt{n} \|X\|_F$, which completes the proof. \square

Separable operators as a subtraction from the identity

At the beginning of this section it was observed that $\|X\|I - X$ lies in \mathbf{S}^+ for all $X \in \mathbf{S}$. One might wonder whether more could be expected of $\|X\|I - X$ than mere positive semidefiniteness. For example, under what conditions is $\|X\|I - X$ a $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operator? The following theorem provides three such conditions. Moreover, the central claim of this section is established by this theorem.

Theorem 7.2 (Ball of separable operators surrounding the identity). *Let $\mathbf{S}_1, \dots, \mathbf{S}_m$ be subspaces of Hermitian operators—all of which contain the identity—and let $n = \dim(\mathbf{S}_{1\dots m})$. The following hold:*

1. $\|P\| I - P$ is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable whenever P is a product operator of the form $P = P_1 \otimes \dots \otimes P_m$ where each $P_i \in \mathbf{S}_i^+$.
2. $(n + 1) \|Q\| I - Q$ is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable whenever Q is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable.
3. $2^{m-1} \sqrt{n} (n + 1) \|X\|_{\mathbb{F}} I - X$ is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable for every $X \in \mathbf{S}_{1\dots m}$.

Proof. The proof of item 1 is an easy but notationally cumbersome induction on m . The base case $m = 1$ was noted at the beginning of Section 7.1. For the general case, it is convenient to let $I_{1\dots m}$, I_{m+1} , and $I_{1\dots m+1}$ denote the identity elements of $\mathbf{S}_{1\dots m}$, \mathbf{S}_{m+1} , and $\mathbf{S}_{1\dots m+1}$, respectively. The induction hypothesis states that the operator $S \stackrel{\text{def}}{=} \|P_{1\dots m}\| I_{1\dots m} - P_{1\dots m}$ is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable. Just as in the base case, we know $S' \stackrel{\text{def}}{=} \|P_{m+1}\| I_{m+1} - P_{m+1}$ lies in \mathbf{S}_{m+1}^+ . Isolating the identity elements, these expressions may be rewritten

$$I_{1\dots m} = \frac{1}{\|P_{1\dots m}\|} (P_{1\dots m} + S)$$

$$I_{m+1} = \frac{1}{\|P_{m+1}\|} (P_{m+1} + S').$$

Taking the Kronecker product of these two equalities and rearranging the terms yields

$$\|P_{1\dots m+1}\| I_{1\dots m+1} - P_{1\dots m+1} = (P_{1\dots m} \otimes S') + (S \otimes P_{m+1}) + (S \otimes S').$$

The right side of this expression is clearly a $(\mathbf{S}_1; \dots; \mathbf{S}_{m+1})$ -separable operator; the proof by induction is complete.

Item 2 is proved as follows. By Carathéodory's Theorem (Fact 1.4), every $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operator Q may be written as a sum of no more than $n + 1$ product operators. In particular,

$$Q = \sum_{j=1}^{n+1} P_{1,j} \otimes \dots \otimes P_{m,j}$$

where each $P_{i,j}$ is an element of \mathbf{S}_i^+ . As each term in this sum is positive semidefinite, it holds that $\|Q\| \geq \|P_{1,j} \otimes \dots \otimes P_{m,j}\|$ for each j . Item 1 implies that the sum

$$\sum_{j=1}^{n+1} \|P_{1,j} \otimes \dots \otimes P_{m,j}\| I - P_{1,j} \otimes \dots \otimes P_{m,j}$$

is also $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable. Naturally, each of the identity terms

$$\|P_{1,j} \otimes \dots \otimes P_{m,j}\| I$$

in this sum may be replaced by $\|Q\|I$ without compromising $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separability, from which it follows that $(n+1)\|Q\|I - Q$ is also $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable.

To prove item 3, apply Theorem 7.1 to obtain $X = X^+ - X^-$ where X^\pm are $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable with $\|X^\pm\| \leq 2^{m-1}\sqrt{n}\|X\|_F$. By item 2, it holds that $(n+1)\|X^+\|I - X^+$ is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable, implying that

$$2^{m-1}\sqrt{n}(n+1)\|X\|_F I - X^+$$

is also $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable. To complete the proof, it suffices to note that adding X^- to this operator yields another $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operator. \square

Norms of positive super-operators

We are now in a position to prove new bounds on the capacity of a positive super-operator to increase the norm of its output relative to that of its input. While these new bounds are not used later in this thesis, it is worthwhile to note them and compare them to previously known bounds of this nature.

The seminal such bound is due to Russo and Dye and states that

$$\|\Phi(X)\| \leq \|\Phi(I)\| \|X\|$$

for all X whenever Φ is positive [RD66]. In particular, if Φ does not increase the operator norm of the identity then Φ does not increase the operator norm of *any* operator. When discussing bounds of this form, it is convenient to make the assumption that $\|\Phi(I)\| \leq 1$, which allows us to state the bound more succinctly as

$$\|\Phi(X)\| \leq \|X\|.$$

Of course, the original bound, which applies to *all* positive Φ , can be recovered from this succinct version via suitable rescaling of Φ by a factor of $\|\Phi(I)\|$.

Consider a subspace $\mathbf{S} \subseteq \mathbf{L}(\mathcal{X})$ that contains the identity. In Bhatia [Bha07] it is shown that

$$\|\Phi(X)\| \leq \sqrt{2}\|X\|$$

for all $X \in \mathbf{S}$ whenever $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is positive on \mathbf{S}^+ . Essentially, the Russo-Dye bound weakens by a factor of $\sqrt{2}$ when we relax the positivity requirement to a mere subspace of the input space. Gurvits and Barnum showed that

$$\|\Phi(X)\| \leq \sqrt{2}\|X\|_F$$

for all X whenever $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ is positive on the unit ball

$$\{X : \|I_{\mathcal{X}} - X\|_F \leq 1\} \subset \mathbf{L}(\mathcal{X})$$

surrounding the identity [GB03]. This bound is considerably weaker than the previous two because the operator norm $\|X\|$ of X is replaced with the Frobenius norm

$\|X\|_{\mathbb{F}}$. On the other hand, it has the advantage that Φ need not be positive on an entire subspace of $\mathbf{L}(\mathcal{X})$. Gurvits and Barnum used this additional flexibility to establish the existence of a ball of multipartite separable quantum states surrounding the completely mixed state.

To the above list of bounds we add one of our own. While the Gurvits-Barnum bound applies only to super-operators $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{Y})$ that are positive on a set that spans all of $\mathbf{L}(\mathcal{X})$ (namely, the unit ball around the identity), the super-operators to which our bound applies might only be positive on a set that does not span all of $\mathbf{L}(\mathcal{X})$. Alas, this flexibility comes at the cost of a multiple of the dimension of the input space. For context, our bound is stated as an extension of Theorem 2.6.3 of Bhatia [Bha07].

Theorem 7.3 (Norms of positive super-operators). *Let $\mathbf{S}_1, \dots, \mathbf{S}_m$ be subspaces of Hermitian operators—all of which contain the identity—and let $n = \dim(\mathbf{S}_{1\dots m})$. Let Φ be a super-operator acting on $\mathbf{S}_{1\dots m}$ with the property that Φ is positive on $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable operators. The following hold:*

1. $\|\Phi(P)\| \leq \|P\| \|\Phi(I)\|$ whenever P is a product operator of the form $P = P_1 \otimes \dots \otimes P_m$ where each $P_i \in \mathbf{S}_i^+$.
2. $\|\Phi(Q)\| \leq (n+1) \|Q\| \|\Phi(I)\|$ whenever Q is $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable.
3. $\|\Phi(X)\| \leq 2^{m-1} \sqrt{n} (n+1) \|X\|_{\mathbb{F}} \|\Phi(I)\|$ for every $X \in \mathbf{S}_{1\dots m}$.

Proof. To prove item 1, we observe that

$$\Phi(\|P\|I - P) = \|P\|\Phi(I) - \Phi(P)$$

is positive semidefinite. As $\Phi(I)$ and $\Phi(P)$ are also positive semidefinite, it follows that $\|\Phi(P)\| \leq \|P\| \|\Phi(I)\|$. Item 2 is proved by the same argument with $(n+1)\|Q\|I - Q$ in place of $\|P\|I - P$.

To prove item 3, apply Theorem 7.1 to obtain $X = X^+ - X^-$ where X^\pm are $(\mathbf{S}_1; \dots; \mathbf{S}_m)$ -separable with $\|X^\pm\| \leq 2^{m-1} \sqrt{n} \|X\|_{\mathbb{F}}$. As $\Phi(X^\pm)$ are positive semidefinite, we have

$$\begin{aligned} \|\Phi(X)\| &= \|\Phi(X^+) - \Phi(X^-)\| \\ &\leq \max(\|\Phi(X^+)\|, \|\Phi(X^-)\|) \\ &\leq \max(\|X^+\|, \|X^-\|) (n+1) \|\Phi(I)\| \\ &\leq 2^{m-1} \sqrt{n} (n+1) \|X\|_{\mathbb{F}} \|\Phi(I)\|. \end{aligned}$$

□

7.2 Ball around the completely noisy channel

In Chapter 6 we noted that a quantum operation Λ is a LOSR operation if and only if $J(\Lambda)$ is $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable. We then established via Theorem 7.2 of the previous section that any operator in the product space $\mathbf{Q}_{1\dots m}$ and close enough to the identity is necessarily a $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operator, from which the existence of a ball of LOSR operations around the completely noisy channel follows.

Theorem 7.4 (Ball around the completely noisy channel). *Let $n = \dim(\mathbf{Q}_{1\dots m})$ and let $k = 2^{m-1}\sqrt{n(n+1)}$. For each operator $A \in \mathbf{Q}_{1\dots m}$ with $\|A\|_F \leq \frac{1}{k}$ there exists an unnormalized LOSR operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ for which*

$$J(\Lambda) = I - A.$$

As a consequence, any quantum operation $\Xi : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$ with

$$\|J(\Xi) - J(\tilde{\Delta})\|_F \leq \frac{1}{kd}$$

is a LOSR operation. Here $\tilde{\Delta}$ denotes the completely noisy channel and $d = \dim(\mathcal{Y}_{1\dots m})$.

Theorem 7.4 establishes a ball of LOSR operations (and hence also of LOSE operations) surrounding the completely noisy channel. However, there seems to be no obvious way to obtain a bigger ball if such a ball is allowed to contain operations that are LOSE but not LOSR. Perhaps a more careful future investigation will uncover such a ball.

The subspace containing the ball of LOSR operations

The ball of Theorem 7.4 is contained within the product space $\mathbf{Q}_{1\dots m}$, which is a strict subspace of the space spanned by all quantum operations $\Phi : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{Y}_{1\dots m})$. Why was attention restricted to this subspace? The answer is that there are no LOSE or LOSR operations Λ for which $J(\Lambda)$ lies outside $\mathbf{Q}_{1\dots m}$. In other words, $\mathbf{Q}_{1\dots m}$ is the *largest* possible space in which to find a ball of LOSR operations. We shall return to this topic in Chapter 10 wherein it is shown that the space $\mathbf{Q}_{1\dots m}$ is generated by the so-called *no-signaling* quantum operations.

Of course, there exist quantum operations arbitrarily close to the completely noisy channel that are not no-signaling operations, much less LOSE or LOSR operations. This fact might seem to confuse the study of, say, the effects of noise on such operations because a completely general model of noise would allow for extremely tiny perturbations that nonetheless turn no-signaling operations into signaling operations. This confusion might even be exacerbated by the fact that separable quantum states, by contrast, are resilient to *arbitrary* noise: any conceivable physical perturbation of the completely mixed state is separable, so long as the perturbation has small enough magnitude.

There is, of course, nothing unsettling about this picture. In any reasonable model of noise, perturbations to a LOSE or LOSR operation occur only on the local operations performed by the parties involved, or perhaps on the state they share. It is easy to see that realistic perturbations such as these always maintain the no-signaling property of these operations. Moreover, any noise *not* of this form could, for example, bestow faster-than-light communication upon spatially separated parties.

Chapter 8

Recognizing LOSE Operations is NP-hard

In this chapter the existence of the ball of LOSR operations around the completely noisy channel is employed to prove that the weak membership problem for LOSE operations is strongly NP-hard. Informally, the weak membership problem asks,

“Given a description of a quantum operation Λ and an accuracy parameter ε , is Λ within distance ε of a LOSE operation?”

This result is achieved in several stages. Section 8.1 reviews a relevant recent result of Kempe *et al.* pertaining to quantum games. In Section 8.2 this result is exploited in order to prove that the weak validity problem—a relative of the weak membership problem—is strongly NP-hard for LOSE operations. Finally, Section 8.3 illustrates how the strong NP-hardness of the weak membership problem for LOSE operations follows from a Gurvits-Gharibian-style application of Liu’s version of the Yudin-Nemirovskii Theorem. It is also noted that similar NP-hardness results hold trivially for LOSR operations, due to the fact that separable quantum states arise as a special case of LOSR operations in which the input space is empty.

8.1 Co-operative quantum games with shared entanglement

Local operations with shared entanglement have been previously studied in the context of two-player co-operative games. In these games, a referee prepares a question for each player and the players each respond to the referee with an answer. The referee evaluates these answers and declares that the players have jointly won or lost the game according to this evaluation. The goal of the players, then, is to coordinate their answers so as to maximize the probability with which the referee

declares them to be winners. In a *quantum* game the questions and answers are quantum states.

In order to differentiate this model from a one-player game, the players are not permitted to communicate with each other after the referee has sent his questions. The players can, however, meet prior to the commencement of the game in order to agree on a strategy. In a quantum game the players might also prepare a shared entangled quantum state so as to enhance the coordination of their answers to the referee.

More formally, a *quantum game* $G = (q, \pi, \mathbf{R}, \mathbf{V})$ is specified by:

- A positive integer q denoting the number of distinct questions.
- A probability distribution π on the question indices $\{1, \dots, q\}$, according to which the referee selects his questions.
- Complex Euclidean spaces $\mathcal{V}, \mathcal{X}_1, \mathcal{X}_2, \mathcal{A}_1, \mathcal{A}_2$ corresponding to the different quantum systems used by the referee and players.
- A set \mathbf{R} of quantum states $\mathbf{R} = \{\rho_i\}_{i=1}^q \subset \mathbf{H}^+(\mathcal{V} \otimes \mathcal{X}_1 \otimes \mathcal{X}_2)$. These states correspond to questions and are selected by the referee according to π .
- A set \mathbf{V} of unitary operators $\mathbf{V} = \{V_i\}_{i=1}^q \subset \mathbf{L}(\mathcal{V} \otimes \mathcal{A}_1 \otimes \mathcal{A}_2)$. These unitaries are used by the referee to evaluate the players' answers.

For convenience, the two players are called *Alice* and *Bob*. The game is played as follows. The referee samples i according to π and prepares the state $\rho_i \in \mathbf{R}$, which is placed in the three quantum registers corresponding to $\mathcal{V} \otimes \mathcal{X}_1 \otimes \mathcal{X}_2$. This state contains the questions to be sent to the players: the portion of ρ_i corresponding to \mathcal{X}_1 is sent to Alice, the portion of ρ_i corresponding to \mathcal{X}_2 is sent to Bob, and the portion of ρ_i corresponding to \mathcal{V} is kept by the referee as a private workspace. In reply, Alice sends a quantum register corresponding to \mathcal{A}_1 to the referee, as does Bob to \mathcal{A}_2 . The referee then applies the unitary operation $V_i \in \mathbf{V}$ to the three quantum registers corresponding to $\mathcal{V} \otimes \mathcal{A}_1 \otimes \mathcal{A}_2$, followed by a standard measurement $\{\Pi_{\text{accept}}, \Pi_{\text{reject}}\}$ that dictates the result of the game.

As mentioned at the beginning of this subsection, Alice and Bob may not communicate once the game commences. But they may meet prior to the commencement of the game in order to prepare a shared entangled quantum state σ . Upon receiving the question register corresponding to \mathcal{X}_1 from the referee, Alice may perform any physically realizable quantum operation upon that register and upon her portion of σ . The result of this operation shall be contained in the quantum register corresponding to \mathcal{A}_1 —this is the answer that Alice sends to the referee. Bob follows a similar procedure to obtain his own answer register corresponding to \mathcal{A}_2 .

For any game G , the *value* $\omega(G)$ of G is the supremum of the probability with which the referee can be made to accept taken over all strategies of Alice and Bob.

Theorem 8.1 (Kempe *et al.* [KKM⁺08]). *There is a fixed polynomial p such that the following promise problem is NP-hard under mapping (Karp) reductions:*

Input. *A quantum game $G = (q, \pi, \mathbf{R}, \mathbf{V})$. The distribution π and the sets \mathbf{R}, \mathbf{V} are each given explicitly: for each $i = 1, \dots, q$, the probability $\pi(i)$ is given in binary, as are the real and complex parts of each entry of the matrices ρ_i and V_i .*

Yes. *The value $\omega(G)$ of the game G is 1.*

No. *The value $\omega(G)$ of the game G is less than $1 - \frac{1}{p(q)}$.*

8.2 Strategies and weak validity

Viewing the two players as a single entity, a quantum game may be seen as a two-message quantum interaction between the referee and the players—a message from the referee to the players, followed by a reply from the players to the referee. The actions of the referee during such an interaction are completely specified by the parameters of the game.

In the language of Part I, the game specifies a one-round measuring co-strategy for the referee represented by some positive semidefinite operators

$$R_{\text{accept}}, R_{\text{reject}} \in \mathbf{H}^+(\mathcal{A}_{1\dots 2} \otimes \mathcal{X}_{1\dots 2}),$$

which are easily computed given the parameters of the game.

In these games, the players implement a one-round non-measuring strategy compatible with $\{R_{\text{accept}}, R_{\text{reject}}\}$ whose representation is given by a positive semidefinite operator

$$P \in \mathbf{H}^+(\mathcal{A}_{1\dots 2} \otimes \mathcal{X}_{1\dots 2}).$$

For any fixed strategy P for the players, Theorem 3.1 tells us that the probability with which the players cause the referee to accept is given by the inner product

$$\Pr[\text{Players win with strategy } P] = \langle R_{\text{accept}}, P \rangle.$$

In any game, the players combine to implement some physical operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots 2}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots 2})$. It is clear that a given super-operator Λ denotes a legal strategy for the players if and only if Λ is a LOSE operation. As the players implement a one-round non-measuring strategy, the representation P of their strategy is given by $P = J(\Lambda)$.

Thus, the problem studied by Kempe *et al.* [KKM⁺08] of deciding whether $\omega(G) = 1$ can be reduced via the formalism of strategies to an optimization problem over the set of LOSE operations:

$$\omega(G) = \sup_{\Lambda \in \text{LOSE}} \langle R_{\text{accept}}, J(\Lambda) \rangle.$$

The following theorem is thus proved.

Theorem 8.2. *The weak validity problem for the set of LOSE [LOSR] operations is strongly NP-hard [NP-complete] under mapping (Karp) reductions:*

Input. *A Hermitian matrix R , a real number γ , and a positive real number $\varepsilon > 0$. The number γ is given explicitly in binary, as are the real and complex parts of each entry of R . The number ε is given in unary, where 1^s denotes $\varepsilon = 1/s$.*

Yes. *There exists a LOSE [LOSR] operation Λ such that $\langle R, J(\Lambda) \rangle \geq \gamma + \varepsilon$.*

No. *For every LOSE [LOSR] operation Λ we have $\langle R, J(\Lambda) \rangle \leq \gamma - \varepsilon$.*

Remark 8.2.1. The hardness result for LOSR operations follows from a simple reduction from separable quantum states to LOSR operations: every separable state may be written as a LOSR operation in which the input space has dimension one. That weak validity for LOSR operations is in NP (and is therefore NP-complete) follows from the fact that all LOSR operations may be implemented with polynomially-bounded shared randomness (Proposition 6.6).

8.3 The Yudin-Nemirovskii Theorem and weak membership

Having established that the weak *validity* problem for LOSE operations is strongly NP-hard, the next step is to follow the leads of Gurvits and Gharibian [Gur02, Gha08] and apply Liu’s version [Liu07] of the Yudin-Nemirovskii Theorem [YN76, GLS88] in order to prove that the weak *membership* problem for LOSE operations is also strongly NP-hard.

The Yudin-Nemirovskii Theorem establishes an oracle-polynomial-time reduction from the weak validity problem to the weak membership problem for any convex set C that satisfies certain basic conditions. One consequence of this theorem is that if the weak validity problem for C is NP-hard then the associated weak membership problem for C is also NP-hard. Although hardness under mapping reductions is preferred, any hardness result derived from the Yudin-Nemirovskii Theorem in this way is only guaranteed to hold under more inclusive oracle reductions.

The basic conditions that must be met by the set C in order for the Yudin-Nemirovskii Theorem to apply are

- (i) C is bounded,
- (ii) C contains a ball, and
- (iii) the size of the bound and the ball are polynomially related to the dimension of the vector space containing C .

It is simple to check these criteria against the set of LOSE operations. For condition (i), an explicit bound is established by Proposition 1.3. The remaining conditions then follow from Theorem 7.4.

The following theorem is now proved. (As in Remark 8.2.1, the analogous result for LOSR operations follows from a straightforward reduction from separable quantum states.)

Theorem 8.3. *The weak membership problem for the set of LOSE [LOSR] operations is strongly NP-hard [NP-complete] under oracle (Cook) reductions:*

Input. *A Hermitian matrix $X \in \mathbf{Q}_{1\dots m}$ and a positive real number $\varepsilon > 0$. The real and complex parts of each entry of X are given explicitly in binary. The number ε is given in unary, where 1^s denotes $\varepsilon = 1/s$.*

Yes. *$X = J(\Lambda)$ for some LOSE [LOSR] operation Λ .*

No. *$\|X - J(\Lambda)\| \geq \varepsilon$ for every LOSE [LOSR] operation Λ .*

Chapter 9

Characterizations of Local Operations

Characterizations of LOSE and LOSR operations are presented in this chapter. These characterizations are reminiscent of the well-known characterizations of bipartite and multipartite separable quantum states due to Horodecki *et al.* [HHH96, HHH01].

Specifically, it is proven in Section 9.1 that Λ is a LOSR operation if and only if $\varphi(J(\Lambda)) \geq 0$ for every linear functional φ that is positive on the cone of $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operators. (This cone is defined in Definitions 6.4 and 6.2.) This characterization of LOSR operations is proved by a straightforward application of the fundamental Separation Theorems of convex analysis (Facts 1.5 and 1.6).

More interesting is the characterization of LOSE operations presented in Section 9.2: Λ is a LOSE operation if and only if $\varphi(J(\Lambda)) \geq 0$ for every linear functional φ that is *completely* positive on that *same cone* of $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operators. Prior to the present work, the notion of complete positivity was only ever considered in the context wherein the underlying cone is the positive semidefinite cone. Indeed, what it even *means* for a super-operator or functional to be “completely” positive on some cone other than the positive semidefinite cone must be clarified before any nontrivial discussion can occur.

9.1 Characterization of local operations with shared randomness

The characterization of LOSR operations presented herein is an immediate corollary of the following simple proposition.

Proposition 9.1. *Let $K \subset \mathbb{R}^n$ be any closed convex cone. A vector $x \in \mathbb{R}^n$ is an element of K if and only if $\varphi(x) \geq 0$ for every linear functional $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ that*

is positive on K .

Proof. The “only if” part of the proposition is immediate: as x is in K , any linear functional positive on K must also be positive on x . For the “if” part of the proposition, suppose that x is not an element of K . The Separation Theorem (Fact 1.5) implies that there exists a vector $h \in \mathbb{R}^n$ such that $\langle h, y \rangle \geq 0$ for all $y \in K$, yet $\langle h, x \rangle < 0$. (To apply Fact 1.5, consider a small open ball that is disjoint from K and contains x .) Let $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ be the linear functional given by $\varphi : z \mapsto \langle h, z \rangle$. It is clear that φ is positive on K , yet $\varphi(x) < 0$. \square

Corollary 9.1.1 (Characterization of LOSR operations). *A quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots m})$ is an m -party LOSR operation if and only if $\varphi(J(\Lambda)) \geq 0$ for every linear functional $\varphi : \mathbf{L}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}) \rightarrow \mathbb{C}$ that is positive on the cone of $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operators.*

Proof. In order to apply Proposition 9.1, it suffices to note the following:

- The space $\mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$ is isomorphic to \mathbb{R}^n for $n = \dim(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})^2$.
- The $(\mathbf{Q}_1; \dots; \mathbf{Q}_m)$ -separable operators form a closed convex cone within the space $\mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$.

While Proposition 9.1 only gives the desired result for real linear functionals $\varphi : \mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}) \rightarrow \mathbb{R}$, it is trivial to construct a complex extension functional $\varphi' : \mathbf{L}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}) \rightarrow \mathbb{C}$ that agrees with φ on $\mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$. \square

9.2 Characterization of local operations with shared entanglement

In order to state our characterization of LOSE operations, it is convenient to parameterize the shorthand notation \mathbf{Q}_i of Definition 6.2, which denotes the subspace of Hermitian operators X with $X = J(\Psi)$ for some trace-preserving super-operator Ψ , or a scalar multiple thereof.

Definition 9.2 (Parameterization $\mathbf{Q}_i(\mathcal{E}_i)$ of the shorthand notation \mathbf{Q}_i). For each $i = 1, \dots, m$ and each complex Euclidean space \mathcal{E}_i , let

$$\mathbf{Q}_i(\mathcal{E}_i) \subset \mathbf{H}(\mathcal{A}_i \otimes \mathcal{X}_i \otimes \mathcal{E}_i)$$

denote the subspace of operators $J(\Psi)$ for which

$$\Psi : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{A}_i)$$

is a trace-preserving super-operator, or a scalar multiple thereof. In particular, the parameter space \mathcal{E}_i is always tensored with the *input* space of Ψ , as opposed to the output space. \square

Theorem 9.3 (Characterization of LOSE operations). *A quantum operation*

$$\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots m})$$

is an m -party LOSE operation if and only if $\varphi(J(\Lambda)) \geq 0$ for every linear functional

$$\varphi : \mathbf{L}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}) \rightarrow \mathbb{C}$$

with the property that the super-operator $(\varphi \otimes \mathbb{1}_{\mathcal{E}_{1\dots m}})$ is positive on the cone of $(\mathbf{Q}_1(\mathcal{E}_1); \dots; \mathbf{Q}_m(\mathcal{E}_m))$ -separable operators for all choices of complex Euclidean spaces $\mathcal{E}_1, \dots, \mathcal{E}_m$.

Remark 9.3.1. The positivity condition of Theorem 9.3 bears striking resemblance to the familiar notion of complete positivity of a super-operator. With this resemblance in mind, a linear functional φ obeying the positivity condition of Theorem 9.3 is said to be *completely positive on the $(\mathbf{Q}_1(\mathcal{E}_1); \dots; \mathbf{Q}_m(\mathcal{E}_m))$ -separable family of cones*. In this sense, Theorem 9.3 represents what seems, to the knowledge of the author, to be the first application of the notion of complete positivity to a family of cones other than the positive semidefinite family.

Moreover, for any fixed choice of complex Euclidean spaces $\mathcal{E}_1, \dots, \mathcal{E}_m$ there exists a linear functional φ for which the super-operator $(\varphi \otimes \mathbb{1}_{\mathcal{E}_{1\dots m}})$ is positive on the cone of $(\mathbf{Q}_1(\mathcal{E}_1); \dots; \mathbf{Q}_m(\mathcal{E}_m))$ -separable operators, and yet φ is *not* completely positive on this family of cones. This curious property is a consequence of the fact the set of LOSE operations with finite entanglement is not a closed set. By contrast, complete positivity (in the traditional sense) of a super-operator $\Phi : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{A})$ is assured whenever $(\Phi \otimes \mathbb{1}_{\mathcal{Z}})$ is positive for a space \mathcal{Z} with dimension at least that of \mathcal{X} . (See, for example, Bhatia [Bha07] for a proof of this fact.)

The proof of Theorem 9.3 employs the following helpful identity involving the Choi-Jamiołkowski representation for super-operators. This identity is proven by straightforward calculation.

Proposition 9.4. *Let $\Psi : \mathbf{L}(\mathcal{X} \otimes \mathcal{E}) \rightarrow \mathbf{L}(\mathcal{A})$ and let $Z \in \mathbf{L}(\mathcal{E})$. Then the super-operator $\Lambda : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{A})$ defined by $\Lambda(X) = \Psi(X \otimes Z)$ for all X satisfies $J(\Lambda) = \text{Tr}_{\mathcal{E}}((I_{\mathcal{A} \otimes \mathcal{X}} \otimes Z^{\top}) J(\Psi))$.*

Proof. Let $\{e_1, \dots, e_{\dim(\mathcal{X})}\}$ and $\{f_1, \dots, f_{\dim(\mathcal{E})}\}$ denote the standard bases of \mathcal{X} and \mathcal{E} , respectively, and let $z_{k,l} \in \mathbb{C}$ be the coefficients of Z in the basis $\{f_k f_l^* : k, l = 1, \dots, \dim(\mathcal{E})\}$ of $\mathbf{L}(\mathcal{E})$. We have

$$\begin{aligned} \text{Tr}_{\mathcal{E}}((I_{\mathcal{A} \otimes \mathcal{X}} \otimes Z^{\top}) J(\Psi)) &= \sum_{i,j=1}^{\dim(\mathcal{X})} \sum_{k,l=1}^{\dim(\mathcal{E})} \Psi(e_i e_j^* \otimes f_k f_l^*) \otimes e_i e_j^* \cdot \underbrace{\text{Tr}(Z^{\top} f_k f_l^*)}_{z_{k,l}} \\ &= \sum_{i,j=1}^{\dim(\mathcal{X})} \Psi\left(e_i e_j^* \otimes \sum_{k,l=1}^{\dim(\mathcal{E})} z_{k,l} f_k f_l^*\right) \otimes e_i e_j^* = J(\Lambda). \end{aligned}$$

□

The following technical lemma is also employed in the proof of Theorem 9.3.

Lemma 9.5. *Let $\Psi : \mathbf{L}(\mathcal{X} \otimes \mathcal{E}) \rightarrow \mathbf{L}(\mathcal{A})$, let $Z \in \mathbf{L}(\mathcal{E})$, and let $\varphi : \mathbf{L}(\mathcal{A} \otimes \mathcal{X}) \rightarrow \mathbb{C}$. Then the super-operator $\Lambda : \mathbf{L}(\mathcal{X}) \rightarrow \mathbf{L}(\mathcal{A})$ defined by $\Lambda(X) = \Psi(X \otimes Z)$ for all X satisfies*

$$\varphi(J(\Lambda)) = \langle \bar{Z}, (\varphi \otimes \mathbb{1}_{\mathcal{E}})(J(\Psi)) \rangle.$$

Proof. Let H be the unique operator satisfying $\varphi(X) = \langle H, X \rangle$ for all X and note that the adjoint $\varphi^* : \mathbb{C} \rightarrow \mathbf{L}(\mathcal{A} \otimes \mathcal{X})$ satisfies $\varphi^*(1) = H$. Then

$$\begin{aligned} \langle \bar{Z}, (\varphi \otimes \mathbb{1}_{\mathcal{E}})(J(\Psi)) \rangle &= \langle \varphi^*(1) \otimes \bar{Z}, J(\Psi) \rangle = \langle H \otimes \bar{Z}, J(\Psi) \rangle \\ &= \langle H, \text{Tr}_{\mathcal{E}}((I_{\mathcal{A} \otimes \mathcal{X}} \otimes Z^{\top}) J(\Psi)) \rangle = \varphi(J(\Lambda)). \end{aligned}$$

□

Proof of Theorem 9.3. For the “only if” part of the theorem, let Λ be any LOSE operation with finite entanglement and let $\Psi_1, \dots, \Psi_m, \sigma$ be such that

$$\Lambda : X \mapsto (\Psi_{1\dots m})(X \otimes \sigma).$$

Let φ be any linear functional on $\mathbf{L}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$ that satisfies the stated positivity condition. Lemma 9.5 implies

$$\varphi(J(\Lambda)) = \langle \bar{\sigma}, (\varphi \otimes \mathbb{1}_{\mathcal{E}_{1\dots m}})(J(\Psi_{1\dots m})) \rangle \geq 0.$$

A standard continuity argument establishes the desired implication when Λ is a finitely approximable LOSE operation.

For the “if” part of the theorem, suppose that $\Xi : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots m})$ is a quantum operation that is *not* a LOSE operation. The Separation Theorem (Fact 1.6) implies that there is a Hermitian operator $H \in \mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$ such that $\langle H, J(\Lambda) \rangle \geq 0$ for all LOSE operations Λ , yet $\langle H, J(\Xi) \rangle < 0$. Let

$$\varphi : \mathbf{L}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}) \rightarrow \mathbb{C}$$

be the linear functional given by $\varphi : X \mapsto \langle H, X \rangle$.

It remains to verify that φ satisfies the desired positivity condition. Toward that end, let $\mathcal{E}_1, \dots, \mathcal{E}_m$ be arbitrary complex Euclidean spaces. By convexity, it suffices to consider only those $(\mathbf{Q}_1(\mathcal{E}_1); \dots; \mathbf{Q}_m(\mathcal{E}_m))$ -separable operators that are product operators. Choose any such operator and note that, up to a scalar multiple, it has the form $J(\Psi_{1\dots m})$ where each $\Psi_i : \mathbf{L}(\mathcal{X}_i \otimes \mathcal{E}_i) \rightarrow \mathbf{L}(\mathcal{A}_i)$ is a quantum operation. The operator

$$(\varphi \otimes \mathbb{1}_{\mathcal{E}_{1\dots m}})(J(\Psi_{1\dots m}))$$

is positive semidefinite if and only if it has a nonnegative inner product with every density operator in $\mathbf{H}^+(\mathcal{E}_{1\dots m})$. As σ ranges over all such operators, so does $\bar{\sigma}$. Moreover, every such σ —together with Ψ_1, \dots, Ψ_m —induces a LOSE operation Λ defined by $\Lambda : X \mapsto (\Psi_{1\dots m})(X \otimes \sigma)$. Lemma 9.5 and the choice of φ imply

$$0 \leq \varphi(J(\Lambda)) = \langle \bar{\sigma}, (\varphi \otimes \mathbb{1}_{\mathcal{E}_{1\dots m}})(J(\Psi_{1\dots m})) \rangle,$$

and so φ satisfies the desired positivity condition. □

Chapter 10

No-Signaling Operations

At the end of Chapter 7 it was claimed that the product space $\mathbf{Q}_{1\dots m}$ is spanned by Choi-Jamiołkowski representations of no-signaling operations. It appears as though this fact has yet to be noted explicitly in the literature, so a proof is offered in this chapter.

More accurately, two characterizations of no-signaling operations are presented in Section 10.2, each of which is expressed as a condition on Choi-Jamiołkowski representations of super-operators. It then follows immediately that $\mathbf{Q}_{1\dots m}$ is spanned by Choi-Jamiołkowski representations of no-signaling operations.

Finally, Section 10.3 provides an example of a so-called *separable* no-signaling operation that is not a LOSE operation, thus ruling out an easy “short cut” to the ball of LOSR operations revealed in Theorem 7.4.

10.1 Formal definition of a no-signaling operation

Intuitively, a quantum operation Λ is no-signaling if it cannot be used by spatially separated parties to violate relativistic causality. Put another way, an operation Λ jointly implemented by several parties is no-signaling if those parties cannot use Λ as a “black box” to communicate with one another.

In order to facilitate a formal definition for no-signaling operations, the shorthand notation for Kronecker products from Chapter 1 must be extended: if $K \subseteq \{1, \dots, m\}$ is an arbitrary index set with $K = \{k_1, \dots, k_n\}$ then we write

$$\mathcal{X}_K \stackrel{\text{def}}{=} \mathcal{X}_{k_1} \otimes \cdots \otimes \mathcal{X}_{k_n}$$

with the convention that $\mathcal{X}_\emptyset = \mathbb{C}$. As with the original shorthand, a similar notation also applies to operators, sets of operators, and super-operators. The notation \overline{K} refers to the set of indices *not* in K , so that K, \overline{K} is a partition of $\{1, \dots, m\}$.

Definition 10.1 (No-signaling operation). A quantum operation $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots m})$ is an m -party no-signaling operation if for each index set $K \subseteq \{1, \dots, m\}$ we have

$$\mathrm{Tr}_{\mathcal{A}_K}(\Lambda(\rho)) = \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(\sigma))$$

whenever

$$\mathrm{Tr}_{\mathcal{X}_K}(\rho) = \mathrm{Tr}_{\mathcal{X}_K}(\sigma).$$

□

What follows is a brief argument that Definition 10.1 captures the meaning of a no-signaling operation—a more detailed discussion of this condition can be found in Beckman *et al.* [BGNP01]. If Λ is no-signaling and ρ, σ are locally indistinguishable to a coalition K of parties (for example, when $\mathrm{Tr}_{\mathcal{X}_K}(\rho) = \mathrm{Tr}_{\mathcal{X}_K}(\sigma)$) then clearly the members of K cannot perform a measurement on their portion of the output that might allow them to distinguish $\Lambda(\rho)$ from $\Lambda(\sigma)$ (that is, $\mathrm{Tr}_{\mathcal{A}_{\bar{K}}}(\Lambda(\rho)) = \mathrm{Tr}_{\mathcal{A}_{\bar{K}}}(\Lambda(\sigma))$). For otherwise, the coalition K would have extracted information—a signal—from the other parties that would allow it to distinguish ρ from σ .

Conversely, if there exist input states ρ, σ such that $\mathrm{Tr}_{\mathcal{X}_K}(\rho) = \mathrm{Tr}_{\mathcal{X}_K}(\sigma)$ and yet $\mathrm{Tr}_{\mathcal{A}_{\bar{K}}}(\Lambda(\rho)) \neq \mathrm{Tr}_{\mathcal{A}_{\bar{K}}}(\Lambda(\sigma))$ then there exists a measurement that allows the coalition K to distinguish these two output states with nonzero bias, which implies that signaling must have occurred.

It is not hard to see that every LOSE operation is also a no-signaling operation. Conversely, much has been made of the fact that there exist no-signaling operations that are not LOSE operations—this is so-called “super-strong” nonlocality, exemplified by the popular “nonlocal box” discussed in Section 10.3.

10.2 Two characterizations of no-signaling operations

In this section it is shown that the product space $\mathbf{Q}_{1\dots m}$ is spanned by Choi-Jamiołkowski representations of no-signaling operations. (Recall from Definition 6.2 that each $\mathbf{Q}_i \subset \mathbf{H}(\mathcal{A}_i \otimes \mathcal{X}_i)$ denotes the subspace of Hermitian operators $J(\Phi)$ for which $\Phi : \mathbf{L}(\mathcal{X}_i) \rightarrow \mathbf{L}(\mathcal{A}_i)$ is a trace-preserving super-operator, or a scalar multiple thereof.) Indeed, that fact is a corollary of the following characterizations of no-signaling operations.

Theorem 10.2 (Two characterizations of no-signaling operations). *Let $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots m}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots m})$ be a quantum operation. The following are equivalent:*

1. Λ is a no-signaling operation.
2. $J(\Lambda)$ is an element of $\mathbf{Q}_{1\dots m}$.

3. For each index set $K \subseteq \{1, \dots, m\}$ there exists an operator $Q \in \mathbf{H}^+(\mathcal{A}_{\overline{K}} \otimes \mathcal{X}_{\overline{K}})$ with $\text{Tr}_{\mathcal{A}_K}(J(\Lambda)) = Q \otimes I_{\mathcal{X}_K}$.

Remark 10.2.1. Membership in the set of no-signaling operations may be verified in polynomial time by checking the linear constraints in Item 3 of Theorem 10.2. While the number of such constraints grows exponentially with m , this exponential growth is not a problem because the number of parties m is always $O(\log n)$ for $n = \dim(\mathbf{Q}_{1\dots m})$. (This logarithmic bound follows from the fact that each space \mathbf{Q}_i has dimension at least two and the total dimension n is the product of the dimensions of each of the m different spaces.)

The partial trace condition of Item 3 of Theorem 10.2 is quite plainly suggested by Theorem 3.2 (Characterization of strategies). Moreover, essential components of the proofs presented for two of the three implications claimed in Theorem 10.2 appear in a 2001 paper of Beckman *et al.* [BGNP01]. The following theorem, however, establishes the third implication and appears to be new.

Theorem 10.3. *A Hermitian operator $X \in \mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$ is in $\mathbf{Q}_{1\dots m}$ if and only if for each index set $K \subseteq \{1, \dots, m\}$ there exists a Hermitian operator $Q \in \mathbf{H}(\mathcal{A}_{\overline{K}} \otimes \mathcal{X}_{\overline{K}})$ with $\text{Tr}_{\mathcal{A}_K}(X) = Q \otimes I_{\mathcal{X}_K}$.*

Proof. The “only if” portion of the theorem is straightforward—only the “if” portion is proven here. The proof proceeds by induction on m . The base case $m = 1$ is trivial. Proceeding directly to the general case, let $s = \dim(\mathbf{H}(\mathcal{A}_{m+1} \otimes \mathcal{X}_{m+1}))$ and let $\{E_1, \dots, E_s\}$ be a basis of $\mathbf{H}(\mathcal{A}_{m+1} \otimes \mathcal{X}_{m+1})$. Let $X_1, \dots, X_s \in \mathbf{H}(\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m})$ be the unique operators satisfying

$$X = \sum_{j=1}^s X_j \otimes E_j.$$

It shall be proven that $X_1, \dots, X_s \in \mathbf{Q}_{1\dots m}$. The intuitive idea is to exploit the linear independence of E_1, \dots, E_s in order to “peel off” individual product terms in the decomposition of X .

Toward that end, for each fixed index $j \in \{1, \dots, s\}$ let H_j be a Hermitian operator for which the real number $\langle H_j, E_i \rangle$ is nonzero only when $i = j$. Define a linear functional $\varphi_j : E \mapsto \langle H_j, E \rangle$ and note that

$$(\mathbb{1}_{\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}} \otimes \varphi_j)(X) = \sum_{i=1}^s \varphi_j(E_i) X_i = \varphi_j(E_j) X_j.$$

Fix an arbitrary partition K, \overline{K} of the index set $\{1, \dots, m\}$. By assumption, $\text{Tr}_{\mathcal{A}_K}(X) = Q \otimes I_{\mathcal{X}_K}$ for some Hermitian operator Q . Apply $\text{Tr}_{\mathcal{A}_K}$ to both sides of the above identity, then use the fact that $\text{Tr}_{\mathcal{A}_K}$ and φ_j act upon different spaces to

obtain

$$\begin{aligned}
& \varphi_j(E_j) \operatorname{Tr}_{\mathcal{A}_K}(X_j) \\
&= \operatorname{Tr}_{\mathcal{A}_K}((\mathbb{1}_{\mathcal{A}_{1\dots m} \otimes \mathcal{X}_{1\dots m}} \otimes \varphi_j)(X)) \\
&= (\mathbb{1}_{\mathcal{A}_{\overline{K}} \otimes \mathcal{X}_{1\dots n}} \otimes \varphi_j)(\operatorname{Tr}_{\mathcal{A}_K}(X)) \\
&= (\mathbb{1}_{\mathcal{A}_{\overline{K}} \otimes \mathcal{X}_{\overline{K}}} \otimes \varphi_j)(Q) \otimes I_{\mathcal{X}_K},
\end{aligned}$$

from which it follows that $\operatorname{Tr}_{\mathcal{A}_K}(X_j)$ is a product operator of the form $R \otimes I_{\mathcal{X}_K}$ for some Hermitian operator R . As this identity holds for all index sets K , it follows from the induction hypothesis that $X_j \in \mathbf{Q}_{1\dots m}$ as desired.

Now, choose a maximal linearly independent subset $\{X_1, \dots, X_t\}$ of $\{X_1, \dots, X_s\}$ and let Y_1, \dots, Y_t be the unique Hermitian operators satisfying

$$X = \sum_{i=1}^t X_i \otimes Y_i.$$

A similar argument shows $Y_1, \dots, Y_t \in \mathbf{Q}_{m+1}$, which completes the induction. \square

Proof of Theorem 10.2. Item 3 implies item 2. This implication follows immediately from Theorem 10.3.

Item 2 implies item 1. The proof of this implication borrows heavily from the proof of Theorem 2 in Beckman *et al.* [BGNP01].

Fix any partition K, \overline{K} of the index set $\{1, \dots, m\}$. Let $s = \dim(\mathbf{L}(\mathcal{X}_{\overline{K}}))$ and $t = \dim(\mathbf{L}(\mathcal{X}_K))$ and let $\{\rho_1, \dots, \rho_s\}$ and $\{\sigma_1, \dots, \sigma_t\}$ be bases of $\mathbf{L}(\mathcal{X}_{\overline{K}})$ and $\mathbf{L}(\mathcal{X}_K)$, respectively, that consist entirely of density operators. Given any two operators $X, Y \in \mathbf{L}(\mathcal{X}_{1\dots m})$ let $x_{j,k}, y_{j,k} \in \mathbb{C}$ be the unique coefficients of X and Y respectively in the product basis $\{\rho_j \otimes \sigma_k\}$. Then $\operatorname{Tr}_{\mathcal{X}_K}(X) = \operatorname{Tr}_{\mathcal{X}_K}(Y)$ implies

$$\sum_{k=1}^t x_{j,k} = \sum_{k=1}^t y_{j,k}$$

for each fixed index $j = 1, \dots, s$.

As $J(\Lambda) \in \mathbf{Q}_{1\dots m}$, it is possible to write

$$J(\Lambda) = \sum_{l=1}^n J(\Phi_{1,l}) \otimes \dots \otimes J(\Phi_{m,l})$$

where n is a positive integer and $\Phi_{i,l} : \mathbf{L}(\mathcal{X}_i) \rightarrow \mathbf{L}(\mathcal{A}_i)$ satisfies $J(\Phi_{i,l}) \in \mathbf{Q}_i$ for each of the indices $i = 1, \dots, m$ and $l = 1, \dots, n$. In particular, as each $\Phi_{i,l}$ is (a scalar multiple of) a trace-preserving super-operator, it holds that for each index $l = 1, \dots, n$ there exists $a_l \in \mathbb{R}$ with $\operatorname{Tr}(\Phi_{K,l}(\sigma)) = a_l$ for all

density operators σ . Then

$$\begin{aligned}\mathrm{Tr}_{\mathcal{A}_K}(\Lambda(X)) &= \sum_{l=1}^n a_l \cdot \sum_{j=1}^s \left(\sum_{k=1}^t x_{j,k} \right) \cdot \Phi_{\overline{K},l}(\rho_j) \\ &= \sum_{l=1}^n a_l \cdot \sum_{j=1}^s \left(\sum_{k=1}^t y_{j,k} \right) \cdot \Phi_{\overline{K},l}(\rho_j) = \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(Y))\end{aligned}$$

as desired.

Item 1 implies item 3. This implication is essentially a multi-party generalization of Theorem 8 in Beckman *et al.* [BGNP01]. The proof presented here differs from theirs in some interesting but non-critical details.

Fix any partition K, \overline{K} of the index set $\{1, \dots, m\}$. To begin, observe that

$$\mathrm{Tr}_{\mathcal{X}_K}(X) = \mathrm{Tr}_{\mathcal{X}_K}(Y) \implies \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(X)) = \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(Y))$$

for *all* operators $X, Y \in \mathbf{L}(\mathcal{X}_{1\dots m})$ —not just density operators. (This observation follows from the fact that $\mathbf{L}(\mathcal{X}_{1\dots m})$ is spanned by the density operators—a fact used in the above proof that item 2 implies item 1.)

Now, let $s = \dim(\mathcal{X}_{\overline{K}})$ and $t = \dim(\mathcal{X}_K)$ and let $\{e_1, \dots, e_s\}$ and $\{f_1, \dots, f_t\}$ be the standard bases of $\mathcal{X}_{\overline{K}}$ and \mathcal{X}_K respectively. If c and d are distinct indices in $\{1, \dots, t\}$ and $Z \in \mathbf{L}(\mathcal{X}_{\overline{K}})$ is any operator then

$$\mathrm{Tr}_{\mathcal{X}_K}(Z \otimes f_c f_d^*) = Z \otimes \mathrm{Tr}(f_c f_d^*) = 0_{\mathcal{X}_{\overline{K}}} = \mathrm{Tr}_{\mathcal{X}_K}(0_{\mathcal{X}_{1\dots m}})$$

and hence

$$\mathrm{Tr}_{\mathcal{A}_K}(\Lambda(Z \otimes f_c f_d^*)) = \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(0_{\mathcal{X}_{1\dots m}})) = \mathrm{Tr}_{\mathcal{A}_K}(0_{\mathcal{A}_{1\dots m}}) = 0_{\mathcal{A}_{\overline{K}}}.$$

(Here a natural notation for the zero operator was used implicitly.) Similarly, if $\rho \in \mathbf{L}(\mathcal{X}_K)$ is any density operator then

$$\mathrm{Tr}_{\mathcal{A}_K}(\Lambda(Z \otimes f_c f_c^*)) = \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(Z \otimes \rho))$$

for each fixed index $c = 1, \dots, t$. Employing these two identities, one obtains

$$\begin{aligned}\mathrm{Tr}_{\mathcal{A}_K}(J(\Lambda)) &= \sum_{a,b=1}^s \sum_{c=1}^t \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(e_a e_b^* \otimes f_c f_c^*)) \otimes (e_a e_b^* \otimes f_c f_c^*) \\ &= \sum_{a,b=1}^s \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(e_a e_b^* \otimes \rho)) \otimes e_a e_b^* \otimes \left(\sum_{c=1}^t f_c f_c^* \right) = J(\Psi) \otimes I_{\mathcal{X}_K}\end{aligned}$$

where the quantum operation Ψ is defined by $\Psi : X \mapsto \mathrm{Tr}_{\mathcal{A}_K}(\Lambda(X \otimes \rho))$. As $J(\Psi) \otimes I_{\mathcal{X}_K}$ is a product operator of the desired form, the proof that item 1 implies item 3 is complete. □

10.3 A separable no-signaling operation that is not a LOSE operation

Theorem 7.4 establishes a ball of LOSR operations around the completely noisy channel. Was the work of Chapter 7 necessary, or might there be a simpler way to establish the same thing? For example, suppose $\Lambda : \mathbf{L}(\mathcal{X}_1 \otimes \mathcal{X}_2) \rightarrow \mathbf{L}(\mathcal{A}_1 \otimes \mathcal{A}_2)$ is a quantum operation for which the operator $J(\Lambda)$ is $(\mathbf{H}(\mathcal{A}_1 \otimes \mathcal{X}_1); \mathbf{H}(\mathcal{A}_2 \otimes \mathcal{X}_2))$ -separable. Quantum operations with separable Choi-Jamiołkowski representations such as this are called *separable operations* [Rai97]. If an operation is both separable and no-signaling then must it always be a LOSE operation, or even a LOSR operation? An affirmative answer to this question would yield a trivial proof of Theorem 7.4 that leverages existing knowledge of separable balls around the identity operator.

Alas, such a short cut is not to be had: there exist no-signaling operations Λ that are not LOSE operations, yet $J(\Lambda)$ is separable. One example of such an operation is the so-called “nonlocal box” discovered in 1994 by Popescu and Rohrlich [PR94]. This nonlocal box is easily formalized as a two-party no-signaling quantum operation Λ , as in Ref. [BGNP01, Section V.B]. That formalization is reproduced here.

Let $\mathcal{X}_1 = \mathcal{X}_2 = \mathcal{A}_1 = \mathcal{A}_2 = \mathbb{C}^2$, let $\{e_0, e_1\}$ denote the standard bases of both \mathcal{X}_1 and \mathcal{A}_1 , and let $\{f_0, f_1\}$ denote the standard bases of both \mathcal{X}_2 and \mathcal{A}_2 . Write

$$\rho_{ab} \stackrel{\text{def}}{=} e_a e_a^* \otimes f_b f_b^*$$

for $a, b \in \{0, 1\}$. The nonlocal box $\Lambda : \mathbf{L}(\mathcal{X}_{1\dots 2}) \rightarrow \mathbf{L}(\mathcal{A}_{1\dots 2})$ is defined by

$$\begin{aligned} \{\rho_{00}, \rho_{01}, \rho_{10}\} &\xrightarrow{\Lambda} \frac{1}{2}(\rho_{00} + \rho_{11}) \\ \rho_{11} &\xrightarrow{\Lambda} \frac{1}{2}(\rho_{01} + \rho_{10}). \end{aligned}$$

Operators not in $\text{span}\{\rho_{00}, \rho_{01}, \rho_{10}, \rho_{11}\}$ are annihilated by Λ . It is routine to verify that Λ is a no-signaling operation, and this operation Λ is known not to be a LOSE operation [PR94]. To see that $J(\Lambda)$ is separable, write

$$\begin{aligned} E_{a \rightarrow b} &\stackrel{\text{def}}{=} e_b e_a^* \\ F_{a \rightarrow b} &\stackrel{\text{def}}{=} f_b f_a^* \end{aligned}$$

for $a, b \in \{0, 1\}$. Then for all $X \in \mathbf{L}(\mathcal{X}_{1\dots 2})$ it holds that

$$\begin{aligned} \Lambda(X) &= \frac{1}{2} [E_{0 \rightarrow 0} \otimes F_{0 \rightarrow 0}] X [E_{0 \rightarrow 0} \otimes F_{0 \rightarrow 0}]^* + \frac{1}{2} [E_{0 \rightarrow 1} \otimes F_{0 \rightarrow 1}] X [E_{0 \rightarrow 1} \otimes F_{0 \rightarrow 1}]^* \\ &\quad + \frac{1}{2} [E_{0 \rightarrow 0} \otimes F_{1 \rightarrow 0}] X [E_{0 \rightarrow 0} \otimes F_{1 \rightarrow 0}]^* + \frac{1}{2} [E_{0 \rightarrow 1} \otimes F_{1 \rightarrow 1}] X [E_{0 \rightarrow 1} \otimes F_{1 \rightarrow 1}]^* \\ &\quad + \frac{1}{2} [E_{1 \rightarrow 0} \otimes F_{0 \rightarrow 0}] X [E_{1 \rightarrow 0} \otimes F_{0 \rightarrow 0}]^* + \frac{1}{2} [E_{1 \rightarrow 1} \otimes F_{0 \rightarrow 1}] X [E_{1 \rightarrow 1} \otimes F_{0 \rightarrow 1}]^* \\ &\quad + \frac{1}{2} [E_{1 \rightarrow 0} \otimes F_{1 \rightarrow 1}] X [E_{1 \rightarrow 0} \otimes F_{1 \rightarrow 1}]^* + \frac{1}{2} [E_{1 \rightarrow 1} \otimes F_{1 \rightarrow 0}] X [E_{1 \rightarrow 1} \otimes F_{1 \rightarrow 0}]^*, \end{aligned}$$

from which the $(\mathbf{H}(\mathcal{A}_1 \otimes \mathcal{X}_1); \mathbf{H}(\mathcal{A}_2 \otimes \mathcal{X}_2))$ -separability of $J(\Lambda)$ follows. It is interesting to note that the nonlocal box is the smallest possible nontrivial example of such an operation—the number of parties $m = 2$ and the input spaces $\mathcal{X}_1, \mathcal{X}_2$ and output spaces $\mathcal{A}_1, \mathcal{A}_2$ all have dimension two.

Chapter 11

Conclusion

In this thesis we discussed two distinct topics. In Part I we initiated the study of quantum strategies, which are complete descriptions of one party’s actions in an interaction involving the exchange of multiple quantum messages among multiple parties. We saw proofs of three important properties of strategies, and we saw applications of these properties to zero-sum quantum games, complexity theory, and quantum coin-flipping. We also introduced a new norm for super-operators and argued that this norm, which generalizes the familiar diamond norm, captures the operational distinguishability of two quantum strategies in the same sense that the trace norm captures the distinguishability of two quantum states, or the diamond norm captures the distinguishability of two quantum operations.

In Part II we established several properties of local quantum operations, the implementation of which might be assisted by shared entanglement. Specifically, we showed that every quantum operation sufficiently close to the completely noisy operation can be implemented locally using only shared randomness. This fact was used to prove strong NP-hardness of the weak membership problem for local operations with shared entanglement. We then provided algebraic characterizations of the sets of local operations with shared randomness and entanglement in terms of linear functionals that are positive and “completely” positive, respectively, on a certain cone of separable Hermitian operators. Finally, we made explicit for the first time two fundamental characterizations of no-signaling operations, establishing that the spaces spanned by the local operations and by the larger class of no-signaling operations are in fact equal.

We conclude the thesis with some pointers for future research and open problems pertaining to the topics covered.

11.1 Quantum strategies

New properties. Three important properties of quantum strategies were established in Theorems 3.1 (Interaction output probabilities), 3.2 (Characteriza-

tion of strategies), and 3.3 (Maximum output probability). Other simpler and more basic properties were noted in Propositions 2.7, 3.15, 3.16, and 3.17. We also saw several properties of a new distance measure for quantum strategies in Chapter 5. What other properties are held by our representation for quantum strategies?

Simplifying proofs. In Section 4.3 we employed the properties of quantum strategies to provide a simplified proof of Kitaev’s bound for strong quantum coin-flipping. We also noted in Section 1.1.1 that Chiribella *et al.* have provided a short proof by quantum strategies of the impossibility of quantum bit commitment [CDP⁺09b].

Given these examples, one is tempted to believe that the properties of quantum strategies encapsulate many of the critical elements of various proofs involving the exchange of quantum information. It is reasonable to expect, for example, that our formalism could lead to new or alternate security proofs for various quantum cryptographic protocols, or possibly even to proofs that certain classical protocols are secure against quantum attacks.

For an unsolved example, consider the protocol for weak quantum coin-flipping with arbitrarily small bias given in Ref. [Moc07]. The proof that the exhibited protocol has arbitrarily small bias is very complicated. Could this proof be simplified by quantum strategies?

New applications. Of course, the formalism of quantum strategies is by no means limited to simplifications of existing proofs, as illustrated by the new results established in Sections 4.1, 4.2, and 8.2, and by Chiribella, D’Ariano, Perinotti, and other authors as discussed in Section 1.1.1. Surely, there is more to add to this list of new applications.

11.2 Local operations with shared entanglement

Bigger ball of LOSE or LOSR operations. The size of the ball of (unnormalized) LOSR operations established in Theorem 7.4 scales as $\Omega(2^{-m}n^{-3/2})$. Given that this quantity already includes a factor of the dimension n , is it possible to eliminate the explicit dependence on the number of parties m ? (By contrast, for the case of multipartite separable quantum states the dependence on m seems unavoidable [GB03].) Can the exponent on the dimension n be improved?

As mentioned in Section 7.2, it is not clear that there is a ball of LOSE operations that strictly contains any ball of LOSR operations. Does such a larger ball exist?

Completely positive super-operators. As noted in Remark 9.3.1, the characterization of LOSE operations is interesting because it involves linear func-

tionals that are not just positive, but “completely” positive on the family of $(\mathbf{Q}_1(\mathcal{E}_1); \dots; \mathbf{Q}_m(\mathcal{E}_m))$ -separable cones.

Apparently, the study of completely positive super-operators has until now been strictly limited to the context of positive semidefinite input cones. Any question that may be asked of conventional completely positive super-operators might also be asked of this new class of completely positive super-operators.

Entanglement required for approximating LOSE operations. It was mentioned in Chapters 1 and 6 that there exist LOSE operations that cannot be implemented with any finite amount of shared entanglement [LTW08]. The natural question, then, is how much entanglement is necessary to achieve an arbitrarily close approximation to such an operation?

The present author conjectures that for every two-party LOSE operation Λ there exists an ε -approximation Λ' of Λ in which the dimension of the shared entangled state scales as $O(2^{\varepsilon^{-a}} n^b)$ for some positive constants a and b and some appropriate notion of ε -approximation. Here $n = \dim(\mathbf{Q}_{1\dots m})$ is the dimension of the space in which $J(\Lambda)$ lies.

Evidence pertaining to this conjecture can be found in Refs. [CHTW04, KRT08, LTW08]. Moreover, the example in Ref. [LTW08] strongly suggests that the exponential dependence on $1/\varepsilon$ is unavoidable; the pressing open question pertains to the dependence upon n . At the moment, *no upper bound at all* is known for this general class of two-party LOSE operations.

References

- [ABDR04] Andris Ambainis, Harry Buhrman, Yevgeniy Dodis, and Hein Röhrig. Multipart quantum coin flipping. In *Proceedings of the 19th Annual Conference on Computational Complexity*, pages 250–259, 2004. arXiv:quant-ph/0304112v2. 79
- [AHW00] Grigori Amosov, Alexander Holevo, and Reinhard Werner. On some additivity problems in quantum information theory. *Problems of Information Transmission*, 36(4):25–34, 2000. 96
- [AKN98] Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pages 20–30, 1998. arXiv:quant-ph/9806029v1. 27, 83
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998. 6
- [AS98] Sanjeev Arora and Muli Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998. 6
- [Bab85] László Babai. Trading group theory for randomness. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pages 421–429, 1985. 73
- [Bar02] Alexander Barvinok. *A Course in Convexity*, volume 54 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002. 20
- [BCD⁺09a] Alessandro Bisio, Giulio Chiribella, Giacomo Mauro D’Ariano, Stefano Facchini, and Paolo Perinotti. Optimal quantum learning of a unitary transformation. arXiv:0903.0543v1 [quant-ph], 2009. 4
- [BCD⁺09b] Alessandro Bisio, Giulio Chiribella, Giacomo Mauro D’Ariano, Stefano Facchini, and Paolo Perinotti. Optimal quantum tomography for states, measurements, and transformations. *Physical Review Letters*, 102:010404, 2009. arXiv:0806.1172v1 [quant-ph]. 4

- [BGNP01] David Beckman, Daniel Gottesman, Michael Nielsen, and John Preskill. Causal and localizable quantum operations. *Physical Review A*, 64(5):052309, 2001. arXiv:quant-ph/0102043v2. 6, 9, 44, 127, 128, 129, 130, 131
- [Bha07] Rajendra Bhatia. *Positive Definite Matrices*. Princeton University Press, 2007. 110, 113, 114, 124
- [Blu81] Manuel Blum. Coin flipping by telephone. In *Advances in Cryptology—CRYPTO '81*, pages 11–15, 1981. 79
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. 69
- [CDP08a] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Memory effects in quantum channel discrimination. *Physical Review Letters*, 101:180501, 2008. arXiv:0803.3237v3 [quant-ph]. 4
- [CDP08b] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Optimal cloning of unitary transformation. *Physical Review Letters*, 101:180504, 2008. arXiv:0804.0129v2 [quant-ph]. 4
- [CDP08c] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Optimal cloning of unitary transformation. *Europhysics Letters*, 83:30004, 2008. arXiv:0804.0180v2 [quant-ph]. 4
- [CDP08d] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Optimal covariant quantum networks. In Alexander Lvovsky, editor, *Proceedings of the Ninth International Conference on Quantum Communication, Measurement and Computing (QCMC '08)*, pages 47–56. American Institute of Physics, 2008. arXiv:0812.3922v1 [quant-ph]. 4
- [CDP08e] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Quantum circuits architecture. *Physical Review Letters*, 101:060401, 2008. arXiv:0712.1325v1 [quant-ph]. 4
- [CDP09a] Giulio Chiribella, Giacomo Mauro D’Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80:022339, 2009. arXiv:0904.4483v1 [quant-ph]. 4
- [CDP⁺09b] Giulio Chiribella, Giacomo Mauro D’Ariano, Paolo Perinotti, Dirk Schlingemann, and Reinhard F. Werner. A short impossibility proof of quantum bit commitment. arXiv:0905.3801v1 [quant-ph], 2009. 4, 134
- [CGJ09] Richard Cleve, Dmitry Gavinsky, and Rahul Jain. Entanglement-resistant two-prover interactive proof systems and non-adaptive PIRs. *Quantum Information and Computation*, 9:648–656, 2009. arXiv:0707.1729v1 [quant-ph]. 6

- [CHTW04] Richard Cleve, Peter Høyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th Annual Conference on Computational Complexity*, pages 236–249, 2004. arXiv:quant-ph/0404076v1. 6, 135
- [CK09] André Chailloux and Iordanis Kerenidis. Optimal quantum strong coin flipping. arXiv:0904.1511v1 [quant-ph], 2009. 79
- [CSUU06] Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Strong parallel repetition theorem for quantum XOR proof systems. arXiv:quant-ph/0608146v1, 2006. 76
- [CSUU08] Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, pages 282–299, 2008. arXiv:quant-ph/0608146v2. 6, 76
- [dK02] Etienne de Klerk. *Aspects of Semidefinite Programming: Interior Point Algorithms and Selected Applications*, volume 65 of *Applied Optimization*. Kluwer, 2002. 69
- [DLTW08] Andrew Doherty, Yeong-Cherng Liang, Ben Toner, and Stephanie Wehner. The quantum moment problem and bounds on entangled multi-prover games. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity*, pages 199–210, 2008. arXiv:0803.4373v1 [quant-ph]. 6
- [ESW02] Tilo Eggeling, Dirk Schlingemann, and Reinhard Werner. Semicausal operations are semilocalizable. *Europhysics Letters*, 57(6):782–788, 2002. arXiv:quant-ph/0104027v1. 6, 44
- [FK97] Uriel Feige and Joe Kilian. Making games short. In *Proceedings of the Twenty-Ninth annual ACM Symposium on Theory of Computing*, pages 506–516, 1997. 74
- [GB02] Leonid Gurvits and Howard Barnum. Largest separable balls around the maximally mixed bipartite quantum state. *Physical Review A*, 66(6):062311, 2002. arXiv:quant-ph/0204159v2. 6, 110
- [GB03] Leonid Gurvits and Howard Barnum. Separable balls around the maximally mixed multipartite quantum states. *Physical Review A*, 68(4):042312, 2003. arXiv:quant-ph/0302102v1. 7, 110, 113, 134
- [GB05] Leonid Gurvits and Howard Barnum. Better bound on the exponent of the radius of the multipartite separable ball. *Physical Review A*, 72:032322, 2005. arXiv:quant-ph/0409095v4. 7
- [Gha08] Sevag Gharibian. Strong NP-hardness of the quantum separability problem, 2008. arXiv:0810.4507v2 [quant-ph]. 8, 120

- [GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer–Verlag, 1988. 8, 68, 120
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989. 73
- [Gur02] Leonid Gurvits. Quantum matching theory (with new complexity theoretic, combinatorial and topological insights on the nature of quantum entanglement), 2002. arXiv:quant-ph/0201022v2. 6, 8, 120
- [Gut09] Gus Gutoski. Properties of local quantum operations with shared entanglement. *Quantum Information and Computation*, 9(9):739–764, 2009. arXiv:0805.2209v3 [quant-ph]. 1
- [GW05] Gus Gutoski and John Watrous. Quantum interactive proofs with competing provers. In *Proceedings of the 22nd Symposium on Theoretical Aspects of Computer Science (STACS’05)*, volume 3404 of *Lecture Notes in Computer Science*, pages 605–616. Springer, 2005. arXiv:cs/0412102v1 [cs.CC]. 5, 78, 81, 89, 90, 93
- [GW07] Gus Gutoski and John Watrous. Toward a general theory of quantum games. In *Proceedings of the 39th ACM Symposium on Theory of Computing*, pages 565–574, 2007. arXiv:quant-ph/0611234v2. 1, 44, 60
- [Hel69] Carl Helstrom. Quantum detection and estimation theory. *Journal of Statistical Physics*, 1(2):231–252, 1969. 82
- [HHH96] Michal Horodecki, Pawel Horodecki, and Ryszard Horodecki. Separability of mixed states: necessary and sufficient conditions. *Physical Letters A*, 223:1–8, 1996. arXiv:quant-ph/9605038v2. 8, 122
- [HHH01] Michal Horodecki, Pawel Horodecki, and Ryszard Horodecki. Separability of n -particle mixed states: necessary and sufficient conditions in terms of linear maps. *Physical Letters A*, 283:1–7, 2001. arXiv:quant-ph/0006071v1. 9, 122
- [Hil05] Roland Hildebrand. Cones of ball-ball separable elements, 2005. arXiv:quant-ph/0503194v1. 7
- [HJ85] Roger Horn and Charles Johnson. *Matrix Analysis*. Cambridge University Press, 1985. 95
- [IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal

- strategies. In *Proceedings of the 24th Annual Conference on Computational Complexity*, pages 217–228, 2009. arXiv:0810.0693v1 [quant-ph]. 6
- [JJUW09] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. arXiv:0907.4737v2 [quant-ph], 2009. 74
- [Kha79] Leonid Khachiyan. A polynomial time algorithm in linear programming. *Soviet Mathematics Doklady*, 20:191–194, 1979. 68
- [Kit97] Alexei Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191–1249, 1997. 83
- [Kit02] Alexei Kitaev. Quantum coin-flipping. Presentation at the 6th Workshop on *Quantum Information Processing (QIP 2003)*, 2002. 79
- [KKM⁺08] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. Entangled games are hard to approximate, 2008. arXiv:0704.2903v2 [quant-ph]. 6, 8, 119
- [KKMV08] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, and Thomas Vidick. Using entanglement in quantum multi-prover interactive proofs. In *Proceedings of the 23rd Annual Conference on Computational Complexity*, pages 211–222, 2008. arXiv:0711.3715v1 [quant-ph]. 6
- [KM92] Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4:528–552, 1992. 74
- [KM03] Hirotada Kobayashi and Keiji Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *Journal of Computer and System Sciences*, 66(3):429–450, 2003. arXiv:cs/0102013v5 [cs.CC]. 6
- [KRT08] Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. In *Proceedings of the 49th Annual Symposium on Foundations of Computer Science*, pages 457–466, 2008. arXiv:0710.0655v2 [quant-ph]. 6, 135
- [KSV02] Alexei Kitaev, Alexander Shen, and Mikhail Vyalyi. *Classical and Quantum Computation*, volume 47 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002. 89, 98
- [KW00] Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof system. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, pages 608–617, 2000. 74, 76

- [LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992. 74
- [Liu07] Yi-Kai Liu. *The Complexity of the Consistency and N -representability Problems for Quantum States*. PhD thesis, University of California, San Diego, 2007. arXiv:0712.3041v1 [quant-ph]. 8, 120
- [LJ03] Chiu Fan Lee and Neil Johnson. Efficiency and formalism of quantum games. *Physical Review A*, 67:article 022311, 2003. arXiv:quant-ph/0207012v4. 61
- [LTW08] Debbie Leung, Ben Toner, and John Watrous. Coherent state exchange in multi-prover quantum interactive proof systems. arXiv:0804.4118v1 [quant-ph], 2008. 6, 106, 135
- [Moc07] Carlos Mochon. Quantum weak coin-flipping with arbitrarily small bias, 2007. arXiv:0711.4114v1 [quant-ph]. 79, 134
- [MV99] Peter Bro Miltersen and N. Variyam Vinodchandran. Derandomizing Arthur-Merlin games using hitting sets. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, pages 71–80, 1999. 74
- [NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10:073013, 2008. arXiv:0803.4290v1 [quant-ph]. 6
- [PHHH06] Marco Piani, Michal Horodecki, Pawel Horodecki, and Ryszard Horodecki. Properties of quantum non-signalling boxes. *Physical Review A*, 74:012305, 2006. arXiv:quant-ph/0505110v1. 6
- [PR94] Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994. 9, 131
- [Rai97] Eric Rains. Entanglement purification via separable superoperators, 1997. arXiv:quant-ph/9707002v3. 9, 131
- [RD66] B. Russo and H. A. Dye. A note on unitary operators in C^* -algebras. *Duke Mathematical Journal*, 33:413–416, 1966. 113
- [Roc70] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970. 20, 21, 50, 62
- [Röh04] Hein Röhrig. *Quantum Query Complexity and Distributed Computing*. PhD thesis, Centrum voor Wiskunde en Informatica, 2004. 79

- [RW05] Bill Rosgen and John Watrous. On the hardness of distinguishing mixed-state quantum computations. In *Proceedings of the 20th Annual Conference on Computational Complexity*, pages 344–354, 2005. arXiv:cs/0407056v1 [cs.CC]. 83, 89, 96
- [Sha92] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, 1992. 74
- [Sza05] Stanislaw Szarek. The volume of separable states is super-doubly-exponentially small. *Physical Review A*, 72:032304, 2005. arXiv:quant-ph/0310061v2. 7
- [Vil38] Jean Ville. Sur la théorie générale des jeux où intervient l’habileté des joueurs. *Traité du calcul des probabilités et des applications*, IV(2):105–113, 1938. In French. 62
- [vN28] John von Neumann. Zur theorie der gesellschaftspiele. *Mathematische Annalen*, 100(1):295–320, 1928. In German. 61
- [Wat05] John Watrous. Notes on super-operator norms induced by Schatten norms. *Quantum Information and Computation*, 5(1):58–68, 2005. arXiv:quant-ph/0411077v1. 97
- [Wat08] John Watrous. Distinguishing quantum operations having few Kraus operators. *Quantum Information and Computation*, 8(9):819–833, 2008. arXiv:0710.0902v3 [quant-ph]. 84
- [Wat09] John Watrous. Semidefinite programs for completely bounded norms. *Theory of Computing*, 5:217–238, 2009. arXiv:0901.4709v2 [quant-ph]. 53, 69
- [Weh06] Stephanie Wehner. Entanglement in interactive proof systems with binary answers. In *Proceedings of the 23rd Annual Symposium on Theoretical Aspects of Computer Science*, volume 3884 of *Lecture Notes in Computer Science*, pages 162–171. Springer, 2006. arXiv:quant-ph/0508201v2. 6
- [YN76] David Yudin and Arkadi Nemirovskii. Informational complexity and efficient methods for the solution of convex extremal problems. *Èkonomika i Matematicheskie Metody*, 12:357–369, 1976. 8, 120