Discrimination of Quantum Channels

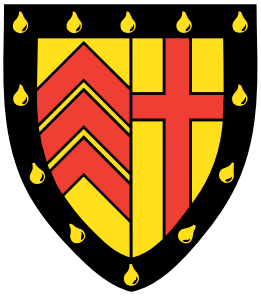

Bjarne Finn Bergh

Clare College, University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

September 2024

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted, or, is being concurrently submitted, for any degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Abstract

This thesis deals with assessing optimal decision strategies and error rates for the task of binary quantum channel discrimination, where the main focus lies on proving optimal asymmetric error exponents (both asymptotically and non-asymptotically) for different variants of this problem.

The task of binary quantum channel discrimination is to distinguish a quantum channel (a quantum device with a quantum input and a quantum output), of which n identical copies of are given, into one of two classes. In the simplest instance the task is just to recognize the channel as one of two possible options. Even when the n copies of the channel are assumed to be independent and identically distributed (i.i.d.), the optimal joint input states will often be entangled, and hence this becomes a highly non-trivial non-i.i.d. hypothesis testing problem, which in certain cases can also be shown to require adaptivity for optimal discrimination performance. This thesis works on exploring the landscape of when adaptivity is helpful, and what can be said about optimal asymmetric error rates in various regimes.

Concretely, first, in the already mentioned simple case of distinguishing between two single options, and for arbitrary (finite) number of channel uses, we prove a relation between the error rates of optimal adaptive and non-adaptive (so called parallel) strategies, with a mostly explicit construction of a "good" parallel strategy (in a certain quantified sense) given an adaptive one. Additionally, we show how optimal parallel strategies and error rates can be computed as a semi-definite program (SDP) that grows only polynomially in n.

Next, we study the problem of composite channel discrimination, where the two classes to distinguish between can now be very general sets of channels, and which to our knowledge has so far not been studied even classically, despite its broad applicability to real-world discrimination tasks. Also here, we establish asymptotic asymmetric error exponents in multiple scenarios, together with some conditions on when they are equal (or not equal) between adaptive and non-adaptive strategies, and hence are able to give conditions under which adaptivity is asymptotically (not) helpful.

Finally, we lift many of these quantum channel discrimination results (and also the information theoretical toolbox they are based on), so far studied only in finite dimensions, to the infinite-dimensional setting of general separable Hilbert spaces, and establish a general theory of quantum channel discrimination also in this setting. For this, we prove (among many other things) the chain rule for the geometric Rényi divergences in infinite dimensions, and then establish relations between error rates of adaptive and parallel strategies similar to the finite-dimensional case. While in finite dimensions we can establish asymptotic equality of the asymmetric error exponent for adaptive and parallel strategies for all pairs of channels, in infinite dimensions we can do this only under an additional finiteness condition, and we discuss this condition in detail.

Preface

The content of this thesis is based on the following four papers:

- Bergh, Salzmann, and Datta (2021), published in collaboration with Robert Salzmann and Nilanjana Datta
- Bergh et al. (2024), published in collaboration with Nilanjana Datta, Robert Salzmann and Mark M. Wilde
- Bergh, Datta, and Salzmann (2023), published in collaboration with Nilanjana Datta and Robert Salzmann
- Bergh et al. (2023), published in collaboration with Jan Kochanowski, Robert Salzmann and Nilanjana Datta

In all of these four papers, I have been the main contributor and author of the paper, and proved all of the key results. Unless explicitly stated otherwise, all proofs included in this thesis are my own and have been obtained independently through original work during my PhD.

Acknowledgements

I would like to deeply thank my supervisor Nilanjana Datta for her dedication, continuous optimism, encouragement, and support during my time as a PhD student. I am particularly thankful for the many hours of discussion, and the constant engagement and availability. I am very grateful for having been able to do a PhD with her, and am very much looking forward to future joint work ahead.

The work in this PhD has also greatly benefited from collaborations with Robert Salzmann, Mark M. Wilde and Jan Kochanowski, all of which I am very much looking forward to working with again. I would like to thank in particular Robert for his constant supply of helpful comments and ideas.

Throughout my PhD I had the pleasure to be able to discuss with Mario Berta, Andreas Bluhm, Josh Cudby, Ángela Capel Cuevas, Marco Fanizza, Hamza Fawzi, Omar Fawzi, Niklas Galke, Lampros Gavalakis, Paul Gondolf, Christoph Hirche, Michael Jabbour, Richard Josza, Ioannis Kontoyiannis, Ludovico Lami, Lauritz van Lujik, Campbell McLauchlan, Milán Mosonyi, Tim Möbus, Mizanur Rahaman, Bartosz Regula, Wilfred Salmon, Samuel O. Scalet, Satvik Singh, Sergii Strelchuk, Peter Wildeman, Albert H. Werner and Andreas Winter, who have in many ways broadened my understanding of various aspects of quantum information theory.

I am also very grateful to my two examiners, Ángela Capel and Milán Mosonyi, whose suggestions have led to some significant improvements of this thesis.

For financial and overall support I am grateful to the Engineering & Physical Sciences Research Council (EPRSC), Cantab Capital Institute for the Mathematics of Information, and Clare College.

My PhD would not have been the same without all the people who have made my time in Cambridge special and enjoyable. There are definitely too many of you to name, but particular thanks go to Capucine, Emily, Ellie, Gaël, Holland, Liv, Megan, Morgan, Pablo and Will. Finally, I would like to thank my family, and in particular my parents, for their unconditional help and support.

Contents

1 Introduction

1.1 The Problem of Quantum Channel Discrimination

One of the most fundamental tasks of statistics and information theory is hypothesis testing, where one has to distinguish between multiple hypotheses based on observed data. This task is ubiquitous in many daily applications: "Does this drug have a measurable effect?", "Are two individual features correlated in the human population?", "How do financial markets react to certain kinds of events?", "Are traces of a fundamental particle visible in experimental data?", just to name a few examples. Generally, hypothesis testing and statistics come in whenever certain observables are random (or sufficiently complex that they can be effectively treated as random), and allow one to say something about the underlying object with a guaranteed (high) probability.

In quantum mechanics, such uncertainty or randomness is inherent; the only way to find out anything about a quantum state is to perform a measurement, whose outcomes will in general be random. Thus, any certification of a quantum process (any answer to the question "Is this quantum thing really what I think it is?"), and most ways of learning something about a quantum device, involve some form of hypothesis testing.

The most prototypical quantum hypothesis testing task is binary quantum state discrimination (Helstrom 1969, Holevo 1972), which involves determining the state of a quantum system, given the side information that it is in one of two possible states. This is done by performing suitable measurements on the state. The task has been studied in both the n-shot regime, in which a finite number (n) of identical copies of the unknown state are available, and in the asymptotic limit in which one assumes that an infinite number of copies are available (i.e., $n \to \infty$). This is also a fundamental primitive of quantum information theory since many other information processing tasks can be reduced to it. Known results for this problem now include optimal decision strategies and the behavior of the error probabilities of misidentification in both these finite and asymptotic regimes (Helstrom 1969, Kholevo 1979, Hiai and Petz 1991, Ogawa and Nagaoka 2000, Nagaoka 2006, Audenaert et al. 2007, Hayashi 2007, Audenaert et al. 2008, Nussbaum and Szkoła 2009, Audenaert, Mosonyi, and Verstraete 2012, Bae and Kwek 2015).

This thesis deals with quantum *channel* discrimination. The basic idea is the same; one wants to distinguish a quantum object between multiple different options, but instead of these objects being quantum states which are then measured, these quantum objects are now quantum channels, i.e. they take an input quantum state before outputting a quantum state to be measured. This might not seem like much of a different problem, since after choosing an input state one gets back a quantum state discrimination problem, and surely there must be a "best" input state to choose. Classically, (at least for the problem where one only has to distinguish between exactly two completely specified channels) this is indeed true, there exists a best input state, and even if one is allowed to use the channel multiple times, there is nothing better to do than to keep using this

input state every time (Hayashi 2009). Quantumly, however this is not true, essentially due to the presence of quantum entanglement, and so an optimal input state for multiple uses of the channel will (in general) be a big globally entangled (i.e. correlated) state, which makes the problem a lot harder to treat (Fang et al. 2020). In fact, in the quantum problem one can show that it is generally beneficial to pick input states one by one based on the previous outputs; this is called adaptivity, or an adaptive strategy (Harrow et al. 2010, Salek, Hayashi, and Winter 2022). These complexities have made it very hard to both gain an understanding of, and also obtain tangible bounds on, (optimal) quantities for quantum channel discrimination, and it is the main contribution of this thesis to make some progress in that regard.

To illustrate our results, let us introduce some aspects of the problem in slightly more detail. Whenever one has to distinguish between two possible hypotheses, there are two possible errors (or error probabilities; we will use these terms exchangeably) that can be incurred: We can infer the first hypothesis when the second is true (*type I error*) or vice versa (*type II error*). It is not hard to see that one can always change the way in which one settles on a hypothesis to reduce one of these errors by increasing the other (such as for example by always claiming the first hypothesis, this will never incur a type-II error), and hence in every case one has to find some kind of trade-off between these two errors. Generally, there are two main ways of doing this, either by treating the two errors symmetrically (i.e. minimizing the average), this is known as the symmetric setting, or by prioritizing one over the other and just requiring the other to be less than some fixed threshold (usually called ε), this is called the asymmetric setting. The symmetric setting might seem more natural, but the asymmetric setting has important applications when there is legitimate reason to consider mistakenly claiming one of the hypotheses to be worse than mistakenly claiming the other hypothesis.

In this thesis, for quantum channel discrimination, we will be exclusively looking at this asymmetric setting. The reasons for this are mostly technical: Since quantum entanglement exists, many asymptotic quantities in quantum information theory are hard to characterize (very roughly speaking, classical correlations are always mixtures of product distributions, whereas this is not true for quantum entanglement, and hence multipartite states on quantum systems can be very tricky to study). This turns out to be true also for the asymptotic error exponents of quantum channel discrimination, yet for the asymmetric error exponent (and pretty much only the asymmetric error exponent) the theory turns out in such a way that it is treatable with the methods currently available (for a slightly more in-depth discussion of why symmetric channel discrimination is hard, see Section 3.4.2 below).

In this asymmetric setting, the key object we consider is the asymmetric error probability (i.e. the minimum type-II error probability under a constraint on the type-I error), or more precisely its negative logarithm, which is called the hypothesis testing relative entropy (see Section 2.4 for a more thorough introduction of this quantity). For a quantum state-discrimination problem between the two quantum states ρ and σ we write it as:

$$D_H^\varepsilon(\rho\|\sigma) = -\log \inf_{0\leq M\leq \mathbb{1}} \{ \operatorname{Tr}(\sigma M) \mid \operatorname{Tr}(\rho M) \geq 1-\varepsilon \} . \tag{1.1.1}$$

We can now express the asymmetric error probability of various channel discrimination setups through this hypothesis testing relative entropy of state discrimination. We already mentioned earlier, that classically there is always a single optimal input state even when the channel is used multiple times, whereas this is not true quantumly, where for optimal strategies channel inputs have to be entangled, or adaptively picked.

We call a strategy with adaptively picked input states an *adaptive strategy*, and one where no adaptivity is present (but input states can still be entangled) a *parallel* strategy. For two quantum channels $\mathcal{E}$ and $\mathcal{F}$ the asymmetric error probability for a parallel strategy with n channel uses can be written as:

$$\sup_{\nu \in \mathcal{D}(A'^n A^n)} D_H^\varepsilon(\mathcal{E}^{\otimes n}(\nu) \| \mathcal{F}^{\otimes n}(\nu)) \,, \tag{1.1.2}$$

see Section 2.1 for a precise definition of all the notation used. This quantity is the main practically relevant object for characterizing channel discrimination with parallel strategies, and generally difficult to treat because of the optimization over a large space of input states. Remarkably though, we show in Theorem 5.5.2 that this quantity (i.e. the minimal asymmetric error probability for every $\varepsilon \in [0, 1]$, and also the corresponding optimal input state) can be phrased as a semi-definite program (SDP) in a way that makes it possible to calculate in time polynomial in n, which is remarkable since the naive complexity of this problem is exponential in n. Hence, this makes the computation of these optimal error probabilities for parallel strategies tractable for any given channels $\mathcal{E}$ and $\mathcal{F}$, even in the regime where n is not very small.

While we cannot show a similar computability result also for adaptive strategies, we can bound the difference there can be between a parallel and an adaptive strategy. For this, for any given adaptive strategy, we construct a parallel one and bound the difference in asymmetric error probability of this new parallel strategy and the given adaptive strategy (Corollary 3.3.1). This bound works in such a way that when the number of channel uses goes to infinity, the exponential decay rate of the asymmetric error probability of adaptive and parallel strategies becomes equal. This purely asymptotic statement was known before (Fang et al. 2020), although our results add significantly more understanding to it, in that we show explicitly what parallel strategies can be used to approximate given adaptive strategies, and also how the errors behave for finite number of channel uses. In particular, this result can be combined with the previous computability result for parallel strategies to obtain a bound on the best possible performance of adaptive strategies for finite number of channel uses, in a way that becomes asymptotically tight (on the level of decay rates) as the number of channel uses goes to infinity.

So far in this introduction, we have only been talking about channel discrimination in a setting where one should discriminate between two options for the channel (which we labeled $\mathcal{E}$ and $\mathcal{F}$). This is the simplest and most prototypical situation, although by far not the most practical one. In many cases, one does not know *exactly* which channel one could be given, quite often also just because of the presence of noise. This leads to the theory of composite hypothesis testing, where the hypotheses are given by sets of possible objects (like in this case two sets of channels, between which we should decide). This task of composite channel discrimination has surprisingly not been studied much, even in the classical literature, and so we are the first to provide some detailed results for the asymmetric error probability and corresponding decay rate when the number of channel uses goes to infinity (we call this the asymmetric error exponent). In the case of classical channels we are able to characterize this exponent depending on some properties of the involved sets (see Table 4.1 for an overview of our results). In particular, we can show that adaptive and parallel strategies are asymptotically equivalent if the sets are convex, but in general differ if the sets are non-convex. For the corresponding quantum problem, we give an expression for the asymptotic asymmetric error probability for

parallel strategies, together with a bound on adaptive strategies, but we cannot fully answer the question whether adaptive and parallel strategies are equivalent for convex sets of hypotheses.

Finally, we also study the problem of quantum channel discrimination in what are called infinite-dimensional systems (concretely, infinite-dimensional separable Hilbert spaces), which can be seen as the quantum analogue of continuous-variable classical probability distributions. While the prototypical quantum system is finite-dimensional, there are many important infinite-dimensional quantum systems (most notably in photonics and quantum optics), and so a theory of quantum channel discrimination is relevant there too (arguably the most relevant example of a quantum channel for quantum communication, a glass fibre, is an infinite-dimensional quantum system). Treating the problem in infinite-dimensions comes with some technical challenges, most notably that the general framework of quantum information theory is substantially less developed in this setting. Nevertheless, we are able to prove key results for simple (i.e. non-composite) channel discrimination also in this infinite-dimensional setting, in particular the relation between adaptive and parallel strategies (Corollary 5.5.1), and hence also a characterization of the asymptotic error exponent (known as the Stein's lemma for quantum channels, Theorem 5.6.1). With this, we can again establish the asymptotic equivalence of adaptive and parallel strategies (when considering the asymmetric error exponent), but unlike the finite-dimensional case, not quite for all quantum channels, but only those which satisfy an additional condition that a certain key quantity is finite. We subsequently give explicit (and non-classical) examples of channels, both for the cases where this condition is, and is not satisfied, and for which the corresponding quantum channel discrimination problem seems to be non-trivial. Some of these examples illustrate, that it would at least be possible that in infinite dimensions there exist channels for which the asymptotic asymmetric error exponent is not equal between adaptive and parallel strategies, although we also cannot prove that this is indeed the case for any of the channels we construct, and solving that question for these examples, which are chosen somewhat pathologically for our approach, seems to require quite different technical methods.

1.2 Structure of the Thesis and Main Results

In Chapter 2 we introduce all the required background, including previous results on quantum hypothesis testing, and all the different quantum divergences which we will make use of. Besides that, we also prove new technical lemmas about these divergences which are required for our later results.

In Chapter 3 we treat asymmetric quantum channel discrimination with simple hypotheses. After another more in-depth introduction of the problem, we prove our main computability result for the parallel n-shot error exponent (Theorem 3.2.1). Subsequently, we prove our main result regarding the conversion of a finite adaptive strategy into a parallel one (Theorem 3.3.4 and Corollary 3.3.1), and give an example illustrating this result in Section 3.3.3.

In Chapter 4 we extend our treatment to asymmetric quantum channel discrimination with composite hypotheses. We prove various results regarding classical and quantum asymptotic error exponents with composite hypotheses. for an overview of our results in this section see Table 4.1.

Finally, Chapter 5 deals with channel discrimination in infinite-dimensional (separable) Hilbert spaces,

for which we give a separate mathematical introduction of all the key quantities, aiming to highlight the subtleties that can arise in this setting. In general in this chapter, we give infinite-dimensional proofs of all the many small results required for our main statements, unless we could find explicit references. Our main results of the chapter are then the chain rule for the geometric Rényi divergence (Theorem 5.3.6), again a one-shot conversion result between adaptive and parallel strategies (Theorem 5.5.2 and Corollary 5.5.1), and the quantum Stein's lemma for channels in infinite dimensions (Theorem 5.6.1). Finally, as mentioned previously for the equivalence of adaptive and parallel strategies in infinite dimensions, we need an additional condition, which is discussed in detail in Section 5.6.3.

2 Mathematical Introduction and Preliminaries

This chapter introduces the key mathematical objects of this thesis in the finite-dimensional setting (which will be used in Chapters 3 and 4). The infinite-dimensional setting is introduced in Chapter 5. While this section is written with the aim of giving the reader an overview of relevant ideas and includes all required definitions, it can in no way give an exhaustive treatment of all the concepts that are introduced, for which we refer to the relevant literature. Good general treatments of quantum information theory can be found in e.g. Nielsen and Chuang (2010), Khatri and Wilde (2020), and Hayashi (2006).

2.1 Fundamental Objects of Classical and Quantum Information Theory

2.1.1 Probability Distributions and Density Matrices

The fundamental object of quantum information theory is that of a density matrix, a positive semi-definite operator of unit trace. Let us introduce a bit of notation: Throughout this thesis, we write $\mathcal{H}$ for a *complex Hilbert space* (finite-dimensional outside of Chapter 5), and $\mathcal{B}(\mathcal{H})$ for the set of *bounded linear operators* acting on $\mathcal{H}$. We use $\mathbb{1}$ for the *identity operator* on $\mathcal{B}(\mathcal{H})$, and sometimes use a subscript $\mathbb{1}_{\mathcal{H}}$ to emphasize the underlying Hilbert space. We write $\mathscr{P}(\mathcal{H})$ for the set of *positive semi-definite operators* acting on $\mathcal{H}$, and will also write $A \geq 0$ iff $A \in \mathscr{P}(\mathcal{H})$. For $A, B \in \mathscr{P}(\mathcal{H})$, we further write $A \ll B$ if $\operatorname{supp}(A) \subseteq \operatorname{supp}(B)$ and $A \not\ll B$ if $\operatorname{supp}(A) \not\subseteq \operatorname{supp}(B)$. $\mathcal{D}(\mathcal{H}) \subset \mathscr{P}(\mathcal{H})$ then denotes the set of *density matrices*, i.e., the set of positive semi-definite operators with trace one. When talking about density matrices, we will also use the term "quantum state" interchangeably.

In many cases we will study states of multipartite systems, i.e. systems which are composed out of more than one subsystem. We will label different quantum systems by capital Roman letters (A, B, C, etc.) and often use these letters interchangeably with the corresponding Hilbert space or set of density matrices (i.e., we write $\rho \in \mathcal{D}(A)$ instead of $\rho \in \mathcal{D}(\mathcal{H}_A)$). We will also concatenate these letters to mean tensor products of systems, i.e. we write $\rho \in \mathcal{D}(RA)$ for $\rho \in \mathcal{D}(\mathcal{H}_R \otimes \mathcal{H}_A)$, and include a superscript for tensor powers of a system, i.e. we write $\rho \in \mathcal{D}(A^n)$ for $\rho \in \mathcal{D}\left(\mathcal{H}_A^{\otimes n}\right)$.

In classical information theory, the analogous object to the density matrix (in finite dimensions) is a *probability distribution*, i.e. a function $P : \mathcal{X} \to [0, 1]$ on some finite discrete sample space $\mathcal{X}$ that satisfies $\sum_{x \in \mathcal{X}} P(x) = 1$. To simplify and harmonize notation, we will often use quantum notation throughout this thesis even when talking about a purely classical problem, where we will use the following embedding of probability distributions into density matrices: For a discrete sample space $\mathcal{X}$, let $\{|x\rangle\}_{x \in \mathcal{X}}$ be a fixed orthonormal basis of $\mathbb{C}^{|\mathcal{X}|}$. A probability distribution P is then written as the density matrix

$$\rho = \sum_{x \in \mathcal{X}} P(x) \, |x\rangle\langle x| \ . \tag{2.1.1}$$

With slight abuse of notation, we also write the set of all density matrices that are diagonal in this basis as $\mathcal{D}(\mathcal{X})$, i.e. the set $\mathcal{D}(\mathcal{X})$ contains all probability distributions on $\mathcal{X}$ written as diagonal density matrices. Throughout this thesis, we will use the calligraphic letters $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ to denote such classical systems, i.e. a density matrix on these three systems is automatically assumed to be diagonal in some fixed basis, and thus corresponds to a classical probability distribution. These three systems can also be paired with quantum systems to form classical-quantum states, i.e. a density matrix $\rho_{\mathcal{X}A} \in \mathcal{D}(\mathcal{X}A)$ is a state of the form:

$$\rho_{\mathcal{X}A} = \sum_{x \in \mathcal{X}} p(x) \, |x\rangle\langle x| \otimes \rho_A^{(x)} \tag{2.1.2}$$

where $p : \mathcal{X} \to [0, 1]$ is a probability distribution, and $\rho_A^{(x)} \in \mathcal{D}(A)$ for all $x \in \mathcal{X}$.

2.1.2 Quantum Channels

A *quantum channel* (in this thesis often denoted as $\mathcal{E}$ or $\mathcal{F}$) is a completely positive trace preserving linear map between density operators, for example $\mathcal{E} : \mathcal{D}(\mathcal{H}_A) \to \mathcal{D}(\mathcal{H}_B)$. To simplify notation, and analogously to the previously introduced notation for density matrices, we will often write $\mathcal{E} : A \to B$ instead of $\mathcal{E} : \mathcal{D}(\mathcal{H}_A) \to \mathcal{D}(\mathcal{H}_B)$. We further write $\mathrm{CPTP}(A \to B)$ for the set of all quantum channels from $\mathcal{D}(\mathcal{H}_A)$ to $\mathcal{D}(\mathcal{H}_B)$. As previously stated, we denote classical systems using the calligraphic letters $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$, and any density matrix on these systems is assumed to be diagonal in a fixed basis. We assume also any channel acting on (or mapping into) one of such classical systems to respect this structure, and in particular then any channel in $\mathrm{CPTP}(\mathcal{X} \to \mathcal{Y})$ corresponds to a classical channel (i.e. it is equivalent to a stochastic matrix operating between probability distributions, via the mapping between density matrices and probability distributions introduced earlier).

Implicit Identities

Throughout this thesis, we will often be in a setting where a channel acts only on a part of an input state (with the untouched part often being called a "reference system", and often bearing the letter R). In this case, writing the output state as $(\mathrm{id}_R \otimes \mathcal{E})(\nu_{RA})$ (for a channel $\mathcal{E} : A \to B$) is often quite verbose. To simplify notation we will often make the identities implicit and thus write

$$(\mathrm{id}_R \otimes \mathcal{E})(\nu_{RA}) \equiv \mathcal{E}(\nu_{RA}) \equiv \mathcal{E}(\nu) \, . \tag{2.1.3}$$

To avoid any ambiguity with this notation, we will always do the following:

- We will always give labels to the input and output systems of any channel we use, i.e. we will always specify $\mathcal{E} : A \to B$.
- We will always specify the systems on which a state is defined, i.e. we will always specify $\nu_{RA} \in \mathcal{D}(RA)$, although we might replace ν_{RA} with ν in subsequent equations to further simplify notation.

There will then be an implicit identity in $\mathcal{E}(\nu)$ if (and only if) the state ν has any subsystem on which the channel $\mathcal{E}$ does not act.

Choi Matrices

Every quantum channel $\mathcal{E} : A \to B$ is uniquely described by its so-called *Choi matrix*, which is defined as follows. Let $R \cong A$ be another system identical to A, and let $\{|i\rangle_R\}_{i=1}^d$, $\{|i\rangle_A\}_{i=1}^d$ be orthonormal bases of R and A. Then, with the maximally entangled state on RA:

$$|\Phi_{RA}\rangle = \frac{1}{\sqrt{d}} \sum_{i=1}^{d} |i\rangle_R |i\rangle_A \tag{2.1.4}$$

$$\Phi_{RA} = |\Phi_{RA}\rangle\langle\Phi_{RA}| \tag{2.1.5}$$

the Choi matrix of $\mathcal{E}$ is defined as

$$\Gamma_{RB}^{\mathcal{E}} \coloneqq d\,\mathcal{E}(\Phi_{RA}) = d\,(\mathrm{id}_R \otimes \mathcal{E})(\Phi_{RA})\,, \tag{2.1.6}$$

where the factor d is included for convenience. In particular, the action of the channel on any input state $\rho \in \mathcal{D}(A)$ can be reconstructed from the Choi matrix as

$$\mathcal{E}(\rho) = \mathrm{Tr}_R(\Gamma_{RB}^{\mathcal{E}}(\rho_R^T \otimes \mathbb{1}_B))\,, \tag{2.1.7}$$

where we write $\rho = \rho_R$ to emphasize how $\rho^T \otimes \mathbb{1}_B$ becomes an operator on the joint system RB, and the transposition of ρ is performed in the same basis that we used for the definition of the maximally entangled state Φ_{RA} (since transposition is not a basis-invariant operation).

2.1.3 Measurements and POVMs

A general measurement in quantum mechanics can be described by a *positive operator valued measure* (POVM), where we will restrict ourselves to POVMs with finitely many outcomes (corresponding to discrete measures with finitely many atoms). Such a discrete POVM is a finite collection of positive operators $(M_i \in \mathscr{P}(A))_{i=1}^n$, such that $\sum_i M_i = \mathbb{1}$. This specifies a measurement of a quantum state ρ, such that outcome i is obtained with probability $\mathrm{Tr}(\rho M_i)$. To any POVM specified by $(M_i \in \mathscr{P}(A))_{i=1}^n$, we can associate a quantum-classical channel $\mathcal{M} : A \to \mathcal{X}$ (where $|\mathcal{X}| = n$) via

$$\mathcal{M}(\rho_A) = \sum_{i=1}^{n} \mathrm{Tr}(\rho_A M_i)\,|i\rangle\langle i|\,, \tag{2.1.8}$$

where $\{|i\rangle\}_{i=1}^n$ is the fixed orthonormal basis in which classical states are encoded in $\mathcal{X}$. Throughout this thesis, we will usually specify and refer to any measurements and POVMs in terms of their quantum-classical channels.

2.1.4 Distances Between States

Schatten p-norms

For any operator $A \in \mathcal{B}(\mathcal{H})$, and any $p \in (0, \infty)$, the *Schatten p-norm* of A is defined as

$$\|A\|_p = \mathrm{Tr}(|A|^p)^{1/p} = \left(\mathrm{Tr}\left[(A^\dagger A)^{p/2}\right]\right)^{1/p}. \tag{2.1.9}$$

Special cases of this include the one-norm (also called trace norm) for $p = 1$, or the infinity norm for $p \to \infty$, which can also be written as $\|A\|_\infty = \lambda_{\max}(A)$, i.e. the largest singular value of A (or largest eigenvalue if A is diagonalizable).

Trace Distance

The most commonly used distance in quantum information theory is the *trace distance*, which is the quantum analogue of the total variation distance in classical information theory. For two density matrices $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, their trace distance is defined as

$$T(\rho, \sigma) \coloneqq \frac{1}{2}\|\rho - \sigma\|_1 . \tag{2.1.10}$$

The trace distance satisfies the properties of a metric, and also the data-processing inequality, i.e. it is contractive under the action of any quantum channel $\mathcal{E} : \mathcal{D}(\mathcal{H}_A) \to \mathcal{D}(\mathcal{H}_B)$, so

$$T(\rho_A, \sigma_A) \geq T(\mathcal{E}(\rho_A), \mathcal{E}(\rho_B)) \tag{2.1.11}$$

(see, e.g., Khatri and Wilde (2020) for a proof).

Fidelity and Purified Distance

Another often used distance measure (which also has a classical analogue, but is actually a lot more natural in quantum information theory than in classical information theory) is the *fidelity*, which can be seen as a mixed state generalization of the inner product between pure states. For two quantum states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, we define the fidelity as (Uhlmann 1976)

$$F(\rho, \sigma) \coloneqq \mathrm{Tr}\left(\sqrt{\sqrt{\sigma}\rho\sqrt{\sigma}}\right). \tag{2.1.12}$$

Note that this is sometimes also called the square root fidelity due to different conventions on whether to include a square or not. The fidelity itself is not a metric, but one can derive one from it via (Rastegin 2002, 2003, Gilchrist, Langford, and Nielsen 2005, Rastegin 2006)

$$P(\rho, \sigma) \coloneqq \sqrt{1 - F(\rho, \sigma)^2}, \tag{2.1.13}$$

which satisfies the properties of a metric (especially the triangle inequality) and also the data-processing inequality (again, see e.g., Khatri and Wilde (2020)). This quantity was originally introduced as the sine

distance, although is mostly known as the *purified distance*. The trace distance and purified distance satisfy the following relations, which are known as Fuchs-van de Graaf inequalities (Fuchs and van de Graaf 1999)

$$1 - \sqrt{1 - P(\rho, \sigma)^2} \leq T(\rho, \sigma) \leq P(\rho, \sigma)\,. \tag{2.1.14}$$

2.2 Classical and Quantum Entropies and Divergences

2.2.1 Quantum Divergences

Quantum divergences will be the key mathematical objects used in this thesis to prove our results regarding quantum channel discrimination. Generally, we understand a quantum divergence $\mathbf{D}$ to be a function of two (usually not necessarily normalized) positive operators $\mathbf{D} : \mathscr{P}(\mathcal{H}) \times \mathscr{P}(\mathcal{H}) \to \mathbb{R} \cup \{\infty\}$ that also satisfies the following two properties:

- Positivity: Whenever ρ and σ are normalized (i.e. $\rho, \sigma \in \mathcal{D}(\mathcal{H})$), $\mathbf{D}(\rho\|\sigma) \geq 0$.
- Data-Processing Inequality: Whenever ρ is normalized and $\mathcal{E} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{H}')$ is a quantum channel, $\mathbf{D}(\mathcal{E}(\rho)\|\mathcal{E}(\sigma)) \leq \mathbf{D}(\rho\|\sigma)$.

As seen already in the statement of the data-processing inequality, we generally want a quantum divergence to be defined not only on one particular Hilbert space $\mathcal{H}$, but on all (in this section finite-dimensional) Hilbert spaces $\mathcal{H}$. Generally speaking, divergences often quantify the degree to which the two inputs ρ and σ are "different", although not directly in a metric sense (divergences usually don't satisfy the triangle inequality), but in a more "entropic" sense, which is usually related to certain applications (the most famous one being the (quantum) Stein's Lemma for the relative entropy, which is discussed in Section 2.4).

The prototypical, and by far the most important, quantum divergence is the quantum relative entropy, which can be defined as:

$$D(\rho\|\sigma) = \mathrm{Tr}[\rho(\log(\rho) - \log(\sigma))]\,. \tag{2.2.1}$$

Throughout this thesis all logarithms are understood to be base 2.

Formal Definitions and Support Conditions

The equation (2.2.1) is not obviously well-defined whenever ρ and σ are not full-rank. This turns out to be a generic issue with most (not only quantum) divergences. There typically is no actual issue with the first argument having a kernel (in the case of (2.2.1), the fact that the real function $x \mapsto x\log(x)$ can be continuously extended to $x = 0$ implies that also the matrix function $\rho \mapsto \mathrm{Tr}(\rho\log(\rho))$ can be continuously extended to positive semi-definite matrices), although there can be singularities when the second argument σ is not full-rank.

In general, all quantum and classical divergences we consider in this thesis will be given by formulas which are without issues whenever the second argument is a full-rank (i.e. positive definite, and hence invertible) state. Furthermore, all divergences considered here are continuous when mixed with the identity in the

second argument, i.e. in all cases they can be defined for only positive semi-definite ρ, σ as

$$\mathbf{D}(\rho\|\sigma) \coloneqq \lim_{\varepsilon\to 0} \mathbf{D}(\rho\|\sigma + \varepsilon\mathbb{1}) \,, \tag{2.2.2}$$

where $\mathbf{D}$ is any of the divergences considered in this thesis, and this limit is allowed to be ∞. In many cases we will also specify what these limits exactly are equal to. In particular, all divergences considered here are such that whenever $\rho \ll \sigma$, i.e. the support of ρ is contained in the support of σ, then one can ignore the kernel of σ completely, and apply all formulas as if everything was defined on the reduced Hilbert space of the support of σ.

Note that while we use (2.2.2) as a definition for all our divergences, these divergences are usually not continuous (that means jointly continuous). For most of them it is not hard to find two (not full-rank) states ρ, σ, and a sequence of full rank density matrices $(\sigma_n)_n$ such that

$$\lim_{n\to\infty} \mathbf{D}\left(\rho + \frac{1}{n}\mathbb{1} \,\middle\|\, \sigma + \frac{1}{n}\sigma_n\right) \neq \mathbf{D}(\rho\|\sigma) \,, \tag{2.2.3}$$

and similarly also if one was to normalize the two arguments on the left-hand side.

Finally, let us remark that we will define all our divergences on positive semi-definite operators, i.e. both arguments need not be normalized. Note also, that we do not divide by the trace of the first argument in any of our definitions (this is sometimes done in the literature, in particular in certain axiomatic characterizations).

2.2.2 The Quantum Relative Entropy

For $\rho \in \mathcal{D}(\mathcal{H})$ and $\sigma \in \mathscr{P}(\mathcal{H})$ the *(Umegaki) quantum relative entropy* can be defined as (Umegaki 1962)

$$D(\rho\|\sigma) \coloneqq \mathrm{Tr}[\rho(\log\rho - \log\sigma)], \tag{2.2.4}$$

if $\rho \ll \sigma$ and $D(\rho\|\sigma) \coloneqq \infty$ if $\rho \not\ll \sigma$, and it famously satisfies the data-processing inequality (Lindblad 1975). One particular consequence of that is that it is also jointly convex, i.e. for any two finite sets of states $\{\rho_i\}_{i=1}^n$, $\{\sigma_i\}_{i=1}^n \subset \mathcal{D}(\mathcal{H})$ and a probability distribution $\{p_i\}_{i=1}^n$ it holds that

$$D\left(\sum_i p_i\rho_i \,\middle\|\, \sum_i p_i\sigma_i\right) \leq \sum_i p_i D(\rho_i\|\sigma_i) \,. \tag{2.2.5}$$

The Measured Relative Entropy

We can also define the *measured relative entropy* as the maximal classical relative entropy when measuring both states with a POVM. Specifically,

$$D_M(\rho\|\sigma) \coloneqq \sup_{\mathcal{M}\,\mathrm{POVM}} D(\mathcal{M}(\rho)\|\mathcal{M}(\sigma)) \,, \tag{2.2.6}$$

where $\mathcal{M}$ is a POVM (with arbitrarily many outcomes) interpreted as a quantum-classical channel as outlined above. Crucially, the quantum relative entropy and the measured relative entropy are different in general

(this is very well-known, but see e.g. Berta, O. Fawzi, and Tomamichel (2017) for an exhaustive treatment).

The von Neumann and Shannon Entropy

The von Neumann and Shannon entropies can be seen as child quantities of the above-defined quantum relative entropy.

For any density matrix $\rho \in \mathcal{D}(\mathcal{H})$ we can write its *von Neumann entropy* as (Von Neumann 1932)

$$H(\rho) = -D(\rho \| \mathbb{1}_{\mathcal{H}}) = -\operatorname{Tr}(\rho \log(\rho)) \,. \tag{2.2.7}$$

Equivalently, for a classical probability distribution $P : \mathcal{X} \to [0,1]$, we can define its *Shannon entropy* via the above-mentioned embedding of classical probability distributions into quantum states, which leads to (Shannon 1948)

$$H(P) = -\sum_{x \in \mathcal{X}} P(x) \log(P(x)) \,. \tag{2.2.8}$$

Binary Entropy Function

In multiple places below, we will make use of the *binary entropy* function, which is the Shannon entropy of a classical random variable that takes two possible values. It is uniquely specified by the probability $p \in [0,1]$ of one of the values and is given by

$$h(p) \coloneqq -p \log p - (1-p) \log(1-p) \,. \tag{2.2.9}$$

It satisfies

$$0 \leq h(p) \leq 1 \qquad \forall\, p \in [0,1] \,. \tag{2.2.10}$$

Properties of the Quantum Relative Entropy and Measured Relative Entropy

The following two Lemmas establish some continuity properties of the relative entropy which we will use later on. Both statements are well known, but included here for the sake of completeness.

Lemma 2.2.1. *The quantum relative entropy $D(\rho\|\sigma)$ is a lower semi-continuous function of ρ, σ.*

Proof. It is easy to see that

$$D(\rho\|\sigma) = \lim_{\varepsilon \downarrow 0} D(\rho\|\sigma + \varepsilon \mathbb{1}) \,, \tag{2.2.11}$$

and since the function $x \mapsto x \log(x)$ is continuous on $[0,\infty)$, for all $\varepsilon > 0$ the function $D(\rho\|\sigma + \varepsilon\mathbb{1})$ is well-defined and jointly continuous in ρ and σ. Since the logarithm is operator monotone, we have $D(\rho\|\sigma_1) \leq D(\rho\|\sigma_2)$ whenever $\sigma_1 \geq \sigma_2$, and hence

$$\lim_{\varepsilon \downarrow 0} D(\rho\|\sigma + \varepsilon\mathbb{1}) = \sup_{\varepsilon > 0} D(\rho\|\sigma + \varepsilon\mathbb{1}) \tag{2.2.12}$$

and the supremum of lower semi-continuous functions is lower semi-continuous. □

Note that the only property of D that we used is that $D(\rho\|\sigma)$ is fully continuous on sets where the minimum eigenvalue of σ is bounded from below by a positive value, and this property is similarly satisfied for many other divergences in this thesis, which are therefore also all lower semi-continuous.

The following Lemma establishes the continuity of the relative entropy when evaluated on the (potentially measured) outputs of two different channels with the same input state. It uses some concepts introduced later in this chapter (such as the max-divergence from Section 2.5, or the Petz–Rényi divergence from Section 2.3.1), but is still included here since it is in essence a continuity statement for the quantum relative entropy.

Lemma 2.2.2. *Let $\mathcal{E}, \mathcal{F} \in \mathrm{CPTP}(A \to B)$ be such that $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$, and let R be an arbitrary auxiliary system. Then the function*

$$(\mathcal{M}, \nu) \mapsto D(\mathcal{M} \circ \mathcal{E}(\nu)\|\mathcal{M} \circ \mathcal{F}(\nu)) \tag{2.2.13}$$

is continuous on $\mathrm{CPTP}(B \to C) \times \mathcal{D}(R \otimes A)$.

Proof. It is easy to see that for $0 < \alpha < 1$ the Petz-Rényi divergence (defined below in Section 2.3.1)

$$D_\alpha(\rho\|\sigma) \coloneqq \frac{1}{\alpha - 1} \log \mathrm{Tr}\left(\rho^\alpha \sigma^{1-\alpha}\right) \tag{2.2.14}$$

is jointly continuous in ρ and σ. Using the continuity bound from Lemma 2.3.1, together with $D_\alpha(\rho\|\sigma) \geq 0$ and $D_\alpha(\rho\|\sigma) \leq D_{\max}(\rho\|\sigma)$ (see (2.5.3) for the latter) we find

$$|D_\alpha(\rho\|\sigma) - D(\rho\|\sigma)| \leq (1 - \alpha) \log^2 \left(2^{D_{\max}(\rho\|\sigma)} + 2\right). \tag{2.2.15}$$

As $D_{\max}(\mathcal{M} \circ \mathcal{E}(\nu)\|\mathcal{M} \circ \mathcal{F}(\nu)) \leq D_{\max}(\mathcal{E}\|\mathcal{F})$ which is independent of $\mathcal{M}$ and ν, we get that $D_\alpha(\mathcal{M} \circ \mathcal{E}(\nu)\|\mathcal{M} \circ \mathcal{F}(\nu))$ converges to $D(\mathcal{M} \circ \mathcal{E}(\nu)\|\mathcal{M} \circ \mathcal{F}(\nu))$ uniformly in ν and $\mathcal{M}$ as $\alpha \uparrow 1$. Hence, also the limiting function $(\mathcal{M}, \nu) \mapsto D(\mathcal{M} \circ \mathcal{E}(\nu)\|\mathcal{M} \circ \mathcal{F}(\nu))$ is continuous. □

See also Mosonyi and Hiai (2024) for a more exhaustive treatment of similar continuity properties of many quantum divergences.

We say that a divergence $\mathbf{D}$ satisfies the direct-sum property, if

$$\mathbf{D}\left(\bigoplus_{i=1}^n p_i \rho_i \,\middle\|\, \bigoplus_{i=1}^n p_i \sigma_i\right) = \sum_{i=1}^n p_i \mathbf{D}(\rho_i\|\sigma_i). \tag{2.2.16}$$

whenever $\{\rho_i\}_i, \{\sigma_i\}_i \subset \mathcal{H}_i$ are two sets of density matrices and $\{p_i\}_{i=1}^n$ is a probability distribution.

It is easy to see from its definition, that the quantum relative entropy satisfies the direct-sum property. The same is also true for the measured relative entropy:

Lemma 2.2.3. *The measured relative entropy D_M satisfies the data-processing inequality, is lower semi-continuous, and satisfies the direct-sum property.*

Proof. For the data-processing inequality it is easy to see that concatenating a POVM $\mathcal{M}$ (treated as a classical-quantum channel as explained in Section 2.1.2) and an arbitrary quantum channel $\mathcal{E}$ as $\mathcal{M} \circ \mathcal{E}$ yields another POVM, and hence D_M satisfies the data-processing inequality, as the supremum over all POVMs of the form

$\mathcal{M} \circ \mathcal{E}$ can only be smaller than the supremum over all POVMs. D_M is also lower semi-continuous as a supremum of lower semi-continuous functions is lower semi-continuous, and D is lower semi-continuous. For the direct-sum property, let $\mathcal{H} = \bigoplus_{i=1}^n \mathcal{H}_i$, and $\{\rho_i\}_i, \{\sigma_i\}_i \subset \mathcal{D}(\mathcal{H}_i)$ be two sets of density matrices, $\{\mathcal{M}_i\}_i$ be a set of POVMs each on $\mathcal{H}_i$ and $\{p_i\}_i$ be a probability distribution, where all indices are in the range $i = 1, \ldots n$. Define $\rho = \bigoplus_{i=1}^n p_i\rho_i$, $\sigma = \bigoplus_{i=1}^n p_i\sigma_i$ and $\bar{\mathcal{M}} \coloneqq \bigoplus_{i=1}^n \mathcal{M}_i$. It is easy to see that

$$\bar{\mathcal{M}}\left(\bigoplus_{i=1}^n p_i\rho_i\right) = \bigoplus_{i=1}^n p_i\mathcal{M}_i(\rho_i) \tag{2.2.17}$$

and hence

$$\begin{aligned} D_M\left(\bigoplus_{i=1}^n p_i\rho_i \,\middle\|\, \bigoplus_{i=1}^n p_i\sigma_i\right) &\geq D\left(\bar{\mathcal{M}}\left(\bigoplus_{i=1}^n p_i\rho_i\right) \,\middle\|\, \bar{\mathcal{M}}\left(\bigoplus_{i=1}^n p_i\sigma_i\right)\right) && (2.2.18)\\ &= D\left(\bigoplus_{i=1}^n p_i\mathcal{M}_i(\rho_i) \,\middle\|\, \bigoplus_{i=1}^n p_i\mathcal{M}_i(\sigma_i)\right) && (2.2.19)\\ &= \sum_{i=1}^n p_i D(\mathcal{M}_i(\rho_i)\|\mathcal{M}_i(\sigma_i)), && (2.2.20) \end{aligned}$$

which leads to the $\geq$ direction in the direct-sum property after optimizing over the $\mathcal{M}_i$. For the reverse direction, let $\mathcal{M}$ be an arbitrary POVM on $\mathcal{H}$. Then

$$D(\mathcal{M}(\rho)\|\mathcal{M}(\sigma)) = D\left(\sum_i p_i\mathcal{M}(\rho_i) \,\middle\|\, \sum_i p_i\mathcal{M}(\sigma_i)\right) \leq \sum_i p_i D(\mathcal{M}(\rho_i)\|\mathcal{M}(\sigma_i)) \leq \sum_i p_i D_M(\rho_i\|\sigma_i), \tag{2.2.21}$$

where we used the joint convexity of the relative entropy. The claim follows by optimizing over all $\mathcal{M}$. □

2.2.3 Divergences Between Channels

For every given divergence $\mathbf{D}$ between quantum states, one can define an associated *channel divergence* (Leditzky et al. 2018) by performing a (so-called stabilized) maximization over all input states. Concretely, with $\mathcal{E}, \mathcal{F} : A \to B$ being quantum channels, we define

$$\mathbf{D}(\mathcal{E}\|\mathcal{F}) \coloneqq \sup_{\substack{\rho_{RA} \in \mathcal{D}(R\otimes A)\\ R \text{ arbitrary}}} \mathbf{D}\big((\mathrm{id}_R \otimes \mathcal{E})(\rho)\|(\mathrm{id}_R \otimes \mathcal{F})(\rho)\big). \tag{2.2.22}$$

Since $\mathbf{D}$ satisfies the data-processing inequality by definition, the supremum can in fact be restricted to pure states such that the reference system R is isomorphic to the channel input system A. This is a standard argument, but also shown here again as a slightly more general statement we will make use of below:

Lemma 2.2.4. *Let $\mathbf{D}$ be a quantum divergence satisfying the data-processing inequality, and let $\mathcal{S}, \mathcal{T} \subset$*

$\mathrm{CPTP}(A \to B)$ *be two arbitrary sets of channels. Then,*

$$\sup_{\substack{\nu \in \mathcal{D}(R\otimes A)\\ R \text{ arbitrary}}} \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \sup_{\substack{\nu \in \mathcal{D}(R\otimes A)\\ R\cong A}} \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)), \tag{2.2.23}$$

i.e. the size of the system R can be restricted to be isomorphic to A. Furthermore, the supremum can be restricted to pure states.

Proof. Let R be an arbitrary reference system, and consider a state $\nu_{RA} \in \mathcal{D}(R \otimes A)$. Let ν have a purification ν_{PRA}. Furthermore, take the state $\nu_A = \mathrm{Tr}_R \nu_{AR}$ and a purification ν_{SA}, where $S \cong A$ (it is well known that such a purification always exists). Then, as also ν_{PRA} is a purification of ν_A, ν_{PRA} and ν_{SA} are related by a partial isometry $V : PR \to S$. As the channels $\mathcal{E}$ and $\mathcal{F}$ all act as identity on these systems, the partial isometry commutes with them, and so we get from the data-processing inequality that

$$D(\mathcal{E}(\nu_{RA})\|\mathcal{F}(\nu_{RA})) \leq D(\mathcal{E}(\nu_{PRA})\|\mathcal{F}(\nu_{PRA})) = D(\mathcal{E}(\nu_{SA})\|\mathcal{F}(\nu_{SA}))) \tag{2.2.24}$$

and hence the supremum can be restricted to pure states and reference systems isomorphic to A. Note that the set $\mathcal{D}(S \otimes A) = \mathcal{D}(A \otimes A)$ is compact. □

Furthermore, for most divergences $\mathbf{D}$, the supremum in the corresponding channel divergence is also achieved. The following Lemma establishes this for the cases we need in this thesis:

Lemma 2.2.5. *Let $\mathcal{E}, \mathcal{F} \in \mathrm{CPTP}(A \to B)$, and R be an arbitrary auxiliary system. Then, if $\mathbf{D}$ is either D or D_M:*

$$\sup_{\nu \in \mathcal{D}(R\otimes A)} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \max_{\substack{\nu \in \mathcal{D}(R\otimes A)\\ R\cong A}} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) \tag{2.2.25}$$

i.e. the supremum is achieved and R can be chosen isomorphic to A.

Proof. The restriction to R isomorphic to A follows from Lemma 2.2.4. What remains to show is that the supremum is achieved. Consider first the case where $D_{\max}(\mathcal{E}\|\mathcal{F}) = \infty$ ($D_{\max}$ is introduced in Section 2.5). As shown in Wilde et al. (2020), $D_{\max}(\mathcal{E}\|\mathcal{F}) = D_{\max}(\mathcal{E}(\Omega)\|\mathcal{F}(\Omega))$ where $\Omega = \Omega_{RA}$ is a maximally entangled state. It is well known that (in finite dimensions) $D_{\max}(\rho\|\sigma) = \infty \Rightarrow D(\rho\|\sigma) = \infty$ for all states ρ, σ, and hence if $\mathbf{D} = D$ the value of infinity is achieved in this case. It is also not hard to see that if $D(\rho\|\sigma) = \infty$ also $D_M(\rho\|\sigma) = \infty$ (in this case σ will have a smaller support than ρ, so a POVM built from the projection onto the support of σ will achieve infinity). Hence in this case also with $\mathbf{D} = D_M$ the supremum is achieved. So assume $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$. If $\mathbf{D} = D$, then by Lemma 2.2.2, the function is continuous in ν and hence, since the set of density matrices optimized over is compact, the optimum value is achieved. If $\mathbf{D} = D_M$ we have the expression

$$\sup_{\nu \in \mathcal{D}(R\otimes A)} D_M(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \sup_{\nu \in \mathcal{D}(R\otimes A)} \sup_{\mathcal{M}\ \mathrm{POVM}} D(\mathcal{M}\circ\mathcal{E}(\nu)\|\mathcal{M}\circ\mathcal{F}(\nu)) \tag{2.2.26}$$

Again, by Lemma 2.2.2 the expression optimized over is continuous in ν and $\mathcal{M}$. Also, by Berta, O. Fawzi, and Tomamichel (2017) the optimization over $\mathcal{M}$ in the measured relative entropy can be restricted to von

Neumann measurements (i.e. projective rank-1 measurements), and the set of these measurements is compact (as such a measurement is uniquely specified by a unitary matrix). Hence, also here the optimum value is achieved. □

Regularized and Amortized Divergences

We can also define the *regularized* and *amortized* (Wilde et al. 2020) channel divergences as

$$\mathbf{D}^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) \coloneqq \lim_{n\to\infty} \frac{1}{n}\mathbf{D}(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}), \tag{2.2.27}$$

$$\mathbf{D}_A(\mathcal{E}\|\mathcal{F}) \coloneqq \sup_{\substack{\rho,\sigma\in\mathcal{D}(RA)\\ R \text{ arbitrary}}} [\mathbf{D}(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) - \mathbf{D}(\rho\|\sigma)]. \tag{2.2.28}$$

Note that the limit in (2.2.27) only exists if the divergence $\mathbf{D}$ is superadditive on states (i.e. $\mathbf{D}(\rho_1\otimes\rho_2\|\sigma_1\otimes\sigma_2) \geq \sum_{i=1}^2 D(\rho_i\|\sigma_i)$), which is the case every time when we consider $\mathbf{D}^{\mathrm{reg}}$ in this thesis, otherwise one would have to replace the $\lim$ with a $\limsup$ or $\liminf$. Note also, that for the amortized divergence, there is no known way in which the size of the reference system can be restricted.

2.3 Rényi Entropies and Divergences

Rényi entropies and divergences can be seen as a one-parameter generalization of the above mentioned (relative) entropy. In particular, Rényi (1961) showed that all classical functions of two probability distributions that satisfy certain properties (properties satisfied by the relative entropy introduced above) are given by Rényi divergences (see also Tomamichel (2016) for a "quantization" of this axiomatic approach). The classical Rényi divergences form a one-parameter family, parametrized by $\alpha \in [0,1)\cup(1,\infty)$, and are defined for two probability distributions P, Q on a finite sample space as

$$D_\alpha(P\|Q) \coloneqq \frac{1}{\alpha-1}\log\left[\sum_x P(x)^\alpha Q(x)^{\alpha-1}\right], \tag{2.3.1}$$

with the value being possibly ∞ if $\alpha > 1$ and if there exists an x, such that $Q(x) = 0$, but $P(x) > 0$.

$\alpha \to 1$ Limit

The expression (2.3.1) is not well-defined for $\alpha = 1$, although one can calculate the limit $\alpha \to 1$ of it, and this happens to be equal to the previously defined relative entropy, i.e.

$$\lim_{\alpha\to 1} D_\alpha(P\|Q) = \lim_{\alpha\uparrow 1} D_\alpha(P\|Q) = \lim_{\alpha\downarrow 1} D_\alpha(P\|Q) = D(P\|Q)\,. \tag{2.3.2}$$

Hence, the expression (2.3.1) can be continuously extended to $\alpha \in [0,\infty)$. It turns out that this will also be the case for all (classical and quantum) Rényi divergences we introduce in this section and consider in this thesis. To simplify notation (i.e. to be able to write "for all $\alpha > 0$" instead of "for all $\alpha \in (0,1)\cup(1,\infty)$ and

additionally in the limit"), from here on we will consider it understood that for any Rényi divergence $\mathbf{D}_\alpha$ introduced in this section

$$\mathbf{D}_1(\rho\|\sigma) = \lim_{\alpha\to 1} \mathbf{D}_\alpha(\rho\|\sigma)\,, \tag{2.3.3}$$

and while this limit might not need to exist in general, it will for all Rényi divergences considered in this thesis.

Common Properties of Quantum Rényi Divergences

We will subsequently introduce multiple quantum generalizations of this classical Rényi divergence. These are then again quantum divergences that are parametrized by a parameter α which can take values in the positive real line (or maybe only a subset of it). We usually write "Rényi " divergence for the family parametrised by α, and α-Rényi divergence for the member of that family for a specific α. What we require for a divergence to be called a "quantum α-Rényi divergence" is that

1. It is equal to the classical α-Rényi divergence when evaluated on classical states (via the embedding of classical states into density matrices discussed in Section 2.1.1)
2. It satisfies the data-processing inequality (as introduced in Section 2.2.1)

We then call a family of quantum α-Rényi divergences a quantum Rényi divergence.

2.3.1 The Petz–Rényi Divergence

The *Petz–Rényi divergence* is the most immediate quantum generalization of the classical Rényi divergence (2.3.1). For every $\rho \in \mathscr{P}(\mathcal{H})$, full rank $\sigma \in \mathscr{P}(\mathcal{H})$, and $\alpha \in [0,1) \cup (1,\infty)$, it is defined as (Petz 1986):

$$D_\alpha(\rho\|\sigma) \coloneqq \frac{1}{\alpha-1} \log \operatorname{Tr}\left(\rho^\alpha \sigma^{1-\alpha}\right). \tag{2.3.4}$$

For general (non full-rank) σ, it can be defined by the limiting expression specified above. It is monotonically increasing in α whenever $\rho \in \mathcal{D}(\mathcal{H})$ is normalized, and satisfies the data-processing inequality for $\alpha \in [0,1) \cup (1,2]$, but not for any $\alpha > 2$ (see e.g., Khatri and Wilde (2020) and Hiai and Mosonyi (2017)).

Additionally, whenever $\rho \in \mathcal{D}(\mathcal{H})$ is normalized, its limit $\alpha \to 1$ is given by the Umegaki relative entropy (2.2.4), and in fact, one can give explicit bounds on how fast this convergence happens (Tomamichel 2012). For some of our later results, we will require such a bound for both α converging from above and below, which we prove here (this Lemma is a slight refinement of Tomamichel (2012), Lemma 6.3, see Remark 2.3.2 for more details on how they compare):

Lemma 2.3.1 (Petz–Rényi Continuity Bound in α)**.** *Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be normalized quantum states. For $\gamma \in (0,1]$, define*

$$c_\gamma(\rho\|\sigma) \coloneqq \frac{1}{\gamma} \log\left(2^{\gamma D_{1+\gamma}(\rho\|\sigma)} + 2^{-\gamma D_{1-\gamma}(\rho\|\sigma)} + 1\right). \tag{2.3.5}$$

Then, for all $\gamma \in (0,1]$ *and* $\delta \in (0, \frac{\gamma}{2}]$*:*

$$D_{1+\delta}(\rho\|\sigma) \leq D(\rho\|\sigma) + \ln(2)\delta(c_\gamma(\rho\|\sigma))^2 \tag{2.3.6}$$
$$\leq D(\rho\|\sigma) + \delta(c_\gamma(\rho\|\sigma))^2. \tag{2.3.7}$$

Furthermore, if $D(\rho\|\sigma) < \infty$*, then for all* $\gamma \in (0,1]$ *and* $\delta \in (0, \frac{\gamma}{2}]$

$$D_{1-\delta}(\rho\|\sigma) \geq D(\rho\|\sigma) - \ln(2)\delta(c_\gamma(\rho\|\sigma))^2 \cosh(\ln(2)\delta c_\gamma(\rho\|\sigma))\,. \tag{2.3.8}$$

and for all $\delta \in (0, \frac{\log 3}{2c_\gamma(\rho\|\sigma)}]$*:*

$$D_{1-\delta}(\rho\|\sigma) \geq D(\rho\|\sigma) - \ln(2)\cosh(\log(3)/2)\delta(c_\gamma(\rho\|\sigma))^2 \tag{2.3.9}$$
$$\geq D(\rho\|\sigma) - \delta(c_\gamma(\rho\|\sigma))^2\,. \tag{2.3.10}$$

Proof. Note first that for (2.3.6), where we did not explicitly require $D(\rho\|\sigma) < \infty$, we can restrict to that case, as otherwise also $D_{1+\delta}(\rho\|\sigma) = \infty$ and the statement is trivial. Hence, we can assume σ to be invertible (otherwise restrict to the subspace where σ is supported). Additionally, we can also assume ρ to be full-rank on the support of σ, as $D_{1+\delta}(\rho\|\sigma) = \lim_{\varepsilon\to 0} D_{1+\delta}(\rho + \varepsilon\Pi_\sigma\|\sigma)$ for all $\delta \in (-1,\infty]$ (where Π_σ is the projection onto the support of σ), and hence the statement for non-full-rank ρ follows after taking limits. Let $\{|i\rangle\}_i$ be an orthonormal basis of $\mathcal{H}$, then define the operator $X = \rho \otimes (\sigma^{-1})^T$, with the transpose taken in that basis, and the corresponding canonical purification of ρ on $\mathcal{H}\otimes\mathcal{H}$:

$$|\phi\rangle = \sum_i \sqrt{\rho}\,|i\rangle \otimes |i\rangle\,, \tag{2.3.11}$$

where $|i\rangle$ is an orthonormal basis of $\mathcal{H}$. Then, for all $\delta \in [-1,1]$:

$$D_{1+\delta}(\rho\|\sigma) = \frac{1}{\delta}\log\Big(\langle\phi|X^\delta|\phi\rangle\Big) \tag{2.3.12}$$

and

$$D(\rho\|\sigma) = \langle\phi|\log(X)|\phi\rangle\ . \tag{2.3.13}$$

For all $t > 0$ and $\delta \in \mathbb{R}$, the first term in the expansion of t^δ around $\delta = 0$ is $\delta\ln(t)$. Hence let us write $t^\delta = 1 + \delta\ln(t) + r_\delta(t)$ where $r_\delta(t) \coloneqq t^\delta - \delta\ln(t) - 1$. Since $1 + x \leq e^x$ for all $x \in \mathbb{R}$ we see that

$$r_\delta(t) = t^\delta - \delta\ln(t) - 1 \tag{2.3.14}$$
$$\leq t^\delta + e^{-\delta\ln(t)} - 2 \tag{2.3.15}$$
$$= e^{\delta\ln(t)} + e^{-\delta\ln(t)} - 2 \tag{2.3.16}$$
$$= 2(\cosh(\delta\ln(t)) - 1) =: s_\delta(t). \tag{2.3.17}$$

It is easy to see that $s_{-\delta}(t) = s_\delta(t)$ and $s_\delta(t) = s_{\gamma\delta}(t^{1/\gamma})$ for all $\gamma \in \mathbb{R}$. Also, it is easy to verify that $s_\delta(t)$ is

monotonically increasing in t for $t \geq 1$. It also turns out to be concave in t if $\delta \leq \frac{1}{2}$ and $t \geq 3$. To see this, note that

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2} \cosh(\delta \ln(t)) = \frac{\delta}{t^2}[\delta \cosh(\delta \ln(t)) - \sinh(\delta \ln(t))]\,. \tag{2.3.18}$$

For all $\delta < 1$ this function eventually becomes negative as $t \to \infty$, and so we are looking for when it crosses the origin, which turns out to happen at $\ln(t) = \frac{1}{2\delta} \ln\left(1 + \frac{2\delta}{1-\delta}\right)$. The right-hand side of this last expression is increasing as a function of δ, since its derivative is given by

$$\frac{\mathrm{d}}{\mathrm{d}\delta} \frac{1}{2\delta} \ln\left(1 + \frac{2\delta}{1-\delta}\right) = \frac{1}{2\delta^2}\left(\frac{2\delta}{1-\delta^2} - \ln\left(1 + \frac{2\delta}{1-\delta}\right)\right) \tag{2.3.19}$$

and the positivity of this right-hand side can be seen from the well-known bound $\ln(1+x) \leq \frac{x}{\sqrt{1+x}}$. Thus, if we require $\delta \leq \frac{1}{2}$, then $s_\delta(t)$ will be convex for all t that satisfy

$$\ln(t) \geq \frac{1}{2\delta} \ln\left(1 + \frac{2\delta}{1-\delta}\right)\Big|_{\delta=\frac{1}{2}} = \ln(3)\,, \tag{2.3.20}$$

and thus $t \geq 3$.

For all $t > 0$, either t or $1/t$ will be larger than or equal to one, and so we can always use the monotonicity to write

$$s_\delta(t) = s_{-\delta}(t) = s_\delta\left(\frac{1}{t}\right) \leq s_\delta\left(t + \frac{1}{t}\right)\,. \tag{2.3.21}$$

Using this, we get for all $t > 0$ and $\gamma \in (0, 1]$

$$s_\delta(t) = s_{\delta/\gamma}(t^\gamma) \leq s_{\delta/\gamma}\left(t^\gamma + t^{-\gamma}\right) \leq s_{\delta/\gamma}\left(t^\gamma + t^{-\gamma} + 1\right)\,. \tag{2.3.22}$$

It is easy to see that the real function $x + 1/x$ has a global minimum for $x > 0$ at $x = 1$ and hence the argument on the right hand side is guaranteed to be larger than 3. Hence we can use the concavity of $s_{\delta/\gamma}$ whenever $\delta \leq \gamma/2$, to write

$$\langle\phi|s_\delta(X)|\phi\rangle \leq \langle\phi|s_{\delta/\gamma}(X^\gamma + X^{-\gamma} + \mathbb{1})|\phi\rangle \leq s_{\delta/\gamma}(\langle\phi|X^\gamma + X^{-\gamma} + \mathbb{1}|\phi\rangle) = s_{\delta/\gamma}(2^{\gamma c_\gamma(\rho\|\sigma)}), \tag{2.3.23}$$

where the second inequality is a well-known property of concave functions (Jensen's inequality) and c_γ is defined in (2.3.5). Finally, we use Taylor's theorem with the Lagrange form of the remainder to bound

$$s_\delta(t) = s_0(t) + \frac{\mathrm{d}}{\mathrm{d}\delta} s_\delta(t)\big|_{\delta=0}\delta + \frac{1}{2}\frac{\mathrm{d}^2}{\mathrm{d}\delta^2} s_\delta(t)\big|_{\delta=\xi}\delta^2 \tag{2.3.24}$$

$$= \delta^2(\ln(t))^2 \cosh(\xi \ln(t)) \leq \delta^2(\ln(t))^2 \cosh(\delta \ln(t)), \tag{2.3.25}$$

for all $t > 0$ and some $\xi \in (0, \delta)$, where we have used that $s_0(t) = \frac{\mathrm{d}}{\mathrm{d}\delta} s_\delta(t)|_{\delta=0} = 0$. Hence,

$$\langle\phi|s_\delta(X)|\phi\rangle \leq s_{\delta/\gamma}(2^{\gamma c_\gamma(\rho\|\sigma)}) \leq (\delta \ln(2) c_\gamma(\rho\|\sigma))^2 \cosh(\delta \ln(2) c_\gamma(\rho\|\sigma))\,. \tag{2.3.26}$$

To finally prove our claims, let us start with the case where $\delta > 0$. Then,

$$D_{1+\delta}(\rho\|\sigma) = \frac{1}{\delta}\log\Big(\langle\phi|X^{\delta}|\phi\rangle\Big) = \frac{1}{\delta}\log(1+\delta\ln(2)\,\langle\phi|\log(X)|\phi\rangle + \langle\phi|r_\delta(X)|\phi\rangle) \tag{2.3.27}$$

$$= \frac{1}{\delta}\log(1+\delta\ln(2)D(\rho\|\sigma) + \langle\phi|r_\delta(X)|\phi\rangle) \tag{2.3.28}$$

$$\leq \frac{1}{\delta}\log(1+\delta\ln(2)D(\rho\|\sigma) + \langle\phi|s_\delta(X)|\phi\rangle) \tag{2.3.29}$$

$$= \frac{1}{\delta}\log(1+\delta\ln(2)D(\rho\|\sigma)) + \frac{1}{\delta}\log\left(1+\frac{\langle\phi|s_\delta(X)|\phi\rangle}{1+\delta\ln(2)D(\rho\|\sigma)}\right) \tag{2.3.30}$$

$$\leq D(\rho\|\sigma) + \frac{1}{\delta}\log(1+\langle\phi|s_\delta(X)|\phi\rangle) \tag{2.3.31}$$

$$\leq D(\rho\|\sigma) + \frac{1}{\delta}\log\big(1+(\delta\ln(2)c_\gamma(\rho\|\sigma))^2\cosh(\delta\ln(2)c_\gamma(\rho\|\sigma))\big) \tag{2.3.32}$$

$$\leq D(\rho\|\sigma) + \delta\ln(2)(c_\gamma(\rho\|\sigma))^2, \tag{2.3.33}$$

where the third-to-last inequality uses $\log(1+x) \leq \frac{x}{\ln(2)}$, the fact that δ and $D(\rho\|\sigma)$ are non-negative, and that the logarithm is monotone. The final inequality follows from the fact that $k^2 - \ln\big(1+k^2\cosh(k)\big)$ is monotonically increasing in k when $k \geq 0$ and hence positive. This can be seen by noting that

$$\frac{\mathrm{d}}{\mathrm{d}k}\big[k^2 - \ln\big(1+k^2\cosh(k)\big)\big] = 2k\left[1 - \frac{\cosh(k)+\frac{k}{2}\sinh(k)}{1+k^2\cosh(k)}\right] \tag{2.3.34}$$

and

$$\cosh(k) + \frac{k}{2}\sinh(k) = 1 + \sum_{n=0}^{\infty}\left[\frac{k^{2n+2}}{(2n+2)!} + \frac{k^{2n+2}}{2(2n+1)!}\right] = 1 + \sum_{n=0}^{\infty}\frac{k^{2n+2}}{(2n)!}\frac{2n+3}{2(2n+1)(2n+2)} \tag{2.3.35}$$

$$\leq 1 + \sum_{n=0}^{\infty}\frac{k^{2n+2}}{(2n)!} = 1 + k^2\cosh(k)\,. \tag{2.3.36}$$

If $\delta < 0$, the issue arises that $1+\delta\ln(2)D(\rho\|\sigma)$ no longer has to be greater than 1 (it might in fact be negative) and so we have to apply a slightly different argument. Starting analogously, we get

$$D_{1+\delta}(\rho\|\sigma) = \frac{1}{\delta}\log(1+\delta\ln(2)D(\rho\|\sigma) + \langle\phi|r_\delta(X)|\phi\rangle) \tag{2.3.37}$$

$$\geq \frac{1}{\delta}\log(1+\delta\ln(2)D(\rho\|\sigma) + \langle\phi|s_\delta(X)|\phi\rangle) \tag{2.3.38}$$

$$\geq D(\rho\|\sigma) + \frac{1}{\delta\ln(2)}\,\langle\phi|s_\delta(X)|\phi\rangle \tag{2.3.39}$$

$$\geq D(\rho\|\sigma) + \delta\ln(2)(c_\gamma(\rho\|\sigma))^2\cosh(\delta\ln(2)c_\gamma(\rho\|\sigma)), \tag{2.3.40}$$

where we used $\log(1+x) \leq \frac{x}{\ln(2)}$ in the second inequality. Finally, $\ln(2)\cosh(\ln(3)/2) < 1$, and so if we assume that $|\delta| \leq \frac{\log 3}{2c_\gamma(\rho\|\sigma)}$ then

$$D_{1+\delta}(\rho\|\sigma) \geq D(\rho\|\sigma) + \delta(c_\gamma(\rho\|\sigma))^2\,. \tag{2.3.41}$$

Note that since $c_\gamma(\rho\|\sigma) \geq \frac{\log(3)}{\gamma}$ the condition $|\delta| \leq \gamma/2$ is automatically fulfilled. □

Remark 2.3.2. The previously known bound (Tomamichel 2012, Lemma 6.3) states only an analogue of (2.3.7) and follows from Lemma 2.3.1 after setting $\gamma = 1/2$. Note that in its analogue of (2.3.7), Tomamichel (2012), Lemma 6.3, also has a stronger constraint on the range of δ, which turns out not to be necessary.

2.3.2 The Sandwiched Rényi Divergence

The second quantum generalization of the classical Rényi divergence, which is often considered and used, is the so-called *sandwiched Rényi divergence* (Müller-Lennert et al. 2013, Wilde, Winter, and Yang 2014), which is for $\rho, \sigma \in \mathscr{P}(A)$, σ full rank, defined as

$$\widetilde{D}_\alpha(\rho\|\sigma) := \frac{1}{\alpha - 1} \log \operatorname{Tr}\left(\sigma^{\frac{1-\alpha}{2\alpha}} \rho \sigma^{\frac{1-\alpha}{2\alpha}}\right)^\alpha , \tag{2.3.42}$$

and can be defined for non full-rank σ by a limiting expression as specified for general divergences above. It satisfies the data-processing inequality for all $\alpha \in [1/2, \infty)$, and additionally whenever $\rho \in \mathcal{D}(A)$ is normalized, it is monotonically increasing in α and also converges to the quantum relative entropy for $\alpha \to 1$ (Müller-Lennert et al. 2013, Beigi 2013, Frank and Lieb 2013). One can show that it is equal to the regularized measured Rényi divergence (Mosonyi and Ogawa 2015), i.e.

$$\lim_{n\to\infty} \frac{1}{n} \sup_{\mathcal{M}_n : A^n \to \mathcal{X}} D_\alpha(\mathcal{M}_n(\rho^{\otimes n}) \| \mathcal{M}_n(\sigma^{\otimes n})) = \widetilde{D}(\rho\|\sigma) , \tag{2.3.43}$$

(where the supremum goes also over all classical systems $\mathcal{X}$), and thus for each $\alpha \in [1/2, \infty)$ the sandwiched Rényi divergence is the smallest quantum α-Rényi divergence that satisfies the data-processing inequality and is additive on tensor powers of the same state (this is also called weak additivity).

2.3.3 The Geometric Rényi Divergence

The geometric Rényi divergence is another quantum Rényi divergence, which gets its name due to its similarity to a weighted matrix geometric mean. For A, B positive *definite* matrices, and $\alpha \in \mathbb{R}$, we define the weighted matrix geometric mean as

$$A \#_\alpha B := A^{\frac{1}{2}} \left(A^{-\frac{1}{2}} B A^{-\frac{1}{2}}\right)^\alpha A^{\frac{1}{2}} . \tag{2.3.44}$$

This was defined in Kubo and Ando (1979) originally only for $\alpha \in [0, 1]$, and has many desirable properties of a matrix mean only in this range, but the definition perfectly makes sense for all α, and we will make use of it also for α outside of $[0, 1]$. For $\rho, \sigma \in \mathscr{P}(\mathcal{H})$, $\rho \ll \sigma$ and $\alpha \in (0, \infty)$ then define the *geometric trace function* as

$$\widehat{Q}_\alpha(\rho\|\sigma) := \operatorname{Tr}(\sigma \#_\alpha \rho) = \operatorname{Tr}\left(\sigma^{\frac{1}{2}} (\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}})^\alpha \sigma^{\frac{1}{2}}\right) . \tag{2.3.45}$$

Similar to the discussion above, this can be extended to general $\rho, \sigma \in \mathscr{P}(\mathcal{H})$ by setting (Matsumoto 2018)

$$\widehat{Q}_\alpha(\rho\|\sigma) \coloneqq \lim_{\varepsilon\to 0} \widehat{Q}_\alpha(\rho\|\sigma + \varepsilon\mathbb{1}) \,. \tag{2.3.46}$$

This limit can potentially be infinite, and will be so if $\alpha > 1$ and $\rho \not\ll \sigma$, whereas for $\alpha \in (0, 1]$ the limit is always finite and explicit expressions for it can be found in Matsumoto (2018), Hiai and Mosonyi (2017), and Katariya and Wilde (2020). For $\alpha \in (0, 1) \cup (1, \infty)$ one then defines the *geometric Rényi divergence* as (Matsumoto 2018, Petz and Ruskai 1998, Fang and H. Fawzi 2021, Katariya and Wilde 2020)

$$\widehat{D}_\alpha(\rho\|\sigma) \coloneqq \frac{1}{\alpha - 1} \log \widehat{Q}_\alpha(\rho\|\sigma) \,. \tag{2.3.47}$$

This reduces to the classical α-Rényi divergence for commuting states and satisfies the data-processing inequality for $\alpha \in (0, 1) \cup (1, 2]$ (Matsumoto 2018). Further, it is known (Matsumoto 2018, Tomamichel 2016, Katariya and Wilde 2020) that for $\rho \in \mathcal{D}(\mathcal{H})$ and $\sigma \in \mathscr{P}(\mathcal{H})$

$$\lim_{\alpha\to 1} \widehat{D}_\alpha(\rho\|\sigma) = \widehat{D}(\rho\|\sigma) \coloneqq \operatorname{Tr}\left(\rho \log\left(\rho^{\frac{1}{2}}\sigma^{-1}\rho^{\frac{1}{2}}\right)\right), \tag{2.3.48}$$

the Belavkin-Staszewski relative entropy (Belavkin and Staszewski 1982).

The geometric Rényi divergence was originally introduced by Matsumoto (2018) coming from a slightly different angle, namely trying to find the largest quantum Rényi divergence for each α, going through a construction he called reverse tests. For $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, we say that a tuple $(\mathcal{X}, P, Q, \Lambda)$ of a classical sample space $\mathcal{X}$, two probability distributions P, Q on $\mathcal{X}$, and a classical-quantum channel $\Lambda : \mathcal{X} \to \mathcal{H}$, forms a reverse test of (ρ, σ), if $\Lambda(P) = \rho$, $\Lambda(Q) = \sigma$. Then, consider the following expression

$$\widehat{D}^{\max}_\alpha(\rho\|\sigma) \coloneqq \inf_{\mathcal{X},P,Q,\Lambda} \{\, D_\alpha(P\|Q) \mid (\mathcal{X}, P, Q, \Lambda) \text{ form a reverse test of } (\rho, \sigma) \,\} \,, \tag{2.3.49}$$

where $D_\alpha(P\|Q)$ is the classical α-Rényi divergence, i.e. we look for the smallest classical α-Rényi divergence between two probability distributions which are such that one can recover the states ρ and σ from them via a classical-quantum channel. This construction ensures that:

1. $\widehat{D}^{\max}_\alpha$ reduces to the classical α-Rényi divergence when evaluated on classical probability distributions. In this case one optimizing reverse test is to pick P and Q as exactly these two classical distributions, and Λ the direct embedding of the classical system $\mathcal{X}$ into the Hilbert space $\mathcal{H}$ as explained in Section 2.1.

2. $\widehat{D}^{\max}_\alpha$ satisfies the data-processing inequality for all $\alpha \in (0, \infty)$, since, for any quantum channel $\mathcal{E}$, whenever $(\mathcal{X}, P, Q, \Lambda)$ forms a reverse test of (ρ, σ), then $(\mathcal{X}, P, Q, \mathcal{E} \circ \Lambda)$ forms a reverse test of $(\mathcal{E}(\rho), \mathcal{E}(\sigma))$, and since these reverse tests have identical Ps and Qs, they achieve the same value in the optimization problem.

3. For every $\alpha \in (0, \infty)$, $\widehat{D}^{\max}_\alpha$ is the largest quantum α-Rényi divergence that satisfies the data-processing inequality. This can be seen as follows: Let $\mathbf{D}_\alpha$ be any quantum α-Rényi divergence (as defined in Section 2.3, in particular that that means $\mathbf{D}_\alpha$ satisfies the data-processing inequality), and let

$(\mathcal{X}, P, Q, \Lambda)$ be a reverse test of (ρ, σ) that gets δ-close to achieving the optimal value. Then

$$\mathbf{D}_\alpha(\rho\|\sigma) = \mathbf{D}_\alpha(\Lambda(P)\|\Lambda(Q)) \leq \mathbf{D}_\alpha(P\|Q) = D_\alpha(P\|Q) = \widehat{D}_\alpha^{\max}(\rho\|\sigma) + \delta \xrightarrow{\delta\to 0} \widehat{D}_\alpha^{\max}(\rho\|\sigma) \tag{2.3.50}$$

where we used that all quantum α-Rényi divergences are identical on classical input states.

We hence call $\widehat{D}_\alpha^{\max}(\rho\|\sigma)$ the maximal Rényi divergence. Matsumoto (2018) showed that for $\alpha \in (0, 2]$ (but only for this range of α) this maximal Rényi divergence is equal to the geometric Rényi divergence as defined explicitly above, i.e. for $\alpha \in (0, 2]$, $\widehat{D}_\alpha^{\max}(\rho\|\sigma) = \widehat{D}_\alpha(\rho\|\sigma)$. Hence, these two names are often used interchangeably, although we would like to emphasize that the maximal Rényi divergence satisfies the data-processing inequality for all α, whereas the geometric Rényi divergence (when defined using the explicit formula above) does not for $\alpha > 2$.

One of the key properties of the geometric Rényi divergence is that its channel version satisfies a chain rule (Fang and H. Fawzi 2021, Thm 3.4). If $\rho_{RA}, \sigma_{RA} \in \mathcal{D}(RA)$ are two quantum states, and $\mathcal{E}, \mathcal{F} : A \to B$ are two quantum channels, then it holds that

$$\widehat{D}_\alpha(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) \leq \widehat{D}_\alpha(\rho\|\sigma) + \widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) \tag{2.3.51}$$

where the last term on the right-hand side is the corresponding channel divergence, as defined in (2.2.22).

2.3.4 The Sharp Rényi Divergence

The sharp Rényi divergence is yet another quantum Rényi divergence, defined by H. Fawzi and O. Fawzi (2021) in terms of a convex optimization problem. For $\alpha \in (1, \infty)$, set:

$$Q_\alpha^\#(\rho\|\sigma) = \inf_{A\geq 0} \left\{ \operatorname{Tr} A \;\middle|\; \rho \leq \sigma \#_{\frac{1}{\alpha}} A \right\}, \tag{2.3.52}$$

and then define the sharp Rényi divergence of order α in terms of it as follows:

$$D_\alpha^\#(\rho\|\sigma) := \frac{1}{\alpha - 1} \log\Big(Q_\alpha^\#(\rho\|\sigma)\Big). \tag{2.3.53}$$

H. Fawzi and O. Fawzi (2021) proved that this quantity satisfies the data-processing inequality, and also that it reduces to the classical α-Rényi divergence for commuting states. Further, they showed that this divergence has several desirable computational properties such as an efficient semi-definite programming representation for states and quantum channels, and a crucial chain rule property which can be exploited to obtain results concerning quantum channel discrimination and quantum channel capacities. Note that unlike other Rényi divergences though, the sharp quantum Rényi divergence is not monotonically increasing in α.

H. Fawzi and O. Fawzi (2021) originally, left the question of "*What is* $\lim_{\alpha\to 1} D_\alpha^\#(\rho\|\sigma)$*?*" open, although we were able to show that this limit is given by the Belavkin-Staszewski relative entropy (see Theorem 2.3.9 below).

To address the limit of $\alpha \to 1$ we want to make use of the following alternative characterization of the sharp trace function and sharp Rényi divergence.

Proposition 2.3.3. *For $\alpha \in (1, \infty)$ and $\rho, \sigma \in \mathscr{P}(\mathcal{H})$*

$$Q_\alpha^\#(\rho\|\sigma) = \min_{A \geq \rho} \widehat{Q}_\alpha(A\|\sigma), \tag{2.3.54}$$

$$D_\alpha^\#(\rho\|\sigma) = \min_{A \geq \rho} \widehat{D}_\alpha(A\|\sigma). \tag{2.3.55}$$

Remark 2.3.4. Note that the A in the above expressions are in general unnormalized states and we use the definitions of the geometric trace function and the geometric Rényi divergence as in (2.3.45) and (2.3.47) without additional normalization factors.

For what is going to follow, we prove a slightly stronger version of the above statement, given by the following lemma.

Lemma 2.3.5. *For $\alpha \in (1, a]$, with $a \in (1, \infty)$, and $\rho, \sigma \in \mathscr{P}(\mathcal{H})$,*

$$Q_\alpha^\#(\rho\|\sigma) = \min_{A \geq \rho} \widehat{Q}_\alpha(A\|\sigma), \tag{2.3.56}$$

and the minimization can be restricted to $A \in K$ where K is a compact subset of $\{ A \in \mathscr{P}(\mathcal{H}) \mid A \ll \sigma \}$ depending only on ρ, σ, a, but not on α.

Proof of Lemma 2.3.5. It is easy to see (and shown in H. Fawzi and O. Fawzi (2021)) that if $\rho \not\ll \sigma$, $Q_\alpha^\#(\rho\|\sigma) = +\infty$ since the minimization is infeasible. Moreover, if $\rho \not\ll \sigma$ and $A \geq \rho$ then also $A \not\ll \sigma$, hence $\min_{A \geq \rho} \widehat{Q}_\alpha(A\|\sigma) = +\infty$ and the statement holds. Thus, we can assume $\rho \ll \sigma$. Recall the definition of $Q_\alpha^\#(\rho\|\sigma)$ for $\alpha > 1$ (H. Fawzi and O. Fawzi 2021):

$$Q_\alpha^\#(\rho\|\sigma) = \inf_{A \geq 0} \left\{ \operatorname{Tr} A \;\middle|\; \rho \leq \sigma^{\frac{1}{2}} (\sigma^{-\frac{1}{2}} A \sigma^{-\frac{1}{2}})^{\frac{1}{\alpha}} \sigma^{\frac{1}{2}} \right\}. \tag{2.3.57}$$

H. Fawzi and O. Fawzi (2021) showed that the minimization can be further restricted to $0 \leq A \leq c^{\alpha-1} \operatorname{Tr}(\rho) \Pi_\sigma$, where $c = \|\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}}\|$ is the spectral norm, and Π_σ is the projection onto the support of σ. In the following we will perform several redefinitions of the dummy variable A in this optimization problem, where the inequalities between the different lines always arise from the fact that we have enlarged the space over which is optimized. Let us define $C \coloneqq \operatorname{Tr}(\rho) \sup_{\alpha \in [1,a]} c^{\alpha-1}$, which is always finite. We then get:

$$Q_\alpha^\#(\rho\|\sigma) = \min_{A_1 \geq 0} \left\{ \operatorname{Tr} A_1 \;\middle|\; \rho \leq \sigma^{\frac{1}{2}} (\sigma^{-\frac{1}{2}} A_1 \sigma^{-\frac{1}{2}})^{\frac{1}{\alpha}} \sigma^{\frac{1}{2}},\ A_1 \leq \operatorname{Tr}(\rho) c^{\alpha-1} \Pi_\sigma \right\} \tag{2.3.58}$$

$$\geq \min_{A_1 \geq 0} \left\{ \operatorname{Tr} A_1 \;\middle|\; \rho \leq \sigma^{\frac{1}{2}} (\sigma^{-\frac{1}{2}} A_1 \sigma^{-\frac{1}{2}})^{\frac{1}{\alpha}} \sigma^{\frac{1}{2}},\ A_1 \leq C \Pi_\sigma \right\} \tag{2.3.59}$$

$$\overset{A_2 := \sigma^{-\frac{1}{2}} A_1 \sigma^{-\frac{1}{2}}}{=} \min_{A_2 \geq 0} \left\{ \operatorname{Tr}\left(\sigma^{\frac{1}{2}} A_2 \sigma^{\frac{1}{2}}\right) \;\middle|\; \rho \leq \sigma^{\frac{1}{2}} A_2^{\frac{1}{\alpha}} \sigma^{\frac{1}{2}},\ A_2 \leq C \sigma^{-\frac{1}{2}} \Pi_\sigma \sigma^{-\frac{1}{2}} \right\} \tag{2.3.60}$$

$$\geq \min_{A_2 \geq 0} \left\{ \operatorname{Tr}\left(\sigma^{\frac{1}{2}} A_2 \sigma^{\frac{1}{2}}\right) \;\middle|\; \rho \leq \sigma^{\frac{1}{2}} A_2^{\frac{1}{\alpha}} \sigma^{\frac{1}{2}},\ A_2 \leq C \|\sigma^{-1}\| \Pi_\sigma \right\} \tag{2.3.61}$$

$$\overset{A_3 := A_2^{\frac{1}{\alpha}}}{=} \min_{A_3 \geq 0} \left\{ \operatorname{Tr}\left(\sigma^{\frac{1}{2}} A_3^\alpha \sigma^{\frac{1}{2}}\right) \;\middle|\; \rho \leq \sigma^{\frac{1}{2}} A_3 \sigma^{\frac{1}{2}},\ A_3^\alpha \leq C \|\sigma^{-1}\| \Pi_\sigma \right\} \tag{2.3.62}$$

where we redefined the dummy variable A multiple times, and understand $\|\sigma^{-1}\|$ as the operator norm of the pseudo-inverse. For $\alpha > 1$, $A \mapsto A^{\frac{1}{\alpha}}$ is operator monotone, and thus $A_3^\alpha \leq C\|\sigma^{-1}\|\Pi_\sigma$ implies $A_3 \leq (C\|\sigma^{-1}\|)^{\frac{1}{\alpha}}\Pi_\sigma$. Thus, we get:

$$Q_\alpha^\#(\rho\|\sigma) \geq \min_{A_3\geq 0}\left\{ \operatorname{Tr}\left(\sigma^{\frac{1}{2}}A_3^\alpha\sigma^{\frac{1}{2}}\right) \,\middle|\, \rho \leq \sigma^{\frac{1}{2}}A_3\sigma^{\frac{1}{2}},\ A_3 \leq (C\|\sigma^{-1}\|)^{\frac{1}{\alpha}}\Pi_\sigma \right\} \tag{2.3.63}$$

$$\overset{A:=\sigma^{\frac{1}{2}}A_3\sigma^{\frac{1}{2}}}{=} \min_{A\geq 0}\left\{ \widehat{Q}_\alpha(A\|\sigma) \,\middle|\, \rho \leq A,\ A \leq (C\|\sigma^{-1}\|)^{\frac{1}{\alpha}}\sigma^{\frac{1}{2}}\Pi_\sigma\sigma^{\frac{1}{2}} \leq (C\|\sigma^{-1}\|)^{\frac{1}{\alpha}}\|\sigma\|\Pi_\sigma \right\}. \tag{2.3.64}$$

With $\tilde{C} := \sup_{\alpha\in[1,a]}(C\|\sigma^{-1}\|)^{\frac{1}{\alpha}}\|\sigma\|$, which exists, and $K := \{\, A \in \mathscr{P}(\mathcal{H}) \mid \rho \leq A \leq \tilde{C}\Pi_\sigma \,\}$, which is compact, we have for all $\alpha \in (1, a]$:

$$Q_\alpha^\#(\rho\|\sigma) \geq \min_{A\in K} \widehat{Q}_\alpha(A\|\sigma)\,. \tag{2.3.65}$$

For the reverse direction, note that for every $\alpha > 1$ the geometric trace function is the largest α-Rényi trace function that satisfies the data-processing inequality, and hence $Q_\alpha^\#(\rho\|\sigma) \leq \widehat{Q}_\alpha(\rho\|\sigma)$ (alternatively it is easy to see this explicitly, by noting that $\sigma\#_s(\sigma\#_t\rho) = \sigma\#_{st}\rho$ for all $s, t \in \mathbb{R}$, and hence $A = \sigma\#_\alpha\rho$ is feasible in (2.3.52)). Moreover, from the definition (2.3.52) it is easy to see that

$$\min_{A\geq\rho} Q_\alpha^\#(A\|\sigma) = Q_\alpha^\#(\rho\|\sigma) = \min_{A\in K} Q_\alpha^\#(A\|\sigma)\,, \tag{2.3.66}$$

and hence, also

$$Q_\alpha^\#(\rho\|\sigma) \leq \min_{A\in K} \widehat{Q}_\alpha(A\|\sigma)\,. \tag{2.3.67}$$

□

Since the logarithm on $\mathbb{R}$ is monotone, Lemma 2.3.5 implies that also $D_\alpha^\#(\rho\|\sigma) = \min_{A\geq\rho}\widehat{D}_\alpha(A\|\sigma)$ for all $\alpha > 1$. This completes the proof of Proposition 2.3.3. Note, that for $\alpha > 2$ the geometric divergence does not satisfy the data-processing inequality, while the sharp divergence does.

The $\alpha \to 1$ limit

We show in Theorem 2.3.9 at the end of this section that the limit $\alpha \to 1$ of the sharp divergence is the Belavkin-Staszewski relative entropy. The key step in the proof is to show that $\frac{\mathrm{d}}{\mathrm{d}\alpha}Q_\alpha^\#(\rho\|\sigma)|_{\alpha=1}$ exists and is equal to $\frac{\mathrm{d}}{\mathrm{d}\alpha}\widehat{Q}_\alpha(\rho\|\sigma)|_{\alpha=1}$. Conditions for such an exchange of minimization and taking a derivative have been established in the literature. The following statement directly fits our purposes, and seems to have been among the earliest of such theorems, although the authors seem to leave some parts of the proof up to the reader. Later generalizations which imply this statement (and also with more explicit proofs) can for example be found in Clarke (1975), Theorem 2.1, and also in the references therein.

Theorem 2.3.6 (Dem'yanov and Malozemov (1971), Theorem 2.2.1). *Let $U \subset \mathbb{R}^n$ be open, $K \subset \mathbb{R}^m$ compact, and $f\colon U \times K \to \mathbb{R}$ continuous and also that $\nabla_u f(u,k)$ is (jointly) continuous. Then, the function $g(u) = \min_{k\in K} f(u,k)$ has for every $u \in U$ a one-sided directional derivative along every $v \in \mathbb{R}^n$, which can*

be computed as

$$\lim_{t \searrow 0} \frac{g(u+vt) - g(u)}{t} = \min_{k \in R(u)} \langle \nabla_u f(u,k), v \rangle, \tag{2.3.68}$$

where $R(u) = \{\, k \in K \mid f(u,k) = g(u) \,\}$.[1]

The following lemmas establish the properties of $\widehat{Q}_\alpha$ which are necessary in order to apply Theorem 2.3.6.

Lemma 2.3.7. *Let* Ω *be a subset of* $\mathbb{R}^n$ *and* $h \colon \Omega \times [0,\infty) \to [0,\infty)$ *be a jointly continuous function. Then, the corresponding matrix function* $h \colon \Omega \times \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{H})$ *is jointly continuous on* $\Omega \times \mathscr{P}(\mathcal{H})$.

Proof. For a fixed $\alpha \in \Omega$, the continuity of the matrix function $h(\alpha, \rho)$ in ρ follows from the continuity of $h(\alpha, x)$ in x, for $x \in [0,\infty)$, by Horn and Johnson (1991), Theorem 6.2.37. To see joint continuity, let $\rho_n \to \rho$ be a converging sequence in $\mathscr{P}(\mathcal{H})$ and $\alpha_n \to \alpha$ be a converging sequence in Ω, as $n \to \infty$. Then, taking operator norms,

$$\|h(\alpha_n, \rho_n) - h(\alpha, \rho)\| \leq \|h(\alpha_n, \rho_n) - h(\alpha, \rho_n)\| + \|h(\alpha, \rho_n) - h(\alpha, \rho)\| \,. \tag{2.3.69}$$

The second term on the right hand side goes to zero as $n \to \infty$ by the continuity of $h(\alpha, \rho)$ in ρ, as established above. For the first term, note that $h(\alpha_n, \rho_n)$ and $h(\alpha, \rho_n)$ commute and the operator norm can be evaluated as

$$\max_{\lambda \in \operatorname{spec}(\rho_n)} |h(\alpha_n, \lambda) - h(\alpha, \lambda)| \,. \tag{2.3.70}$$

Since the limit $\alpha_n \to \alpha$ exists, α_n is eventually bounded, so there exists a compact subset K of Ω, such that $\alpha_n \in K$ for all large enough n. Analogously, the limit $\rho_n \to \rho$ exists, and hence there exists a compact subset K' of $[0,\infty)$ such that $\operatorname{spec}(\rho_n) \subset K'$ for all large enough n. Since $K \times K'$ is compact, h is in fact uniformly continuous on $K \times K'$. Hence, (2.3.70) goes to zero as $n \to \infty$. This shows that

$$\lim_{n \to \infty} \|h(\alpha_n, \rho_n) - h(\alpha, \rho)\| = 0 \,, \tag{2.3.71}$$

which completes the proof. □

Lemma 2.3.8. *For* $\sigma \in \mathscr{P}(\mathcal{H})$ *fixed, the functions* $f_\rho(\alpha) = f(\alpha, \rho) = \widehat{Q}_\alpha(\rho \| \sigma)$ *and* $f'_\rho(\alpha) = f'(\alpha, \rho) = \frac{\mathrm{d}}{\mathrm{d}\alpha} \widehat{Q}_\alpha(\rho \| \sigma)$ *are jointly continuous on* $(0,\infty) \times \{\, A \in \mathscr{P}(\mathcal{H}) \mid A \ll \sigma \,\}$.

Proof. We have:

$$\widehat{Q}_\alpha(\rho \| \sigma) = \operatorname{Tr}\left(\sigma^{\frac{1}{2}} \left(\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}} \right)^\alpha \sigma^{\frac{1}{2}} \right) \tag{2.3.72}$$

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} \widehat{Q}_\alpha(\rho \| \sigma) = \operatorname{Tr}\left(\sigma^{\frac{1}{2}} \log\left(\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}} \right) \left(\sigma^{-\frac{1}{2}} \rho \sigma^{-\frac{1}{2}} \right)^\alpha \sigma^{\frac{1}{2}} \right) . \tag{2.3.73}$$

It is easy to check that the real functions $(\alpha, x) \mapsto x^\alpha$ and $(\alpha, x) \mapsto x^\alpha \log(x)$ are jointly continuous on $(0,\infty) \times [0,\infty)$. Also, for $A \in \{\, A \in \mathscr{P}(\mathcal{H}) \mid A \ll \sigma \,\}$ the expression $\sigma^{-\frac{1}{2}} A \sigma^{-\frac{1}{2}}$ is continuous in A

[1] In Dem'yanov and Malozemov (1971) the theorem is phrased with a maximum instead of a minimum, but it is easy to see that this is equivalent upon setting $f \mapsto -f$.

(remember that the inverses are all pseudo-inverses). Hence, the continuity follows from Lemma 2.3.7, the continuity of the matrix product and the continuity of the trace. □

Theorem 2.3.9. *For $\rho \in \mathcal{D}(\mathcal{H})$ and $\sigma \in \mathscr{P}(\mathcal{H})$:*

$$\lim_{\alpha \to 1} D_\alpha^\#(\rho\|\sigma) = \widehat{D}(\rho\|\sigma) \tag{2.3.74}$$

the Belavkin-Staszewski relative entropy.

Proof. As the sharp Rényi divergence, $D_\alpha^\#$, is only defined for $\alpha > 1$, all there is to prove is a limit from above. Let us fix $\rho \in \mathcal{D}(\mathcal{H})$ and $\sigma \in \mathscr{P}(\mathcal{H})$. We can restrict our proof to the case $\rho \ll \sigma$, since otherwise $D_\alpha^\#(\rho\|\sigma) = \widehat{D}_\alpha(\rho\|\sigma) = \infty$ for all $\alpha > 1$ and the statement clearly holds. By Lemma 2.3.5 there exists a compact set $K \subseteq \{\, A \in \mathscr{P}(\mathcal{H}) \mid A \ll \sigma \,\}$[2] such that for $\alpha \in (1, 2]$,

$$Q_\alpha^\#(\rho\|\sigma) = \min_{A \in K} \widehat{Q}_\alpha(A\|\sigma)\,. \tag{2.3.75}$$

For $\alpha \in (0, \infty)$,[3] we define

$$f(\alpha, A) \coloneqq \widehat{Q}_\alpha(A\|\sigma) \tag{2.3.76}$$

$$g(\alpha) \coloneqq \min_{A \in K} f(\alpha, A) = \min_{A \in K} \widehat{Q}_\alpha(A\|\sigma)\,. \tag{2.3.77}$$

By Lemma 2.3.8, $f(\alpha, A)$ is jointly continuous on $(0, \infty) \times K$ and has a jointly continuous derivative in α. Hence by Theorem 2.3.6, $g(\alpha)$ has the the following one-sided derivatives:

$$\lim_{\alpha \searrow 1} \frac{g(\alpha) - g(1)}{\alpha - 1} = \min_{A \in R(1)} \frac{\mathrm{d}}{\mathrm{d}\alpha} f(\alpha, A)\Big|_{\alpha=1} \tag{2.3.78}$$

$$\lim_{\alpha \nearrow 1} \frac{g(\alpha) - g(1)}{\alpha - 1} = -\min_{A \in R(1)} \left(-\frac{\mathrm{d}}{\mathrm{d}\alpha} f(\alpha, A)\Big|_{\alpha=1} \right), \tag{2.3.79}$$

where $R(\alpha) \coloneqq \{\, A \in K \mid f(\alpha, A) = \min_{B \in K} f(\alpha, B) \,\}$. Recall that the sharp Rényi divergence is only defined for $\alpha > 1$ and so all we really need here is the limit from above (2.3.78). However, establishing that this is also equal to the limit from below makes things slightly simpler, as we are then able to directly use the chain rule later in (2.3.82).

For $\alpha = 1$ the set $R(\alpha)$ only contains ρ, since

$$\min_{A \geq \rho} \widehat{Q}_1(A\|\sigma) = \min_{A \geq \rho} \mathrm{Tr}(A) = \mathrm{Tr}(\rho), \tag{2.3.80}$$

and $A \geq \rho$, together with $\mathrm{Tr}(A) = \mathrm{Tr}(\rho)$, imply that $A = \rho$. Hence, the two one-sided derivatives are equal

[2] K depends on ρ and σ, but since they are fixed in the entire proof we do not make this dependence explicit.
[3] To apply Theorem 2.3.6 we require an open interval in α which includes 1.

and

$$\frac{\mathrm{d}}{\mathrm{d}\alpha}g(\alpha)\Big|_{\alpha=1} = \frac{\mathrm{d}}{\mathrm{d}\alpha}f(\alpha,\rho)\Big|_{\alpha=1} = \frac{\mathrm{d}}{\mathrm{d}\alpha}\widehat{Q}_\alpha(\rho\|\sigma)\Big|_{\alpha=1}. \tag{2.3.81}$$

We further have $g(\alpha) = Q_\alpha^\#(\rho\|\sigma)$ for $\alpha \in (1,2]$, and $g(1) = \mathrm{Tr}(\rho) = 1$, so

$$\begin{aligned}\lim_{\alpha\searrow 1} D_\alpha^\#(\rho\|\sigma) &= \lim_{\alpha\searrow 1}\frac{\log Q_\alpha^\#(\rho\|\sigma)}{\alpha-1} = \lim_{\alpha\searrow 1}\frac{\log g(\alpha)-\log g(1)}{\alpha-1}\\ &= \frac{\mathrm{d}}{\mathrm{d}\alpha}\log g(\alpha)\Big|_{\alpha=1} = \frac{\frac{\mathrm{d}}{\mathrm{d}\alpha}g(\alpha)\big|_{\alpha=1}}{g(1)} = \frac{\mathrm{d}}{\mathrm{d}\alpha}g(\alpha)\Big|_{\alpha=1}.\end{aligned} \tag{2.3.82}$$

As the same argument also gives

$$\lim_{\alpha\to 1}\widehat{D}_\alpha(\rho\|\sigma) = \frac{\mathrm{d}}{\mathrm{d}\alpha}\widehat{Q}_\alpha(\rho\|\sigma)\Big|_{\alpha=1}, \tag{2.3.83}$$

we get by using (2.3.81)

$$\lim_{\alpha\to 1} D_\alpha^\#(\rho\|\sigma) = \lim_{\alpha\to 1}\widehat{D}_\alpha(\rho\|\sigma) = \widehat{D}(\rho\|\sigma). \tag{2.3.84}$$

□

Relations between different quantum Rényi divergences

All of the introduced quantum Rényi divergences are generally different whenever the two states ρ, and σ don't commute. Regarding relations between them, in the range $\alpha \in [1/2, 2]$ one can show the following (see e.g. Khatri and Wilde (2020) and references therein for the original arguments):

$$\widetilde{D}_\alpha(\rho\|\sigma) \le D_\alpha(\rho\|\sigma) \le \widehat{D}_\alpha(\rho\|\sigma). \tag{2.3.85}$$

2.4 Hypothesis Testing and Error Exponents

Hypothesis testing is one of the most fundamental tasks in statistics and information theory. Throughout this thesis we will only be dealing with binary hypothesis testing, i.e. where one has to decide between exactly two different hypotheses, usually called the null hypothesis and the alternative hypothesis, which we will denote as H_0 and H_1 respectively. The essence of hypothesis testing is that one is presented with some data (i.e. samples from a probability distribution) and the task is to identify which of the two hypotheses produced this data. In general, given a hypothesis, the data one obtains from it will not be deterministic, and so also one will not always be able to perfectly identify the hypothesis given the data, but one can still try to find an identification strategy that (in a certain sense) minimizes the probability of making an error. Given a fixed procedure through which one claims H_0 or H_1 to be true (we call this a decision strategy), one can define the

following two errors associated with this strategy:

$$\alpha = \mathbb{P}[\text{claim } H_1 | H_0 \text{ is true}] \qquad \text{called the type-I error,} \tag{2.4.1}$$

$$\beta = \mathbb{P}[\text{claim } H_0 | H_1 \text{ is true}] \qquad \text{called the type-II error.} \tag{2.4.2}$$

Throughout this thesis we will use the phrases "type-I/II error" and "type-I/II error probability" interchangeably. The exact values of these errors of course depend on both the decision strategy, and also the actual hypotheses and what data they produce with what probability. It is easy to see that there almost always exists a trade-off between the two errors, for example one could choose a strategy that always claims H_0, and thus never makes a type-I error. Hence, when one talks about minimizing these two errors, one has to specify in what sense one wants to achieve this. There are typically two main ways to minimizes errors:

- The so-called symmetric setting, which corresponds to minimizing the average (or sometimes also the p-weighted average) of the two errors, i.e.

$$p_{\text{err}} \coloneqq \frac{1}{2}(\alpha + \beta) \tag{2.4.3}$$

- The so-called asymmetric setting, where one considers one of the two errors more grave than the other (say in a medical setting, where claiming an ill patient to be healthy is usually worse than falsely identifying a healthy patient to be potentially ill, since in this case additional investigations would then usually be used to confirm this diagnosis with higher certainty). In this case, one minimizes only one of the errors (typically the type-II error, although this is arbitrary, since one can also exchange the two hypotheses), and only requires that the other error (typically then the type-I error) stays below a certain threshold, which is usually called ε.

Throughout this thesis we will almost exclusively focus on the asymmetric setting. For a summary on the currently encountered challenges when dealing with the symmetric setting in quantum channel discrimination, see Section 3.4.2.

2.4.1 Simple Classical Hypothesis Testing

In the classical setting, the two hypotheses typically correspond to two probability distributions, from which samples are observed and then form the basis of a decision. The decision strategy is then determined by a subset $B \subset \mathcal{X}$ of the possible sample outcomes, such that whenever the observed sample is in this subset one will claim H_0 (and conversely one will claim H_1 if a sample is observed that is not in the subset). If the null hypothesis H_0 is that our samples are distributed according to a probability distribution P, and the alternative hypothesis is that the samples are distributed according to Q, then we have the following

expression for the two error probabilities

$$\alpha = \sum_{x \notin B} P(x) \tag{2.4.4}$$

$$\beta = \sum_{x \in B} Q(x)\,, \tag{2.4.5}$$

and the corresponding optimal asymmetric error probability for type-I error threshold ε is then given by:

$$\min_{B \subset \mathcal{X}} \left\{ \sum_{x \in B} Q(x) \;\middle|\; \sum_{x \notin B} P(x) \leq \varepsilon \right\}. \tag{2.4.6}$$

2.4.2 Simple Quantum Hypothesis Testing

In simple quantum hypothesis testing (also called quantum state discrimination), instead of being given samples from a classical probability distribution we are given a (usually mixed) quantum state. This might seem counter-intuitive, since the quantum states we are given are fixed, and not drawn from some probability distribution, however we cannot use the state itself to make a decision, but rather have to perform a measurement on it first, whose outcomes will then again be probabilistic. The analogue of the classical decision region is thus an operator $0 \leq M \leq \mathbb{1}_{\mathcal{H}}$, which can be completed to a 2-outcome POVM $\{M, \mathbb{1}_{\mathcal{H}} - M\}$. If we again say that if H_0 is true we are given the quantum state $\rho \in \mathcal{D}(\mathcal{H})$, and if H_1 is true we are given $\sigma \in \mathcal{D}(\mathcal{H})$, then we can calculate the two error probabilities as follows

$$\alpha = \mathrm{Tr}((\mathbb{1}_{\mathcal{H}} - M)\rho) = 1 - \mathrm{Tr}(M\rho)\,, \tag{2.4.7}$$

$$\beta = \mathrm{Tr}(M\sigma)\,, \tag{2.4.8}$$

and we have the following expression for the optimal asymmetric error probability for type-I error threshold ε:

$$\min_{0 \leq M \leq \mathbb{1}_{\mathcal{H}}} \{\, \mathrm{Tr}(M\sigma) \mid 1 - \mathrm{Tr}(M\rho) \leq \varepsilon \,\} = \min_{\substack{0 \leq M \leq \mathbb{1}_{\mathcal{H}} \\ \mathrm{Tr}(M\rho) \geq 1-\varepsilon}} \mathrm{Tr}(M\sigma) \tag{2.4.9}$$

where the minimum is achieved since the set of operators bounded by the identity (in finite dimensions) is compact.

2.4.3 The Hypothesis Testing Relative Entropy

The optimal asymmetric error probability is one of the key quantities considered in this thesis, however we will usually consider not the quantity itself, but the negative logarithm of it:

$$D_H^{\varepsilon}(\rho\|\sigma) \coloneqq -\log \left[\min_{\substack{0 \leq M \leq \mathbb{1}_{\mathcal{H}} \\ \mathrm{Tr}(M\rho) \geq 1-\varepsilon}} \mathrm{Tr}(M\sigma) \right] = \max_{\substack{0 \leq M \leq \mathbb{1}_{\mathcal{H}} \\ \mathrm{Tr}(M\rho) \geq 1-\varepsilon}} -\log(\mathrm{Tr}(M\sigma))\,. \tag{2.4.10}$$

which is known as the hypothesis testing relative entropy (L. Wang and Renner 2012). It has been given this name, since it shares some properties with the quantum relative entropy, in particular it satisfies the data-processing inequality (L. Wang and Renner 2012).

Additionally, it satisfies the following relation with the Umegaki quantum relative entropy:

Lemma 2.4.1 (Upper bound on D_H^ε). *Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be quantum states. Then for all $\varepsilon \in [0, 1)$*

$$D_H^\varepsilon(\rho\|\sigma) \leq \frac{1}{1-\varepsilon}\big(D(\rho\|\sigma) + h(\varepsilon)\big), \tag{2.4.11}$$

where $h(\varepsilon)$ is the binary entropy function.

This is a very well-known consequence of the data-processing inequality of the relative entropy and has been widely used in converse proofs in information theory. A statement and proof using our notation can be found for example in L. Wang and Renner (2012), although the essence of the statement can already be found much earlier, for example in Hiai and Petz (1991), Theorem 2.2 and also Hayashi (2006), (3.30).

2.4.4 Asymptotics and the (Quantum) Stein's Lemma

One often considered case in hypothesis testing is the probability distributions (or quantum states) associated to the hypotheses being independent and identically distributed (i.i.d.), i.e. one obtains multiple independent samples from the same probability distribution. In that case, the asymmetric error probability will decay exponentially in the number of samples (usually denoted as n), and one can study the asymptotic rate of this decay, which is given by

$$\lim_{n\to\infty} \frac{1}{n} D_H^\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n})\,. \tag{2.4.12}$$

In fact, what one often considers is the following

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty} \frac{1}{n} D_H^\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n})\,, \tag{2.4.13}$$

which is the optimal achievable decay rate of the type-II error, such that also the type-I error goes to zero asymptotically (the limit $\varepsilon \to 0$ filters out all type-II error decay rates which would only be possible with type-I error that does not go to zero asymptotically). The quantum Stein's lemma states that this limit is given by the relative entropy between ρ and σ, i.e. (Hiai and Petz 1991)

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty} \frac{1}{n} D_H^\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) = D(\rho\|\sigma)\,. \tag{2.4.14}$$

In fact, this equality is even true if one drops the limit $\varepsilon \to 0$, i.e. for all $\varepsilon \in (0, 1)$:

$$\lim_{n\to\infty} \frac{1}{n} D_H^\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) = D(\rho\|\sigma)\,. \tag{2.4.15}$$

which is known as the strong-converse property (Ogawa and Nagaoka 2000). This means that even if one was to allow a very high type-I error probability (just not exactly 1), then the decay rate of the optimal type-II

error still cannot exceed the relative entropy asymptotically. Note that this is an asymptotic statement, for finite n the type-II error will of course be smaller if one allows for a higher type-I error.

The quantum Stein's Lemma forms the basis of all of asymmetric quantum hypothesis testing, and also gives the quantum relative entropy an operational meaning.

2.5 The (Smoothed) Max-Divergence

Yet another quantum divergence is the so-called quantum max-divergence. It is the quantum analogue of the classical maximum log-likelihood ratio, and defined as follows (Datta 2009):

$$D_{\max}(\rho\|\sigma) \coloneqq \log \inf \{ \lambda \in \mathbb{R} \mid \rho \leq \lambda\sigma \} . \tag{2.5.1}$$

The quantum max-divergence satisfies the data-processing inequality (Datta 2009) and is the unique quantum generalization of the classical $\alpha = \infty$ Rényi divergence. Precisely, if we define for P, Q classical probability distributions

$$D_\infty(P\|Q) \coloneqq \lim_{\alpha\to\infty} D_\alpha(P\|Q) \tag{2.5.2}$$

then $D_{\max}$ reduces to D_∞ when evaluated on classical states, and it also is the only such quantum divergence that also satisfies the data-processing inequality (see Tomamichel (2016) for a proof of this statement). In particular, this implies that for any α-Rényi divergence $\mathbf{D}_\alpha$ and any two quantum states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ we have

$$\mathbf{D}_\alpha(\rho\|\sigma) \leq \widehat{D}_\alpha^{\max}(\rho\|\sigma) \leq \lim_{\alpha\to\infty} \widehat{D}_\alpha^{\max}(\rho\|\sigma) = D_{\max}(\rho\|\sigma) \tag{2.5.3}$$

where we used the maximal α-Rényi divergence from Section 2.3.3 and the fact that it is monotonically increasing in α (which follows from the fact that the classical Rényi divergence is monotonically increasing in α).

Max Channel Divergence

Specifically relevant later will also be that the the max channel divergence has the following simple representation, where the optimal input state is always given by a maximally entangled state (for a proof see Díaz et al. (2018), Definition 19 and Wilde et al. (2020), Lemma 12):

$$D_{\max}(\mathcal{E}\|\mathcal{F}) \coloneqq \sup_{\psi_{RA}} D_{\max}\big((\mathrm{id}_R \otimes \mathcal{E})(\psi)\|(\mathrm{id}_R \otimes \mathcal{F})(\psi)\big) \tag{2.5.4}$$

$$= D_{\max}\big((\mathrm{id}_R \otimes \mathcal{E})(\Phi)\|(\mathrm{id}_R \otimes \mathcal{F})(\Phi)\big), \tag{2.5.5}$$

where $R \simeq A$ and $\Phi = \Phi_{RA}$ is a maximally entangled state. Hence, the max channel divergence is easily computable as a semi-definite program (SDP).

2.5.1 Smoothing and Different Smoothing Conventions

The max-divergence can often be quite large, since it has a fairly strict requirement: ρ has to be dominated by $\lambda\sigma$ everywhere. One can try to relax this to only requiring "almost everywhere", which is usually done by smoothing the state ρ (Renner 2005, Tomamichel 2012), i.e. considering $D_{\max}(\tilde{\rho}\|\sigma)$, where $\tilde{\rho} \in B_\varepsilon(\rho)$ lies within an ε-ball of ρ (in a certain distance, see below). There are different conventions employed throughout the literature, regarding which distance is chosen and what exactly the ball is, with the differences usually being whether one includes only normalized or also sub-normalized states, and whether one uses the trace distance or purified distance as a metric. Throughout this thesis we will be using the purified distance as a metric, and smooth only over normalized states. For normalized states, we can use the purified distance as defined in with the purified distance as defined in (2.1.13), and we will then write the ε-ball in purified distance of normalized states around ρ as $B_\varepsilon(\rho)$ (in some other works this is also written as $B_\varepsilon^{P,\circ}(\rho)$, or $B_\varepsilon^\circ(\rho)$) , and then define the smoothed max-divergence between a normalized $\rho \in \mathcal{D}(\mathcal{H})$, and $\sigma \in \mathscr{P}(\mathcal{H})$ as:

$$D_{\max}^\varepsilon(\rho\|\sigma) := \inf_{\tilde{\rho}\in B_\varepsilon(\rho)} D_{\max}(\tilde{\rho}\|\sigma)\,. \tag{2.5.6}$$

Relation to the Hypothesis Testing Relative Entropy

The smoothed max divergence turns out to be very closely related to the hypothesis testing relative entropy, in particular they are in some sense "asymptotically equivalent". This is made precise by the following lemma:

Lemma 2.5.1 (Anshu et al. 2019, Theorem 4)**.** *Let $\rho \in \mathcal{D}(\mathcal{H})$ and $\sigma \in \mathscr{P}(\mathcal{H})$. Then, for $\varepsilon \in (0,1)$ and $\delta \in (0, 1-\varepsilon^2)$:*

$$D_H^{1-\varepsilon^2-\delta}(\rho\|\sigma) - \log\left(\frac{4(1-\varepsilon^2)}{\delta^2}\right) \leq D_{\max}^\varepsilon(\rho\|\sigma) \leq D_H^{1-\varepsilon^2}(\rho\|\sigma) - \log\left(1-\varepsilon^2\right). \tag{2.5.7}$$

Note that Anshu et al. (2019) defines the smoothed max-divergence using sub-normalized smoothing; however, the proof of this lemma goes through also when restricting to normalized states. Note also that the arXiv version (v1) of Anshu et al. (2019) has a typo in the admissible range of δ, which has been corrected in the published version. For similar bounds, although with a slightly different structure and thus not directly comparable to the ones above, also see Datta et al. (2013).

2.5.2 The Quantum Asymptotic Equipartition Property (AEP)

In classical information theory, it is known as the asymptotic equipartition property (see e.g. Cover and Thomas (2012)), that the only sample strings from an i.i.d. distribution that "happen" (with high probability) are those where the average log probability of the individual sample outcomes is close to the entropy of the distribution, and of those strings that "happen", all "happen roughly equally likely". Analogously, for two classical probability distributions P and Q, the average log-likelihood ratio $\log\left(\frac{P(x)}{Q(x)}\right)$ of n samples drawn from P will concentrate around $D(P\|Q)$ (i.e. values outside of a small region around this will occur with vanishing probability as $n \to \infty$), and all values in a region around $D(P\|Q)$ will occur "roughly equally often". This means that asymptotically for large n, for all sample strings x_n but some with total mass (under

P^n) of at most ε, $\frac{1}{n}\log\left(\frac{P^n(x_n)}{Q^n(x_n)}\right) \approx D(P\|Q)$ and hence there exists a $\tilde{P}_n \approx^\varepsilon P^n$ (which only deviates from P^n on those x_n which have mass at most ε) such that $\frac{1}{n}\log\left(\frac{\tilde{P}_n(x_n)}{Q^n(x_n)}\right) \approx D(P\|Q)$ now for all x_n. Thus, then $\frac{1}{n}D^\varepsilon_{\max}(P^n\|Q^n) \approx D(P\|Q)$, and one can show that indeed

$$\lim_{n\to\infty} D^\varepsilon_{\max}(P^n\|Q^n) = D(P\|Q) \tag{2.5.8}$$

(this explanation was aimed at bridging the gap between the classical origin of the term AEP and the quantum statements below, but was of course in no way rigorous; for a rigorous treatment of the classical asymptotic equipartition property we refer to e.g. Cover and Thomas (2012)). It turns out that this last statement has a quantum analogue, namely the asymptotic equipartition property for the quantum smoothed max-relative entropy, which states that (Tomamichel, Colbeck, and Renner 2009, Tomamichel 2012)

$$\lim_{n\to\infty} \frac{1}{n} D^\varepsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) = D(\rho\|\sigma) \quad \forall \varepsilon \in (0,1)\,. \tag{2.5.9}$$

2.5.3 Non-Asymptotic Bounds for the Smoothed Max Divergence

For the proof of one of our main theorems later (Theorem 3.3.4), we will require a bound on the speed of this convergence. Results in this direction have appeared before in the literature; however, none of the previous results turn out to be directly applicable for our purposes (see Remark 2.5.3 and Remark 2.5.4 below). To prove our result we follow the established path of using quantum Rényi divergences. One key ingredient is the previously established bound on the distance between the Petz–Rényi and the Umegaki relative entropy (Lemma 2.3.1), using which we can establish a bound on the convergence speed in the asymptotic equipartition property. Since we want to apply this to obtain n-shot bounds, where constants are (unfortunately) not completely irrelevant, this Lemma also goes to a bit of length to optimize the appearing constants in the argument.

Lemma 2.5.2 (AEP Convergence Bound). *Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be quantum states, and for $\gamma \in (0,1]$, take $c_\gamma(\rho\|\sigma)$ as in (2.3.5). Then, for all $\varepsilon \in (0,1)$ and $n \in \mathbb{N}$:*

$$\frac{1}{n} D^\varepsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) - \frac{c_\gamma(\rho\|\sigma)}{\sqrt{n}}\left[\frac{\ln(2)\log(3)}{2}\cosh\left(\frac{\log(3)}{2}\right) + \frac{4}{\log(3)}\log\left(\frac{1}{1-\varepsilon}\right)\right], \tag{2.5.10}$$

$$\frac{1}{n} D^\varepsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq D(\rho\|\sigma) + \frac{c_\gamma(\rho\|\sigma)}{\sqrt{n}}\left[\frac{\ln(2)\log(3)}{2} + \frac{4}{\log(3)}\log\left(\frac{1}{\varepsilon}\right)\right] + \frac{1}{n}\log\left(\frac{1}{1-\varepsilon^2}\right), \tag{2.5.11}$$

which implies

$$\frac{1}{n} D^\varepsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) - \frac{4c_\gamma(\rho\|\sigma)}{\sqrt{n}}\log\left(\frac{2}{1-\varepsilon}\right), \tag{2.5.12}$$

$$\frac{1}{n} D^\varepsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq D(\rho\|\sigma) + \frac{4c_\gamma(\rho\|\sigma)}{\sqrt{n}}\log\left(\frac{2}{\varepsilon}\right) + \frac{1}{n}\log\left(\frac{1}{1-\varepsilon^2}\right). \tag{2.5.13}$$

These bounds can be tightened by adding a condition on n to be large enough. Define first the numerical constants

$$K_1 := 2\sqrt{2\ln(2)\cosh\left(\frac{\log(3)}{2}\right)} \leq 2.72, \tag{2.5.14}$$

$$K_2 := 2\sqrt{2\ln(2)} \leq 2.36\,. \tag{2.5.15}$$

Then, if

$$n \geq \log\left(\frac{1}{1-\varepsilon}\right)\left(\frac{8}{\log(3)K_1}\right)^2, \tag{2.5.16}$$

we have the stronger bound

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) - \frac{K_1}{\sqrt{n}}c_\gamma(\rho\|\sigma)\sqrt{\log\left(\frac{1}{1-\varepsilon}\right)}, \tag{2.5.17}$$

and similarly, if

$$n \geq \log\left(\frac{1}{\varepsilon}\right)\left(\frac{8}{\gamma c_\gamma(\rho\|\sigma)K_2}\right)^2, \tag{2.5.18}$$

it holds that

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq D(\rho\|\sigma) + \frac{K_2}{\sqrt{n}}c_\gamma(\rho\|\sigma)\sqrt{\log\left(\frac{1}{\varepsilon}\right)} + \frac{1}{n}\log\left(\frac{1}{1-\varepsilon^2}\right). \tag{2.5.19}$$

Remember that $\gamma c_\gamma(\rho\|\sigma) \geq \log(3)$, and hence (2.5.18) is also always satisfied if

$$n \geq \log\left(\frac{1}{\varepsilon}\right)\left(\frac{8}{\log(3)K_2}\right)^2. \tag{2.5.20}$$

Proof. We start with the proof of (2.5.17). We use X. Wang and Wilde (2019a), Prop. 4, which states that for all $\alpha \in [0,1)$ and all $\varepsilon \in [0,1)$

$$D^{\varepsilon}_{\max}(\rho\|\sigma) \geq D_\alpha(\rho\|\sigma) + \frac{2}{\alpha-1}\log\left(\frac{1}{1-\varepsilon}\right). \tag{2.5.21}$$

Note that the statement in X. Wang and Wilde (2019a) is given for smoothing in trace-distance, but since the trace distance is always less than the purified distance, for fixed ε the trace-distance-smoothed max-divergence is smoothed over a larger ball, and hence also always smaller, which is why the statement for the purified-distance-smoothed max-divergence is implied. We apply this to $\rho^{\otimes n}$ and $\sigma^{\otimes n}$ and use the additivity of the Petz-Rényi relative entropy to get:

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D_\alpha(\rho\|\sigma) + \frac{1}{n}\frac{2}{\alpha-1}\log\left(\frac{1}{1-\varepsilon}\right). \tag{2.5.22}$$

Now, we combine this with (2.3.9) of Lemma 2.3.1 to get

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) + \ln(2)\cosh\left(\frac{\log(3)}{2}\right)c_\gamma(\rho\|\sigma)^2(\alpha-1) + \frac{1}{n}\frac{2}{\alpha-1}\log\left(\frac{1}{1-\varepsilon}\right), \quad (2.5.23)$$

together with the condition $0 \leq 1-\alpha \leq \frac{\log 3}{2c_\gamma(\rho\|\sigma)}$. We are now free to choose α to optimize the right-hand side. It is easy to see that the right-hand side will be maximal if both terms are equal, which is achieved for

$$1-\alpha = \sqrt{\frac{2\log\left(\frac{1}{1-\varepsilon}\right)}{n\ln(2)\cosh\left(\frac{\log(3)}{2}\right)c_\gamma(\rho\|\sigma)^2}} = \frac{4}{K_1 c_\gamma(\rho\|\sigma)}\sqrt{\frac{\log\left(\frac{1}{1-\varepsilon}\right)}{n}}. \quad (2.5.24)$$

Hence we get

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) - \frac{K_1}{\sqrt{n}}c_\gamma(\rho\|\sigma)\sqrt{\log\left(\frac{1}{1-\varepsilon}\right)}, \quad (2.5.25)$$

together with the condition

$$n \geq \log\left(\frac{1}{1-\varepsilon}\right)\left(\frac{4}{K_1\log(3)}\right)^2. \quad (2.5.26)$$

In order to get an expression without a condition on n we have to choose a different α. Choosing

$$1-\alpha = \frac{\log(3)}{2c_\gamma(\rho\|\sigma)\sqrt{n}} \quad (2.5.27)$$

satisfies the condition on $1-\alpha$ for all n and gives:

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) - \frac{c_\gamma(\rho\|\sigma)}{\sqrt{n}}\left[\frac{\ln(2)\log(3)}{2}\cosh\left(\frac{\log(3)}{2}\right) + \frac{4}{\log(3)}\log\left(\frac{1}{1-\varepsilon}\right)\right]. \quad (2.5.28)$$

The simplification (2.5.12) follows from $\log(3) > 1$ and

$$\frac{\ln(2)\log(3)}{2}\cosh\left(\frac{\log(3)}{2}\right) < 1\,. \quad (2.5.29)$$

The second equation (2.5.19) is proved very analogously. We start again with a relation to a Rényi relative entropy for $\alpha > 1$ (Khatri and Wilde 2020, Prop. 4.61):

$$D^{\varepsilon}_{\max}(\rho\|\sigma) \leq D_\alpha(\rho\|\sigma) + \frac{2}{\alpha-1}\log\left(\frac{1}{\varepsilon}\right) + \log\left(\frac{1}{1-\varepsilon^2}\right). \quad (2.5.30)$$

Note that Khatri and Wilde (2020), Prop. 4.61 uses the sandwiched Rényi divergence, which however is always less than the Petz, and hence the above statement is implied. We again apply this to n tensor powers of ρ and σ and combine it with (2.3.6) from Lemma 2.3.1 to obtain:

$$\frac{1}{n}D^{\varepsilon}_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq D(\rho\|\sigma) + \ln(2)c_\gamma(\rho\|\sigma)^2(\alpha-1) + \frac{2}{n(\alpha-1)}\log\left(\frac{1}{\varepsilon}\right) + \frac{1}{n}\log\left(\frac{1}{1-\varepsilon^2}\right), \quad (2.5.31)$$

together with the condition

$$\alpha - 1 \leq \frac{\gamma}{2}\,. \tag{2.5.32}$$

The right hand side is minimized again for

$$\alpha - 1 = \sqrt{\frac{2\log\left(\frac{1}{\varepsilon}\right)}{n\ln(2)c_\gamma(\rho\|\sigma)^2}} = \frac{4}{K_2 c_\gamma(\rho\|\sigma)}\sqrt{\frac{\log\left(\frac{1}{\varepsilon}\right)}{n}}, \tag{2.5.33}$$

which gives

$$\frac{1}{n}D^\varepsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq D(\rho\|\sigma) + \frac{K_2}{\sqrt{n}}c_\gamma(\rho\|\sigma)\sqrt{\log\left(\frac{1}{\varepsilon}\right)} + \frac{1}{n}\log\left(\frac{1}{1-\varepsilon^2}\right), \tag{2.5.34}$$

as well as the condition

$$n \geq \log\left(\frac{1}{\varepsilon}\right)\left(\frac{8}{\gamma c_\gamma(\rho\|\sigma)K_2}\right)^2. \tag{2.5.35}$$

Alternatively, choosing

$$\alpha - 1 = \frac{\log(3)}{2c_\gamma(\rho\|\sigma)\sqrt{n}} \tag{2.5.36}$$

will again give (2.5.11), and the simplified form (2.5.13) follows by the same argument that $\log(3) > 1$ and

$$\frac{\ln(2)\log(3)}{2} < 1\,. \tag{2.5.37}$$

This concludes the proof. □

Remark 2.5.3. Our Lemma 2.5.2 can be seen an extension of Tomamichel (2012), Theorem 6.4, where a similar bound to (2.5.19) is shown (however with a worse constant and different smoothing convention). Note that Tomamichel (2012) uses the notation $S(\rho\|\sigma) \coloneqq -D(\rho\|\sigma)$ and $S^\varepsilon_{\min}(\rho\|\sigma) \coloneqq -D^\varepsilon_{\max}(\rho\|\sigma)$, and hence Tomamichel (2012), Theorem 6.4, while looking like a lower bound, actually is an upper bound on $D^\varepsilon_{\max}$ (although slightly worse than our upper bound (2.5.19)). An equivalent of (2.5.17) is shown in Tomamichel (2012) only for the smoothed conditional min-entropy and not for the (more general) smoothed max-relative entropy.

Remark 2.5.4. Another already existing bound for the AEP convergence is the so-called second-order expansion (Tomamichel and Hayashi 2013, K. Li 2014), which gives a tight asymptotic characterization also of the second-order $\sqrt{n}$ term in the convergence to the relative entropy. There, the second-order term is shown to be proportional to the square root of the relative entropy variance

$$V(\rho\|\sigma) \coloneqq \mathrm{Tr}\big(\rho(\log(\rho) - \log(\sigma) - D(\rho\|\sigma))^2\big)\,. \tag{2.5.38}$$

Later in the proof of Theorem 3.3.4, we want to apply these convergence bounds to a case in which $\rho = \rho_n$ and $\sigma = \sigma_n$ are the states in an adaptive discrimination protocol, and we would like to obtain a bound on the convergence parameter in n. We will see that by using chain rules for the geometric relative Rényi entropy

(or the max-relative entropy) we will be able to show that $c_\gamma(\rho_n\|\sigma_n) = \mathcal{O}(n)$, while we are not aware of any way of obtaining such a bound for $V(\rho_n\|\sigma_n)$, and hence cannot directly use second-order asymptotics to bound the AEP convergence in this application.

2.6 Symmetries and Permutation Invariance

We say that a state $\rho_n \in \mathcal{D}(\mathcal{H}^{\otimes n})$ is *permutation invariant*, if for any permutation of n systems $\pi \in \mathfrak{S}(n)$ and associated unitary operator $P_{\mathcal{H}}(\pi)$ it holds that $P_{\mathcal{H}}(\pi)\rho_n P_{\mathcal{H}}(\pi)^\dagger = \rho_n$. We say that a channel $\mathcal{E}_n \in \mathrm{CPTP}(A^n \to B^n)$ is *permutation covariant* if $\mathcal{E}_n(P_A(\pi)\rho_n P_A(\pi)^\dagger) = P_B(\pi)\mathcal{E}_n(\rho_n)P_B(\pi)^\dagger$ for all input states $\rho_n \in \mathcal{D}(A^n)$ and all permutations $\pi \in \mathfrak{S}_n$. We say that a set of channels $\mathcal{S}_n \subset \mathrm{CPTP}(A^n \to B^n)$ is *closed under permutations*, if for any $\mathcal{E}_n \in \mathcal{S}_n$ and any permutation $\pi \in \mathfrak{S}_n$, also the channel with permuted input and output systems $\rho \mapsto P_B(\pi)\mathcal{E}_n\big(P_A(\pi)\rho P_A(\pi)^\dagger\big)P_B(\pi)^\dagger$ is an element of $\mathcal{S}_n$.

The simplest examples of permutation invariant states are just tensor power states, although the set of permutation invariant states is significantly bigger than that. Similarly, the simplest examples of permutation covariant channels are channels which are tensor powers, i.e. $\mathcal{E}_n = \mathcal{E}^{\otimes n}$, although again the set of permutation covariant channels is significantly bigger than that. The main reason for permutation invariance being important in this thesis is the following Lemma, which establishes that if a sequence of channels is permutation covariant (e.g. because it is an i.i.d. string of channels), then also the optimizing input state will be permutation invariant.

Lemma 2.6.1. *Let $\mathcal{E}_n, \mathcal{F}_n \in \mathrm{CPTP}(A^n \to B^n)$ both be permutation covariant and let $\mathbf{D}$ be D or D_M. Then,*

$$\sup_{\substack{\nu \in \mathcal{D}\left(R\otimes A^{\otimes n}\right)\\ R \text{ arbitrary}}} \mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) = \max_{\substack{\nu \in \mathcal{D}\left(R^{\otimes n}\otimes A^{\otimes n}\right)\\ R\cong A\\ \nu \text{ permut. invariant}}} \mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) \tag{2.6.1}$$

i.e. the supremum is achieved for a permutation invariant state $\nu_{R^nA^n} = \nu_{(RA)^n}$ where R is isomorphic to A. Note that we mean permutation invariant with respect to permutations permuting the n copies of (RA).

Proof. This can be seen as a special case of Leditzky et al. (2018), Proposition II.4, although for this special case we provide a slightly simpler proof. We use the above introduced notation for permutations and associated unitary operators, for $\pi \in \mathfrak{S}_n$ the action of e.g. $P_A(\pi)$ will only permute the n copies of the system A and ignore any additional (reference) systems. Let $\nu = \nu_{R_0A^n} \in \mathcal{D}(R_0 \otimes A^{\otimes n})$ be an arbitrary state, where R_0 is an arbitrary reference system, and let $\pi \in \mathfrak{S}_n$. Then by unitary invariance of $\mathbf{D}$ (which follows from the data-processing inequality):

$$\begin{aligned}\mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) &= \mathbf{D}(P_B(\pi)\mathcal{E}_n(\nu)P_B(\pi)^\dagger\|P_B(\pi)\mathcal{F}_n(\nu)P_B(\pi)^\dagger) &(2.6.2)\\ &= \mathbf{D}\Big(\mathcal{E}_n\Big(P_A(\pi)\nu P_A(\pi)^\dagger\Big)\ \Big\|\ \mathcal{F}_n\Big(P_A(\pi)\nu P_A(\pi)^\dagger\Big)\Big)\,. &(2.6.3)\end{aligned}$$

Define

$$\omega_{PR_0A^n} = \frac{1}{n!}\sum_{\pi\in\mathfrak{S}_n} |\pi\rangle\langle\pi| \otimes P_A(\pi)\nu_{R_0A^n}P_A(\pi)^\dagger\,, \tag{2.6.4}$$

where the first register is classical and stores the permutation π. By the direct sum property we have

$$\mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) = \frac{1}{n!}\sum_{\pi\in\mathfrak{S}_n}\mathbf{D}\Big(\mathcal{E}_n\Big(P_A(\pi)\nu P_A(\pi)^\dagger\Big)\,\Big\|\,\mathcal{F}_n\Big(P_A(\pi)\nu P_A(\pi)^\dagger\Big)\Big) \tag{2.6.5}$$

$$= \mathbf{D}(\mathcal{E}_n(\omega_{PR_0A^n})\|\mathcal{F}_n(\omega_{PR_0A^n}))\,. \tag{2.6.6}$$

Let $\omega_{SPR_0A^n}$ be a purification of $\omega_{PR_0A^n}$. Note that

$$\omega_{A^n} = \frac{1}{n!}\sum_{\pi\in\mathfrak{S}_n} P_A(\pi)\nu_{A^n}P_A(\pi)^\dagger \tag{2.6.7}$$

is permutation invariant, and hence by Christandl et al. (2007), Lemma II.5, there exists a system K isomorphic to A, and a permutation invariant purification $\omega_{(KA)^n} \in \mathcal{D}\left(K^{\otimes n}\otimes A^{\otimes n}\right)$ where the permutations act on $\omega_{(KA)^n}$ by permuting the copies of (KA). Now the two purifications $\omega_{(KA)^n}$ and $\omega_{SPR_0A^n}$ will be related by a partial isometry $V : K^n \to SPR_0$ which commutes with $\mathcal{E}_n$ and $\mathcal{F}_n$ (since they only act on A^n). Hence,

$$\mathbf{D}(\mathcal{E}_n(\omega_{PR_0A^n})\|\mathcal{F}_n(\omega_{PR_0A^n})) \le \mathbf{D}(\mathcal{E}_n(\omega_{SPR_0A^n})\|\mathcal{F}_n(\omega_{SPR_0A^n})) \tag{2.6.8}$$

$$= \mathbf{D}(\mathcal{E}_n(\omega_{(KA)^n})\|\mathcal{F}_n(\omega_{(KA)^n})) \tag{2.6.9}$$

by the data processing inequality and isometric invariance. The fact that the supremum is also achieved follows from the same argument as in Lemma 2.2.5. □

The significance of these restrictions to permutation covariant channels and permutation invariant input states comes from the fact that both terms $\mathcal{E}_n(\nu_n)$ and $\mathcal{F}_n(\nu_n)$ will then be permutation invariant, and this allows us to use:

Lemma 2.6.2 (Berta, Brandao, and Hirche 2021, Lem. 2.4)**.** *Let $\rho_n, \sigma_n \in \mathcal{D}\left(\mathcal{H}^{\otimes n}\right)$ with σ_n permutation invariant. Then,*

$$D(\rho_n\|\sigma_n) - \log \mathrm{poly}(n) \le D_M(\rho_n\|\sigma_n) \le D(\rho_n\|\sigma_n)\,, \tag{2.6.10}$$

Additionally, we have the following lemma (essentially a variant of Berta, Brandao, and Hirche (2021), Lemma 2.5) which allows us to remove a convex hull in the infimum over the first argument, if the states also all lie in a sufficiently small linear subspace. This is for example the case if all states are permutation invariant, as the subspace of permutation invariant density matrices in $\mathcal{D}\left(\mathcal{H}^{\otimes n}\right)$ lies in a linear subspace of $\mathcal{B}\left(\mathcal{H}^{\otimes n}\right)$ whith dimension upper bounded by $(n+1)^{(\dim\mathcal{H})^2}$.

Lemma 2.6.3. *Let $\sigma \in \mathcal{D}(\mathcal{H})$, and let $S \subset \mathcal{W}\cap\mathcal{D}(\mathcal{H})$ be a set of density matrices, where $\mathcal{W}$ is a linear subspace of $\mathcal{B}(\mathcal{H})$. Then,*

$$\inf_{\rho\in\mathcal{C}(S)} D(\rho\|\sigma) \ge \inf_{\rho\in S} D(\rho\|\sigma) - \log(\dim\mathcal{W}+1)\,, \tag{2.6.11}$$

where $\mathcal{C}(S)$ is the convex hull of S.

Proof. By Caratheodory's theorem, we can write any element $\tilde{\rho} \in \mathcal{C}(S)$ as $\sum_{i=1}^{n} p_i \tilde{\rho}_i$, where p_i is a probability distribution, $\tilde{\rho}_i \in S$ and $n = \dim \mathcal{W} + 1$. We can assume $\tilde{\rho}_i \ll \sigma$ for all i, as otherwise $D(\tilde{\rho}\|\sigma) = \infty$ and there is nothing to show. For $\varepsilon > 0$ fixed, let $\rho_i \coloneqq \tilde{\rho}_i + \varepsilon \Pi_\sigma$, where Π_σ is the projection onto the support of σ. Then,

$$
\begin{aligned}
D\left(\sum_i p_i \rho_i \,\middle\|\, \sigma\right) &= \operatorname{Tr}\left[\left(\sum_i p_i \rho_i\right)\left(\log\left(\sum_j p_j \rho_j\right) - \log(\sigma)\right)\right] && (2.6.12)\\
&= \operatorname{Tr}\left[\sum_i p_i \rho_i \left(\log\left(\sum_j p_j \rho_j\right) - \log(\sigma)\right)\right] && (2.6.13)\\
&\geq \operatorname{Tr}\left[\sum_i p_i \rho_i (\log(p_i \rho_i) - \log(\sigma))\right] && (2.6.14)\\
&= \sum_i p_i \operatorname{Tr}[\rho_i (\log \rho_i - \log \sigma + \log p_i)] && (2.6.15)\\
&= \sum_i p_i D(\rho_i \| \sigma) - H(p) \geq \sum_i p_i D(\rho_i \| \sigma) - \log(n)\,, && (2.6.16)
\end{aligned}
$$

where for the first inequality we used the operator monotonicity of the logarithm and that $\sum_j p_j \rho_j \geq p_i \rho_i$ for every i. Now, by the already mentioned continuity of D in the first variable (when restricted to density matrices on the support of σ), we can take the limit $\varepsilon \to 0$ on both sides to get

$$
D(\tilde{\rho} \parallel \sigma) \geq \sum_i p_i D(\tilde{\rho}_i \| \sigma) - \log(n) \geq \inf_{\rho \in S} D(\rho \| \sigma) - \log(n)\,. \tag{2.6.17}
$$

□

2.7 Semi-definite Programs (SDPs)

This section aims to very briefly introduce the fundamental aspects of semi-definite programs used in this thesis. For a more exhaustive treatment of semi-definite programs (in general and in quantum information theory), see e.g. Watrous (2018), Vandenberghe and Boyd (1996), De Klerk (2002), and Anjos and Lasserre (2012). Throughout this section we write the Hilbert-Schmidt inner product using $\langle \cdot, \cdot \rangle$, i.e., $\langle A, B \rangle \coloneqq \operatorname{Tr}\left(A^\dagger B\right)$.

A semi-definite program (SDP) is a convex optimization program, where one optimizes a linear objective function over an affine subspace of the convex cone of positive semi-definite matrices.

Definition 2.7.1 (Primal SDP)**.** For $\mathbb{K} = \mathbb{R}$ or $\mathbb{C}$, $C \in \mathcal{B}\left(\mathbb{K}^n\right)$ and $A_k \in \mathcal{B}\left(\mathbb{K}^n\right)$, $b_k \in \mathbb{K}$ for $k = 1, ..., m$, we call an optimization problem of the form

$$
\begin{aligned}
\underset{X \in \mathcal{B}\left(\mathbb{K}^n\right)}{\text{minimize}} \quad & \langle C, X \rangle \\
\text{subject to} \quad & \langle A_k, X \rangle = b_k \quad \forall k = 1, ..., m, \\
& X \geq 0
\end{aligned} \tag{2.7.1}
$$

a semi-definite program (SDP) in primal standard form in dimension n with m constraints.

Note that it is pure convention to write the primal optimization problem as a minimization, one can replace $C \mapsto -C$ to turn the minimization into a maximization.

We say that the SDP is feasible, if there exists an X that satisfies the constraints of the problem. We also write $\mathcal{S} \subset \mathscr{P}(\mathbb{K}^n)$ for the set of feasible X.

Duality

To a primal SDP as in (2.7.1) one can associate a dual optimization problem

$$\begin{aligned} \underset{S \in \mathcal{B}(\mathbb{K}^n), \underline{y} \in \mathbb{K}^m}{\text{maximize}} \quad & \sum_{k=1}^{m} b_k^* y_k \\ \text{subject to} \quad & \sum_{k=1}^{m} y_k A_k + S = C, \\ & S \geq 0 \end{aligned} \tag{2.7.2}$$

It is known as *weak duality* that the optimal value of the dual problem is always at most as large as the optimal value of the primal problem. Under some additional conditions also *strong duality* holds, which states that the optimum value of the primal and dual problem are equal. One such condition under which strong duality holds (known as Slater's Condition) is that there exists a strictly feasible solution for one of the two problems, i.e. an $X > 0$ that satisfies the constraints of the primal problem, or an $S > 0$ that satisfies the constraints of the dual problem.

Complex and Real SDPs

In this thesis we will want to treat complex SDPs, while the SDPs in the computer science literature are usually considered to be real. Fortunately, there is a fairly straightforward way to map between them:

Lemma 2.7.2 (J. Wang 2024)**.** *A primal complex SDP in dimension n with m constraints is equivalent to a primal real SDP in dimension $\tilde{n} = 2n$ with $\tilde{m} = 2m$ constraints.*

Theoretical Complexity

SDPs are generally considered efficiently solvable, given that one can ensure that a feasible solution exists. The following lemma states some rigorous results regarding the complexity of solving SDPs. Note that (being a result from theoretical computer science) the following lemma assumes the bit-model of computation, and hence that all coefficients in the problem can be represented as integers. If one starts with an SDP where C, A_k and b_k are specified by rational numbers, one can rescale to obtain such an integer SDP.

Lemma 2.7.3 (see e.g. De Klerk (2002), Section 1.9 and references therein, or also e.g. Jiang et al. (2020) for more recent developments and algorithms)**.** *Consider a real SDP in primal standard form (2.7.1) in dimension n with m constraints, where C and the A_k and b_k contain only integers. Given an $R > 0$ and the knowledge*

that either the problem is infeasible, or there exists a feasible X such that $\|X\|_\infty \leq R$, then for every (rational) $\varepsilon > 0$ there exists an algorithm with runtime polynomial in $n, m, \log(R), \log(1/\varepsilon)$, and the bit-length of the input, that either:

- *Returns a guarantee that the feasible set $\mathcal{S}$ does not contain a ball of radius ε (in $\|\cdot\|_\infty$)*
- *Returns an X that is at most ε away from the feasible set $\mathcal{S}$ (in $\|\cdot\|_\infty$), and that satisfies $|\langle C, X\rangle - p^*| \leq \varepsilon$, where p^* is the optimal value of the problem.*

When dealing with SDPs in this thesis, we will generally disregard the aspects of representing numbers through bits and issues of numerical accuracy, with the essence of Lemma 2.7.3 then being that given some representation of an SDP on a computer, there exists an algorithm with at most $\mathcal{O}(\text{poly}(n, m, \log(R), \log(1/\varepsilon)))$ fundamental steps (i.e. addition/multiplication etc.) that approximates a solution up to ε (or certifies that no such solution exists).

3 Simple Binary Quantum Channel Discrimination

In Section 2.4 we introduced simply binary quantum hypothesis testing, where the task was to distinguish two different quantum states (i.e. density matrices). The task of binary quantum channel discrimination is similar, but now instead of being given a density matrix, we are given quantum channels, so the task is the following: Given an unknown quantum channel as a black box and the side information that it is one of two possible channels, determine the channel's identity (Chiribella, D'Ariano, and Perinotti 2008, Duan, Feng, and Ying 2009, Hayashi 2009, Harrow et al. 2010). The additional complexity here comes from the fact that, on top of finding the best measurement to perform on the output of the channel, we also have to figure out which quantum states to send as input to the channel.

Even more so, when we are given access to more than one copy of the channel (say we are given n identical black boxes, each of which can be used once), then there are different strategies (sometimes also called protocols, we will use these two terms interchangeably) in which we could set up our decision experiment – the so-called *parallel* and *adaptive* strategies. In a parallel strategy one prepares a joint state, usually entangled between the input systems of all the n copies of the channel and an additional reference (or memory) system. This state is then fed as input to all the n channels at once (with the state of the reference system being left undisturbed). In an adaptive strategy, on the other hand, one prepares a state of the input system of a single copy of the channel (again usually entangled with a reference system) which is fed into the first copy of the channel, with the state of the reference system being left undisturbed. The input to the next use of the channel is then chosen depending on the output of the first channel and the state of our reference system. This is done, most generally, by subjecting the latter to an arbitrary quantum operation (or channel), which we call a preparation operation. This step is repeated for each successive use of the channel until all the n black boxes have been used. See Figure 3.2 for a depiction of an adaptive strategy. Adaptive strategies are also sometimes called sequential, which however should not be confused with the setting of sequential hypothesis testing (Martínez Vargas et al. 2021, Y. Li, Tan, and Tomamichel 2022, Y. Li, Hirche, and Tomamichel 2022), where samples (i.e., states or channels) can be requested one by one.

One particularly interesting question is whether and to what degree adaptive strategies give an advantage over parallel ones. Note that any parallel strategy can be written as an adaptive strategy by taking all but one channel input as part of the reference system, and then choosing each preparation operation such that it extracts the next part of the joint input state for the next channel use and replaces it by the output of the previous channel use. However, the converse is not true, and so adaptive strategies are more general. Parallel strategies are conceptually a lot simpler than adaptive ones – aside from the measurement, everything is specified just by the joint input state – in contrast to adaptive strategies, in which after each channel use we can perform an arbitrary quantum operation to prepare the input to the next use of the channel. It is thus interesting to determine to what degree parallel strategies can still be optimal. It is known that in certain cases

adaptive strategies can give an advantage over parallel ones. In Harrow et al. (2010) the authors constructed an example in which an adaptive strategy with only two channel uses could be used to discriminate the channels with certainty, which is not possible with a parallel strategy, even if arbitrarily many channel uses are allowed.

Interestingly, *asymptotically*, there are multiple known cases in which adaptive strategies give no advantage over parallel ones, i.e., the optimal exponential decay rate of the error probability per channel use is the same in the asymptotic limit. For example, this is the case both in the symmetric and asymmetric settings when the channels are classical (Hayashi 2009) or classical-quantum (Wilde et al. 2020, Salek, Hayashi, and Winter 2022). For arbitrary quantum channels, the recently shown chain rule for the quantum relative entropy (Fang et al. 2020) and the characterization of asymmetric channel discrimination in terms of amortized relative entropy (Wilde et al. 2020, X. Wang and Wilde 2019b), also imply that in the asymmetric setting, in the asymptotic limit where we require the type I error to vanish, adaptive strategies give no advantage over parallel ones (i.e., the optimal asymptotic exponential decay rate of the type II error per channel use is the same for parallel and adaptive strategies; see Section 3.1.1 below for more details). This is in contrast to the symmetric setting in which the example of Salek, Hayashi, and Winter (2022) shows that there can be an advantage of adaptive strategies, also in the asymptotic limit.

In our main result of this section, we move from these purely asymptotic results to statements comparing adaptive and parallel strategies for finite n. Our main result (Corollary 3.3.1 and Theorem 3.3.4), relates the type I and type II errors of an arbitrary adaptive strategy with those of a suitably chosen parallel strategy. Specifically, given an adaptive strategy with n channel uses, we construct a parallel strategy with m channel uses (where m can be chosen at will), such that for arbitrary fixed type I errors of the two strategies, we find an explicit bound on the difference between their type II error decay rates. This difference goes to zero in a suitably chosen asymptotic limit $m, n \to \infty$ if also the type I error vanishes, hence also implying the known asymptotic equivalence in the asymmetric case. Our result answers the following interesting question which is of practical relevance: *Given an adaptive strategy involving n uses of the channel, if one instead employs a parallel strategy with m channel uses, how much worse are the errors going to be?*

Note that the asymptotic results obtained in Wilde et al. (2020) and X. Wang and Wilde (2019b) do not purely come from finite-length considerations, and hence results analogous to ours are not directly obtainable from these references.

Our main result becomes particularly interesting considering that we additionally show that one can optimize over all parallel strategies in time polynomial in the number of channel uses (Theorem 3.2.1), whereas optimizing over all adaptive strategies seems to require exponential time. Hence, we can also use our theorem to give a bound related to the following practically relevant question: *Given that one computed the optimal parallel strategy involving m uses of the channel, how much better could the errors of the best adaptive strategy with n channel uses possibly be?*

3.1 Parallel and Adaptive Strategies

In this section we will more formally introduce channel discrimination and the two different kinds of strategies we consider. Note that for any chosen channel discrimination strategy, at the end, once all black-box channels have been used, one will be left with some kind of quantum state which will (hopefully) depend on the identity of the black-box, and so the problem then reduces to the state-discrimination problem of discriminating these two possible quantum states. As discussed in Section 2.4, for this problem one will then perform a binary POVM measurement that optimizes errors in a certain way. As already mentioned, we will be exclusively focusing on the asymmetric setting in this thesis, and so the quantity of interest will be the hypothesis testing relative entropy (as defined in (2.4.10)), which includes the optimization over this POVM measurement. In particular, we will usually be interested in the hypothesis testing relative entropy divided by the number of channel uses n, as this is the analogue of the previously used division by the number of received samples, and it behaves appropriately as n becomes large. We will then specify in this section how different channel discrimination strategies lead to different expressions involving this hypothesis testing relative entropy.

Throughout this section, we will assume that we have to discriminate between the hypotheses, where we are given n copies of an unknown channel that is either $\mathcal{E}$ or $\mathcal{F}$, where $\mathcal{E}, \mathcal{F} : A \to B$ are two quantum channels, and so we can specify the hypotheses as

$$H_0 : \text{we are given } \mathcal{E}^{\otimes n}\,, \tag{3.1.1}$$

$$H_1 : \text{we are given } \mathcal{F}^{\otimes n}\,. \tag{3.1.2}$$

Parallel Strategies

As stated in the introductory section, the simple and naive approach to quantum channel discrimination is to use a parallel discrimination strategy, where the input state to all the channels is fixed at the beginning and does not (adaptively) depend on channel outputs. A parallel strategy is depicted in Figure 3.1, and it is specified by a joint input state $\nu_{RA^n} \in \mathcal{D}\left(RA^n\right)$. On the output side, we are then left with the state $\mathcal{E}^{\otimes n}(\nu_{RA^n})$ or $\mathcal{F}^{\otimes n}(\nu_{RA^n})$ (remember our convention to not write out identity channels on reference systems as explained in Section 2.1.2). The corresponding asymmetric error exponent per channel use (which as stated above includes the state-discrimination part of the problem), is then given by

$$\frac{1}{n} D_H^{\varepsilon}(\mathcal{E}^{\otimes n}(\nu_{RA^n}) \| \mathcal{F}^{\otimes n}(\nu_{RA^n}))\,, \tag{3.1.3}$$

and thus the optimal achievable asymmetric error exponent (per channel use) with a parallel strategy is given by:

$$e_P(n, \varepsilon, \mathcal{E}, \mathcal{F}) \coloneqq \frac{1}{n} D_H^{\varepsilon}(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n}) = \sup_{\nu \in \mathcal{D}(RA^n)} \frac{1}{n} D_H^{\varepsilon}(\mathcal{E}^{\otimes n}(\nu) \| \mathcal{F}^{\otimes n}(\nu))\,. \tag{3.1.4}$$

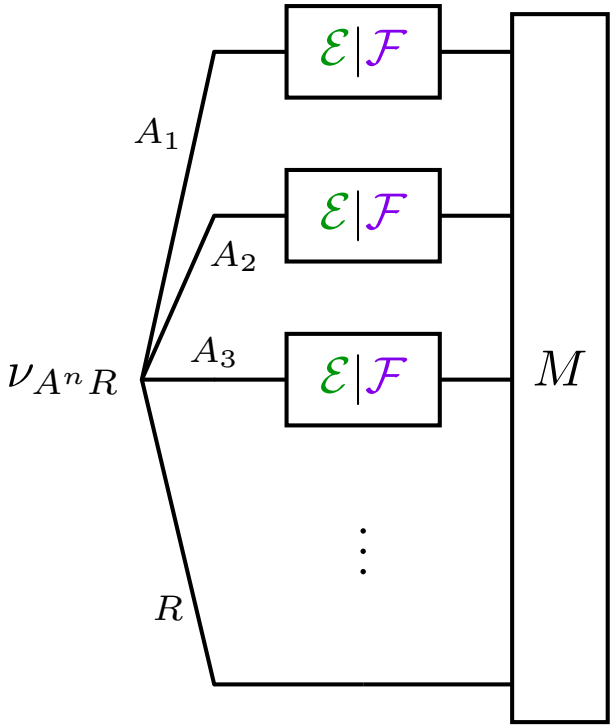

Figure 3.1: Illustration of a general parallel strategy with n uses of the black-box channel, and a joint binary POVM measurement $\{M, \mathbb{1} - M\}$ at the end.

Adaptive Strategies

The most general possible channel discrimination strategy (that uses the channels in a causal order, see Section 3.4.3 for a discussion on possibly more general strategies) will choose input states based on the outputs of previous channel uses. This is called an adaptive strategy (Cooney, Mosonyi, and Wilde 2016). A general adaptive channel discrimination strategy with n uses of the black-box channel ($\mathcal{E}$ or $\mathcal{F}$), can be fully specified by an initial state $\rho_1 = \sigma_1 \in \mathcal{D}(RA)$, a set of $n-1$ CPTP maps $\Lambda_i : RB \to RA$, that transform the state before it is fed into the next black-box channel, and a final binary POVM $\{M, \mathbb{1} - M\}$ on RB. We will assume the size of reference system R to be fixed and identical throughout the protocol (this is without loss

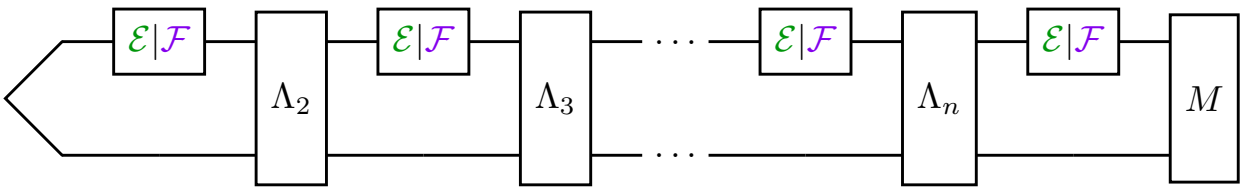

Figure 3.2: Illustration of a general adaptive protocol with n uses of the black-box channel. The top row makes use of the given black-box $\mathcal{E}|\mathcal{F}$, which is either $\mathcal{E}$ or $\mathcal{F}$, while the bottom row depicts the memory system R. At various stages in the protocol, the green states ρ occur if the channel is $\mathcal{E}$ and the purple states σ occur if the channel is $\mathcal{F}$.

of generality). The protocol consists of alternating applications of the black-box channel and the preparation CPTP maps Λ_i (see Figure 3.2). We define for $i \in \{2, \dots, n\}$:

$$\rho_i \coloneqq \Lambda_i(\mathcal{E}(\rho_{i-1})), \quad \sigma_i \coloneqq \Lambda_i(\mathcal{F}(\sigma_{i-1})), \tag{3.1.5}$$

and so the final state before the action of the POVM will be $\mathcal{E}(\rho_n)$ if the channel is $\mathcal{E}$ and $\mathcal{F}(\sigma_n)$ if the channel is $\mathcal{F}$. The distinguishability of the channels then comes down to the distinguishability of the two states $\mathcal{E}(\rho_n)$ and $\mathcal{F}(\sigma_n)$, and so the asymmetric error exponent per channel use will be given by $\frac{1}{n} D_H^\varepsilon(\mathcal{E}(\rho_n) \| \mathcal{F}(\sigma_n))$. We write the optimal adaptive type II error rate for a given finite number of channel uses n as

$$e_A(n, \varepsilon, \mathcal{E}, \mathcal{F}) \coloneqq \sup_{\substack{\rho_1 \in \mathcal{D}(RA) \\ \{\Lambda_i : RB \to RA\}_{i=2}^n}} \frac{1}{n} D_H^\varepsilon(\mathcal{E}(\rho_n) \| \mathcal{F}(\sigma_n)) \tag{3.1.6}$$

$$= \sup_{\substack{\rho_1 \in \mathcal{D}(RA) \\ \{\Lambda_i : RB \to RA\}_{i=2}^n}} \frac{1}{n} D_H^\varepsilon\left(\left[\bigcirc_{i=2}^n (\mathcal{E} \circ \Lambda_i) \right] (\mathcal{E}(\rho_1)) \,\middle\|\, \left[\bigcirc_{i=2}^n (\mathcal{F} \circ \Lambda_i) \right] (\mathcal{F}(\rho_1)) \right) \tag{3.1.7}$$

where the supremum is over every initial input state $\rho_1 = \sigma_1$ and all subsequent input preparation CPTP maps $\{\Lambda_i\}_{i=2}^n$, and the $\bigcirc_{i=2}^n$ in the last line denotes successive function (left-)composition. It was shown in Katariya and Wilde (2021), Prop. 3, that this is computable as a semi-definite program.

Adaptive Strategies include Parallel Strategies

Every parallel strategy is also an adaptive strategy. This can be seen as follows: For a given parallel strategy, specified by an input state ν_{RA^n}, define the system R' as $R' \cong RA^{n-1}B^n$ which will be the reference system of the adaptive strategy we construct, and let ω_A and ω_B be two arbitrary constant states on A and B (these will be placeholders for parts of the reference system that are not yet, or no longer used, since we assumed the size of the reference system to stay constant throughout the protocol). We write $A^n = A_1^n$ for all n copies of the A system, and A_2^n for the $(n-1)$ copies without the first one. We then identify $\nu_{RA^n} = \nu_{RA_1^n} = \nu_{RA_2^n A_1}$, and pick the starting state of our adaptive strategy as $\rho_1 = \sigma_1 = \nu_{R' A_2^n A_1} \otimes \omega_B^n$ which is a state on $\mathcal{D}(R'A)$. After the application of the first black box, we will then have a state on $\mathcal{D}\left(RA^{n-1}BB^n\right)$, and we will choose the preparation operation Λ_2 to do the following: First, discard one the B systems (they all hold the placeholder state ω_B), and add one A system in the placeholder state ω_A, finally pick one of the A systems that is left from the original state ν and put this into the input register for the next channel, and put everything else (which will have system signature $RA^{n-1}B^n$) into the reference slot. This can then be repeated n-times, where we always keep all the B systems that are channel outputs, and only discard those systems that still hold the placeholder state. It is easy to see that this procedure leads to the state $\Lambda^{\otimes n}(\nu_{RA^n}) \otimes \omega_A^{n-1}$, at the end, where Λ is either $\mathcal{E}$ or $\mathcal{F}$, and discriminating this is obviously equivalent to just discriminating $\Lambda^{\otimes n}(\nu_{RA^n})$, which is the output of the parallel strategy. Thus, every parallel strategy can be written as an adaptive strategy, however the converse is not true, as the example from Harrow et al. (2010) shows that there are things one can do with an adaptive strategy that are not possible with a parallel one.

3.1.1 Asymptotic Equality in the Asymmetric Setting

For general adaptive strategies, it has been shown in Wilde et al. (2020) and X. Wang and Wilde (2019b) that asymptotically, when the number of channel uses goes to infinity, the best exponential decay rate of the type II error (per channel use) such that the type I error still goes to zero is given by the amortized Umegaki channel divergence, i.e.

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty} e_A(n,\varepsilon,\mathcal{E},\mathcal{F}) = \lim_{\varepsilon\to 0}\lim_{n\to\infty}\sup_{\rho_1,\{\Lambda_i\}_i}\frac{1}{n}D_H^\varepsilon(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) = D^A(\mathcal{E}\|\mathcal{F}) \tag{3.1.8}$$

where

$$D^A(\mathcal{E}\|\mathcal{F}) \coloneqq \sup_{\rho,\sigma\in\mathcal{D}(R\otimes A)}\big[D(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) - D(\rho\|\sigma)\big]. \tag{3.1.9}$$

Note that the dimension of the reference system R in this last supremum can be arbitrarily large.

The chain rule from Fang et al. (2020) states that this amortized divergence is in fact equal to the regularized channel divergence, i.e.

$$D^A(\mathcal{E}\|\mathcal{F}) = D^{\text{reg}}(\mathcal{E}\|\mathcal{F}) \coloneqq \lim_{n\to\infty}\frac{1}{n}\sup_{\nu\in\mathcal{D}(R^{\otimes n}\otimes A^{\otimes n})} D(\mathcal{E}^{\otimes n}(\nu)\|\mathcal{F}^{\otimes n}(\nu)), \tag{3.1.10}$$

the latter of which can be achieved by parallel protocols, as a consequence of X. Wang and Wilde (2019b), Theorem 3. Note that the reference system R in the latter optimization can be chosen isomorphic to A. Hence, asymptotically (and in the regime in which the type I error goes to zero) adaptive strategies offer no advantage over parallel ones. In Section 3.3 we provide a finite n version of this statement, by giving explicit bounds on how much the error probabilities of adaptive and parallel strategies can differ for a finite number of channel uses.

3.2 Computing n-Shot Error Exponents

Before we show in the next section how the n-shot error exponents of adaptive and parallel strategies are related, let us study how and to what degree one can actually compute these n-shot error exponents. As already mentioned, both the adaptive and parallel error rates e_P and e_A can be phrased in terms of a semi-definite program (X. Wang and Wilde 2019b, Katariya and Wilde 2021), although the size of this program grows with the dimension of the Hilbert space of the joint channels $\mathcal{E}^{\otimes n}$ and $\mathcal{F}^{\otimes n}$ and is thus exponential in n. In this section we show that for the parallel error rate e_P one can use the permutation invariance of the problem to actually reduce the size of the SDP to something polynomial in n, and thus also calculate the asymptotic error exponent (and the optimal parallel strategy) in time $\mathcal{O}(\text{poly}(n))$. This procedure does not work for the adaptive error rate e_A, as the flow of causality breaks the permutation invariance of this problem (applying permutations to adaptive strategies can lead to non-physical strategies where outputs of later channel uses are used as inputs for previous channel uses), and so we are not aware of any way to calculate the adaptive error rate e_A in time $\mathcal{O}(\text{poly}(n))$. Our main result is the following:

Theorem 3.2.1. *Let $\mathcal{E}, \mathcal{F} : A \to B$ be two quantum channels such that $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$. Then, for a fixed $\varepsilon \in [0, 1)$, the quantity*

$$e_P(n, \varepsilon, \mathcal{E}, \mathcal{F}) = \frac{1}{n} D_H^\varepsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) = \frac{1}{n} \sup_{\nu \in \mathcal{D}(R^{\otimes n} \otimes A^{\otimes n})} D_H^\varepsilon(\mathcal{E}^{\otimes n}(\nu)\|\mathcal{F}^{\otimes n}(\nu)) \tag{3.2.1}$$

can be computed up to an additive error $\delta > 0$ in time $\mathcal{O}(\mathrm{poly}(n) \log(\frac{1}{\delta}))$ as $n \to \infty$ and $\delta \to 0$.

Throughout this section, the system R is understood to be isomorphic to A. The starting point of our proof is the following Lemma:

Lemma 3.2.2 (X. Wang and Wilde 2019b, Prop. 2)**.** *The quantity $2^{-D_H^\varepsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n})}$ can be expressed as the following semi-definite program over operators $\Omega_{R^n B^n} \in \mathcal{B}\left(\mathcal{H}_R^{\otimes n} \otimes \mathcal{H}_B^{\otimes n}\right)$ and $\rho_{R^n} \in \mathcal{B}\left(\mathcal{H}_R^{\otimes n}\right)$*

$$\begin{aligned} \underset{\Omega_{R^n B^n},\, \rho_{R^n}}{\text{minimize}} \quad & \mathrm{Tr}\left(\Omega_{R^n B^n} \Gamma^{\mathcal{F}^{\otimes n}}_{R^n B^n}\right) \\ \text{subject to} \quad & \mathrm{Tr}\left(\Omega_{R^n B^n} \Gamma^{\mathcal{E}^{\otimes n}}_{R^n B^n}\right) \geq 1 - \varepsilon, \\ & 0 \leq \Omega_{R^n B^n} \leq \rho_{R^n} \otimes \mathbb{1}_{B^n}, \\ & \mathrm{Tr}(\rho_{R^n}) = 1, \end{aligned} \tag{3.2.2}$$

where $\Gamma^{\mathcal{E}}_{RB}$ is the Choi matrix of $\mathcal{E}$ and it is easy to see that $\Gamma^{\mathcal{E}^{\otimes n}}_{R^n B^n} = (\Gamma^{\mathcal{E}}_{RB})^{\otimes n}$.

The operators $\Omega_{R^n B^n}$ and ρ_{R^n} in this SDP have a number of parameters exponential in n, but we will show that we can use the permutation symmetry of this problem to rephrase it as an SDP polynomial in n.

Throughout this section we use $\mathcal{H}$ to denote a general Hilbert space, which we will then choose to be either $\mathcal{H}_R$, $\mathcal{H}_B$, or $\mathcal{H}_R \otimes \mathcal{H}_B$ in different situations. We will denote any objects that depend on the chosen Hilbert space with a subscript or superscript $\mathcal{H}$ (such as $P_{\mathcal{H}}$ in the following paragraph), where we then replace the subscript or superscript with R, B, or RB whenever we choose a specific Hilbert space; e.g., we write P_R, P_B, or P_{RB}. The next few subsections establish all the necessary properties for the proof of Theorem 3.2.1.

3.2.1 Permutation Invariant Operators

Remember our notation for permutations and permutation invariant objects: For any permutation $\pi \in \mathfrak{S}_n$ (where $\mathfrak{S}_n$ is the symmetric group) we write $P_{\mathcal{H}}(\pi)$ for the permutation matrix corresponding to the action of π on $\mathcal{H}^{\otimes n}$ by permuting the n copies of $\mathcal{H}$. It is then easy to see that these permutation matrices are unitary and $P_{\mathcal{H}}(\pi)^\dagger = P_{\mathcal{H}}(\pi^{-1})$. For any operator $X \in \mathcal{B}\left(\mathcal{H}^{\otimes n}\right)$ we also write the group average as

$$\overline{X} := \frac{1}{|\mathfrak{S}_n|} \sum_{\pi \in \mathfrak{S}_n} P_{\mathcal{H}}(\pi) X P_{\mathcal{H}}(\pi)^\dagger, \tag{3.2.3}$$

and the subspace of all permutation invariant operators as

$$\mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) := \left\{A \in \mathcal{B}\left(\mathcal{H}^{\otimes n}\right) \mid P_{\mathcal{H}}(\pi) A P_{\mathcal{H}}(\pi)^\dagger = A,\ \forall \pi \in \mathfrak{S}_n\right\}. \tag{3.2.4}$$

Such subspaces of permutation invariant operators will be useful, since it is not hard to see that the optimizers in Lemma 3.2.2 can always be chosen from within them:

Lemma 3.2.3. *The minimum in the SDP of Lemma 3.2.2 can be restricted to permutation invariant operators* $\Omega_{R^nB^n} \in \mathrm{End}^{\mathfrak{S}_n}((RB)^n)$, $\rho_{R^n} \in \mathrm{End}^{\mathfrak{S}_n}(R^n)$.

Proof. Let $(\Omega_{R^nB^n}, \rho_{R^n})$ be feasible for the optimization problem; i.e., they satisfy the constraints of (3.2.2). Let $\pi \in \mathfrak{S}_n$ be any permutation; then

$$\mathrm{Tr}\Big(P_R(\pi)\,\rho_{R^n}\,P_R(\pi)^\dagger\Big) = \mathrm{Tr}(\rho_{R^n}) \tag{3.2.5}$$

since the $P_R(\pi)$ are unitary, and thus also $\mathrm{Tr}(\rho_{R^n}) = \mathrm{Tr}(\overline{\rho_{R^n}})$. Similarly, we get

$$0 \leq P_{RB}(\pi)\,\Omega_{R^nB^n}\,P_{RB}(\pi)^\dagger \leq P_R(\pi)\rho_{R^n}P_R(\pi)^\dagger \otimes \mathbb{1}_{B^n} \tag{3.2.6}$$

since positivity is preserved under unitary conjugation, and we also used

$$P_{RB}(\pi) = (\mathbb{1}_{R^n} \otimes P_B(\pi))(P_R(\pi) \otimes \mathbb{1}_{B^n}) \tag{3.2.7}$$

which is immediate from the definition. This again implies that

$$0 \leq \overline{\Omega_{R^nB^n}} \leq \overline{\rho_{R^n}} \otimes \mathbb{1}_{B^n}\,. \tag{3.2.8}$$

By the cyclicity of the trace, and the fact that the Choi matrix of $\mathcal{E}^{\otimes n}$ is the tensor product of the Choi matrix of $\mathcal{E}$ (i.e. $\Gamma^{\mathcal{E}^{\otimes n}}_{R^nB^n} = (\Gamma^{\mathcal{E}}_{RB})^{\otimes n}$), we also see that

$$\mathrm{Tr}\Big(\overline{\Omega_{R^nB^n}}\Gamma^{\mathcal{E}^{\otimes n}}_{R^nB^n}\Big) = \mathrm{Tr}\Big(\Omega_{R^nB^n}\Gamma^{\mathcal{E}^{\otimes n}}_{R^nB^n}\Big)\,, \tag{3.2.9}$$

$$\mathrm{Tr}\Big(\overline{\Omega_{R^nB^n}}\Gamma^{\mathcal{F}^{\otimes n}}_{R^nB^n}\Big) = \mathrm{Tr}\Big(\Omega_{R^nB^n}\Gamma^{\mathcal{F}^{\otimes n}}_{R^nB^n}\Big)\,. \tag{3.2.10}$$

Hence, also $(\overline{\Omega_{R^nB^n}}, \overline{\rho_{R^n}})$ are feasible for the optimization problem (i.e., they satisfy the constraints) and also achieve the exact same value as $(\Omega_{(RB)^n}, \rho_{R^n})$ does. The group averages are elements of $\mathrm{End}^{\mathfrak{S}_n}((RB)^n)$ and $\mathrm{End}^{\mathfrak{S}_n}(R^n)$ respectively, and hence we can restrict the optimization in (3.2.2) to such permutation invariant operators. □

While the dimension of the subspace of these permutation invariant operators is polynomial in n, this does not yet show computability in $\mathrm{poly}(n)$ time using standard SDP solvers, as the problem is not phrased using matrices of size $\mathrm{poly}(n)$. In order to rephrase our problem in this way we follow the approach of O. Fawzi, Shayeghi, and Ta (2022), where a similar statement has been shown for the sharp Rényi divergence $D^{\#}_\alpha(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n})$. The key idea is to construct a suitable basis of the permutation invariant subspaces and then rephrase the SDP as one in which we minimize over (now only $O(\mathrm{poly}(n))$ many) basis coefficients.

Basis of Permutation Invariant Operator Subspace

As explained in O. Fawzi, Shayeghi, and Ta (2022), page 7352 (see also de Klerk, Pasechnik, and Schrijver (2007) and Anjos and Lasserre (2012)), one can construct such an orthogonal basis $C_r^{\mathcal{H}}$, $r \in \{1, \ldots, m_{\mathcal{H}}\}$ of $\mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ (where orthogonality is with respect to the Hilbert-Schmidt inner product, and $m_{\mathcal{H}} = \dim \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$) as follows: Let $\{|i\rangle\}_{i=1}^{d_{\mathcal{H}}}$ be an orthonormal basis of $\mathcal{H}$. Then, for any multi-index $\mathbf{i} \in \{1, \ldots, d_{\mathcal{H}}\}^n$ define the associated vector $|\mathbf{i}\rangle = \bigotimes_{k=1}^n |\mathbf{i}_k\rangle$ on the tensor-product system $\mathcal{H}^{\otimes n}$, and it is immediate that the set of all these vectors forms an orthonormal basis of $\mathcal{H}^{\otimes n}$. For any pair of such multi-indices $(\mathbf{i}, \mathbf{j})$ – this pair should be thought of as indexing a matrix element of an operator on $\mathcal{H}^{\otimes n}$ – we write the group orbit of these indices under the action of $\mathfrak{S}_n$ as

$$O(\mathbf{i}, \mathbf{j}) \coloneqq \{ \, (\pi(\mathbf{i}), \pi(\mathbf{j})) \mid \pi \in \mathfrak{S}_n \, \} \, , \tag{3.2.11}$$

where $\pi(\mathbf{i})$ permutes the components of $\mathbf{i}$, i.e., $\pi(\mathbf{i})_k = \mathbf{i}_{\pi^{-1}(k)}$ for $k \in \{1, \ldots, n\}$. There are exactly $m_{\mathcal{H}} = \dim \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ such group orbits, which we will label as $O_r^{\mathcal{H}}$, $r \in \{1, \ldots, m_{\mathcal{H}}\}$, and a representative element of each such orbit (i.e., a pair $(\mathbf{i}, \mathbf{j}) \in O_r^{\mathcal{H}}$) can be efficiently computed given r, as shown in O. Fawzi, Shayeghi, and Ta (2022), page 7352. This corresponds to the intuition that for $A \in \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ we require $A_{\mathbf{ij}} = A_{\pi(\mathbf{i})\pi(\mathbf{j})}$ for any $\pi \in \mathfrak{S}_n$; hence the matrix elements of A have to be constant on each group orbit, and so the number of group orbits is equal to the dimension of $\mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$. The basis elements $C_r^{\mathcal{H}} \in \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ are now defined by specifying their matrix elements as

$$(C_r^{\mathcal{H}})_{\mathbf{ij}} \coloneqq \begin{cases} 1 & \text{if } (\mathbf{i}, \mathbf{j}) \in O_r^{\mathcal{H}} \\ 0 & \text{otherwise.} \end{cases} \qquad r = 1, \ldots, m_{\mathcal{H}}. \tag{3.2.12}$$

It follows immediately from the definition that the $C_r^{\mathcal{H}}$ are orthogonal, and since we have $m_{\mathcal{H}} = \dim \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ of them, they form a basis.

As explained in O. Fawzi, Shayeghi, and Ta (2022), page 7353, for any matrix $A^{\otimes n} \in \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$, its coefficients with respect to the basis $C_r^{\mathcal{H}}$ can be computed straightforwardly from a description of A by picking a representative of each orbit, and these coefficients are hence computable in $\mathcal{O}(\mathrm{poly}(n))$ time. Specifically, one can show that for any r and any pair of indices in the group orbit $(\mathbf{i}, \mathbf{j}) \in O_r^{\mathcal{H}}$, the corresponding basis coefficient is given by

$$\gamma_r \coloneqq \prod_{k=1}^n A_{\mathbf{i}_k \mathbf{j}_k} \, , \tag{3.2.13}$$

i.e., we get

$$A^{\otimes n} = \sum_{r=1}^{m_{\mathcal{H}}} \gamma_r C_r^{\mathcal{H}} \, . \tag{3.2.14}$$

This can be applied to the Choi matrix $\Gamma_{R^n B^n}^{\mathcal{E}^{\otimes n}} = (\Gamma_{RB}^{\mathcal{E}})^{\otimes n}$, and hence we write

$$(\Gamma_{RB}^{\mathcal{E}})^{\otimes n} = \sum_{r=1}^{m_{RB}} \gamma_r^{\mathcal{E}} C_r^{RB} \tag{3.2.15}$$

and similarly when $\mathcal{E}$ is replaced by $\mathcal{F}$.

Algebra Isomorphism

While the previously introduced basis allows for an efficient parametrization of the elements of the permutation invariant operator subspace, we need to also ensure that we can efficiently treat the positivity condition in the SDP. This can be done via the following algebra isomorphisms: For any finite-dimensional Hilbert space $\mathcal{H}$, there exists a $*$-algebra isomorphism from $\mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ to block-diagonal matrices (Gijswijt 2005, Theorem 1)

$$\phi_{\mathcal{H}} : \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) \to \bigoplus_{i=1}^{t_{\mathcal{H}}} \mathbb{C}^{m_i \times m_i}, \tag{3.2.16}$$

where

$$t_{\mathcal{H}} \leq (n+1)^{d_{\mathcal{H}}}, \tag{3.2.17}$$

$$d_{\mathcal{H}} = \dim(\mathcal{H}), \tag{3.2.18}$$

$$\sum_{i=1}^{t_{\mathcal{H}}} m_i^2 = \dim(\mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})) \leq (n+1)^{d_{\mathcal{H}}^2}. \tag{3.2.19}$$

Introducing the notation $M_{\mathcal{H}} \coloneqq \sum_{i=1}^{t_{\mathcal{H}}} m_i$, for any $A \in \mathrm{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$ we have that $\phi_{\mathcal{H}}(A) \in \mathbb{C}^{M_{\mathcal{H}} \times M_{\mathcal{H}}}$, and we write $[\![\phi_{\mathcal{H}}(A)]\!]_i$ for the i-th block of $\phi_{\mathcal{H}}(A)$. Crucially, by O. Fawzi, Shayeghi, and Ta (2022), Lemma 3.3, Polak (2019), Prop. 2.4.4,

$$A \geq 0 \Leftrightarrow \phi_{\mathcal{H}}(A) \geq 0 \Leftrightarrow [\![\phi_{\mathcal{H}}(A)]\!]_i \geq 0 \ \forall i \in \{1, \ldots, t_{\mathcal{H}}\}. \tag{3.2.20}$$

From Gijswijt (2005), Theorem 1 it also follows that $\phi_{\mathcal{H}}$ preserves orthogonality (i.e., if $\mathrm{Tr}\left(A^\dagger B\right) = 0$ then also $\mathrm{Tr}\left(\phi_{\mathcal{H}}(A)^\dagger \phi_{\mathcal{H}}(B)\right) = 0$) and that $\phi_{\mathcal{H}}$ is unital (i.e. $\phi_{\mathcal{H}}(\mathbb{1}_{\mathcal{H}^{\otimes n}}) = \mathbb{1}_{M_{\mathcal{H}}}$).

Efficient Calculation of Coefficients

After expanding our states in the above basis of the permutation invariant subspace and using the above introduced algebra isomorphism, our SDP will contain some coefficients including these two objects. Here we show that all of these coefficients can be computed efficiently.

Lemma 3.2.4 (O. Fawzi, Shayeghi, and Ta (2022), pages 7352-7353, see also Litjens, Polak, and Schrijver (2017))**.** *For every $r \in \{1, \ldots, m_{\mathcal{H}}\}$, and $i \in \{1, \ldots, t_{\mathcal{H}}\}$, $[\![\phi_{\mathcal{H}}(C_r^{\mathcal{H}})]\!]_i$ can be computed in time $\mathcal{O}(\mathrm{poly}(n))$, and hence also $\phi_{\mathcal{H}}(C_r^{\mathcal{H}})$ can be computed in time $\mathcal{O}(\mathrm{poly}(n))$.*

We will require a slight extension of this result, that also applies to a joint isomorphism ϕ_{RB} only being applied to a basis element C_R^r of a subsystem:

Corollary 3.2.5. *For every $r \in \{1, \ldots, m_R\}$, $\phi_{RB}(C_r^R \otimes \mathbb{1}_{B^n})$ can be computed in time $\mathcal{O}(\mathrm{poly}(n))$.*

Proof. As $\mathbb{1}_{B^n} = \mathbb{1}_B^{\otimes n}$, the coefficients of $\mathbb{1}_{B^n}$ with respect to the $\{C_r^B\}$ basis can be computed efficiently just as for the Choi matrices above. Additionally, from the construction of the $C_r^{\mathcal{H}}$, there is an obvious one-to-one mapping $C_r^R \otimes C_{r'}^B = C_{r''}^{RB}$ for some r'' for every pair of (r, r'), and the result then follows from Lemma 3.2.4. □

Lemma 3.2.6. *For every $r \in \{1, ..., m_{\mathcal{H}}\}$, $\mathrm{Tr}\left(C_r^{\mathcal{H}}\right)$ can be computed in time $\mathcal{O}(\mathrm{poly}(n))$.*

Proof. As explained above, for every $r \in \{1, ..., m_{\mathcal{H}}\}$ one can efficiently find a representative $(\mathbf{i}, \mathbf{j})$ of the orbit $O_r^{\mathcal{H}}$. Given the definition of $C_r^{\mathcal{H}}$, its trace is then given by counting the number of different $k \in \{1, ..., n\}$ where $\mathbf{i}_k = \mathbf{j}_k$. □

Lemma 3.2.7. *For every $r \in \{1, ..., m_{\mathcal{H}}\}$, $\mathrm{Tr}\left((C_r^{\mathcal{H}})^\dagger C_r^{\mathcal{H}}\right)$ can be computed in time $\mathcal{O}(\mathrm{poly}(n))$.*

Proof. From the construction of the $C_r^{\mathcal{H}}$, it is immediate that the norm $\mathrm{Tr}\left((C_r^{\mathcal{H}})^\dagger C_r^{\mathcal{H}}\right)$ is given by the size of the group orbit, which is equal to the number of distinct permutations of an element $(\mathbf{i}, \mathbf{j})$ in the orbit. Concretely, if $\ell \in \{1, \ldots, d_{\mathcal{H}}^2\}$ indexes all pairs of single-system indices (i, j), $i, j \in \{1, \ldots, d_{\mathcal{H}}\}$, then given $(\mathbf{i}, \mathbf{j}) \in O_r^{\mathcal{H}}$, let K_ℓ denote the number of occurrences of $\ell = (i, j)$ in the multi-index $(\mathbf{i}, \mathbf{j})$, i.e., the number of different k values ($k \in \{1, \ldots, n\}$) for which $\mathbf{i}_k = i, \mathbf{j}_k = j$. The size of the group orbit is then given by

$$\mathrm{Tr}\left((C_r^{\mathcal{H}})^\dagger C_r^{\mathcal{H}}\right) = |O_r^{\mathcal{H}}| = \binom{n}{K_1, \ldots, K_{d_{\mathcal{H}}^2}} \tag{3.2.21}$$

and this multinomial coefficient can be calculated in time $\mathcal{O}(d_{\mathcal{H}}^2)$; see, e.g., Araujo, Sansão, and Vale-Cardoso (2021). □

3.2.2 Additive and Multiplicative Error

The SDP in Lemma 3.2.2 does not actually calculate $D_H^\varepsilon(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n})$ directly, but only $2^{-D_H^\varepsilon(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n})}$, and so standard SDP error bounds (i.e. Lemma 2.7.3) directly apply only to this latter quantity. Here, we show that due to the scaling of $D_H^\varepsilon(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n})$ with n, one can recover a similar error scaling also for this quantity (we will show in the proof of Theorem 3.2.1 below explicitly that $D_H^\varepsilon(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n})$ satisfies the conditions of this Lemma):

Lemma 3.2.8.
Let R_n be the solution of an SDP that can be computed up to additive error δ in time $\mathcal{O}(\mathrm{poly}(n) \log(1/\delta))$. Then, if $-\log(R_n) = \mathcal{O}(\mathrm{poly}(n))$, also $-\log(R_n)$ can be calculated up to additive error δ in time $\mathcal{O}(\mathrm{poly}(n) \log(1/\delta))$.

Proof. By assumption, we can compute R_n up to additive error δ' in time $\mathcal{O}(\mathrm{poly}(n) \log(1/\delta'))$. Now, for any given $\delta > 0$, let us pick the error we want to allow in the solution of R_n as $\delta' := (2^\delta - 1)R_n$. This error then propagates to $-\log(R_n)$ via

$$-\log\left(R_n + \delta'\right) = -\log\left(R_n 2^\delta\right) = -\log(R_n) - \delta \tag{3.2.22}$$

and hence this leads to an additive error of δ for the problem of calculating $-\log(R_n)$.[1] It is easy to see that $\log\left(2^{\delta}-1\right) = \log(\delta) + \mathcal{O}(1)$ as $\delta \to 0$, which implies $\log(1/\delta') = -\log(R_n) - \log\left(2^{\delta}-1\right) = \mathcal{O}(\text{poly}(n)) + O(\log(1/\delta)) = O(n\log(1/\delta))$ and hence $-\log(R_n)$ can be calculated up to an error δ in time $\mathcal{O}(\text{poly}(n)\log(1/\delta')) = \mathcal{O}(\text{poly}(n)\log(1/\delta))$.

□

3.2.3 Proof of Theorem 3.2.1 (Poly-Time Computation of Optimal Parallel Strategy)

We now have all the required tools to complete the proof of Theorem 3.2.1.

Proof of Theorem 3.2.1. We start with Lemma 3.2.2, and its expression of $2^{-D_H^{\varepsilon}(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n})}$ as an SDP. By Lemma 3.2.3 this SDP will always have optimizers in the permutation invariant operator subspaces. Remember that $\{C_r^{\mathcal{H}}\}_{r\in\{1,\dots,m_{\mathcal{H}}\}}$ is a basis for $\text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$, and hence we can expand $\Omega_{R^nB^n} \in \text{End}^{\mathfrak{S}_n}(R^nB^n)$ and $\rho_{R^n} \in \text{End}^{\mathfrak{S}_n}(R^n)$ as follows:

$$\Omega_{R^nB^n} = \sum_{r=1}^{m_{RB}} y_r C_r^{RB}, \tag{3.2.23}$$

$$\rho_{R^n} = \sum_{r=1}^{m_R} z_r C_r^{R}, \tag{3.2.24}$$

where $\{y_r \in \mathbb{C}\}_{r=1}^{m_{RB}}$ and $\{z_r \in \mathbb{C}\}_{r=1}^{m_R}$ are the respective basis coefficients. Note that since the C_r are not necessarily Hermitian, the coefficients are not necessarily real. Since optimizing over elements of a vector space is obviously equivalent to optimizing over their basis coefficients, we can rephrase our SDP from Lemma 3.2.2 as follows:

$$\begin{aligned}
&\underset{\{y_r\}_{r=1}^{m_{RB}},\{z_r\}_{r=1}^{m_R}}{\text{minimize}} \quad \sum_{r=1}^{m_{RB}} y_r(\gamma_r^{\mathcal{F}})^* \operatorname{Tr}\left((C_r^{RB})^\dagger C_r^{RB}\right) \\
&\quad \text{subject to} \\
&\sum_{r=1}^{m_{RB}} y_r(\gamma_r^{\mathcal{E}})^* \operatorname{Tr}\left((C_r^{RB})^\dagger C_r^{RB}\right) \geq 1-\varepsilon, \\
&\sum_{r=1}^{t_R} z_r \operatorname{Tr}\left(C_r^R\right) = 1\,, \\
&0 \leq \sum_{r=1}^{m_{RB}} y_r [\![\phi_{RB}(C_r^{RB})]\!]_i \leq \sum_{r=1}^{m_R} z_r [\![\phi_{RB}(C_r^R \otimes \mathbb{1}_{B^n})]\!]_i \quad \forall i \in \{1,\dots,t_{RB}\}
\end{aligned} \tag{3.2.25}$$

where we also used (3.2.20) and (3.2.15).

We will show that this is an SDP in $\text{poly}(n)$ variables with $\text{poly}(n)$ constraints by casting it into standard form. To simplify notation we will write the Hilbert-Schmidt inner product using $\langle\cdot,\cdot\rangle$, i.e., $\langle A, B\rangle :=$

[1] The signs here are chosen as the errors appear in practice for our application: the SDP (3.2.2) is a minimization, so any numerical approximation will be larger than the true value and the additive error hence positive, which then translates to a value smaller than the true value for our original problem (3.2.1).

$\operatorname{Tr}\big(A^\dagger B\big)$ and also write $M = M_{RB}$ (i.e., $M = M_{\mathcal{H}}$ as below (3.2.19) with $\mathcal{H} = \mathcal{H}_R \otimes \mathcal{H}_B$). For $s \in \mathbb{C}$, consider the following $(2M+1) \times (2M+1)$ block-diagonal matrix

$$X \coloneqq \left[\sum_{r=1}^{m_{RB}} y_r \phi_{RB}(C_r^{RB})\right] \oplus \left[\sum_{r'=1}^{m_R} z_{r'} \phi_{RB}(C_{r'}^{R} \otimes \mathbb{1}_{B^n}) - \sum_{r=1}^{m_{RB}} y_r \phi_{RB}(C_r^{RB})\right] \oplus s \tag{3.2.26}$$

where s is a slack variable that turns the one inequality constraint into an equality constraint (see further below). From here on, we write $(\,\cdot\,) \oplus (\,\cdot\,) \oplus (\,\cdot\,)$ to specify a block-diagonal matrix with block sizes equal to the ones in (3.2.26). The constraint $X \geq 0$ now implies

$$0 \leq \sum_{r=1}^{m_{RB}} y_r [\![\phi_{RB}(C_r^{RB})]\!]_i \leq \sum_{r=1}^{m_R} z_r [\![\phi_{RB}(C_r^{R} \otimes \mathbb{1}_{B^n})]\!]_i \tag{3.2.27}$$

for $i = 1, \ldots, t_{RB}$. Additionally, since ϕ_{RB} preserves orthogonality, we can recover the coefficients y_r and z_r from X by taking inner products with suitable operators. Specifically, with

$$\tilde{Y}_r \coloneqq \phi_{RB}(C_r^{RB}) \oplus 0 \oplus 0\,, \tag{3.2.28}$$

$$Y_r \coloneqq \frac{\tilde{Y}_r}{\langle \tilde{Y}_r, \tilde{Y}_r \rangle}\,, \tag{3.2.29}$$

for $r \in \{1, \ldots, m_{RB}\}$, and with

$$\tilde{Z}_r \coloneqq \phi_{RB}(C_r^{R} \otimes \mathbb{1}_{B^n}) \oplus \phi_{RB}(C_r^{R} \otimes \mathbb{1}_{B^n}) \oplus 0\,, \tag{3.2.30}$$

$$Z_r \coloneqq \frac{\tilde{Z}_r}{\langle \tilde{Z}_r, \tilde{Z}_r \rangle}\,, \tag{3.2.31}$$

for $r \in \{1, \ldots, m_R\}$, we have

$$y_r = \langle Y_r, X \rangle, \tag{3.2.32}$$

$$z_r = \langle Z_r, X \rangle. \tag{3.2.33}$$

Hence we can rephrase the expressions in (3.2.25) as inner products with X. Specifically,

$$\sum_{r=1}^{m_{RB}} y_r (\gamma_r^{\mathcal{E}})^* \operatorname{Tr}\Big((C_r^{RB})^\dagger C_r^{RB}\Big) = \left\langle \sum_{r=1}^{m_{RB}} Y_r (\gamma_r^{\mathcal{E}})^* \operatorname{Tr}\Big((C_r^{RB})^\dagger C_r^{RB}\Big), X \right\rangle \tag{3.2.34}$$

and one can do the same with $\mathcal{E}$ replaced by $\mathcal{F}$.

We want to transform our SDP into one where we optimize over X, and for that we need to impose the necessary block-diagonal structure of X (including the block-diagonal substructure that comes from its parts being images of ϕ_{RB}). For this, consider the following linear space

$$\left\{ [\phi_{RB}(\Omega)] \oplus [\phi_{RB}(\rho \otimes \mathbb{1}_{B^n}) - \phi_{RB}(\Omega)] \oplus s \;\middle|\; \Omega \in \operatorname{End}^{\mathfrak{S}_n}(R^n B^n),\ \rho \in \operatorname{End}^{\mathfrak{S}_n}(R^n),\ s \in \mathbb{C} \right\} \tag{3.2.35}$$

and define $\mathcal{A}$ to be a set of matrices that form a basis of the orthogonal complement of this linear space, where we take the orthogonal complement within $\mathbb{C}^{(2M+1)\times(2M+1)}$. It is then easy to see that $|\mathcal{A}| \leq \dim(\mathbb{C}^{(2M+1)\times(2M+1)}) = (2M+1)^2 = \mathcal{O}(\text{poly}(n))$.

With $S \coloneqq (0 \oplus 0 \oplus 1)$, we can then introduce the following SDP

$$\begin{aligned}
&\underset{X}{\text{minimize}} \quad \left\langle \sum_{r=1}^{m_{RB}} Y_r (\gamma_r^{\mathcal{F}})^* \operatorname{Tr}\Big((C_r^{RB})^\dagger C_r^{RB} \Big), X \right\rangle \\
&\text{subject to} \\
&\left\langle \left(\sum_{r=1}^{m_{RB}} Y_r (\gamma_r^{\mathcal{E}})^* \operatorname{Tr}\Big((C_r^{RB})^\dagger C_r^{RB} \Big) \right) - S, X \right\rangle = 1 - \varepsilon, \\
&X \in \mathbb{C}^{(2M+1)\times(2M+1)}, \\
&X \geq 0, \\
&\left\langle \sum_{r=1}^{t_R} Z_r \operatorname{Tr}(C_r^R), X \right\rangle = 1, \\
&\langle A, X \rangle = 0 \qquad \forall A \in \mathcal{A}
\end{aligned} \tag{3.2.36}$$

which is in standard form with $\mathcal{O}(\text{poly}(n))$ many constraints and matrices of size $\mathcal{O}(\text{poly}(n))$.

This new SDP is equivalent to (3.2.25), which can be seen as follows: From (3.2.26) and the calculations thereafter it follows that for every feasible $(\{y_r\}_{r=1}^{m_{RB}}, \{z_r\}_{r=1}^{m_R})$ in (3.2.25) there is a corresponding feasible X in (3.2.36) which achieves the same value. Conversely, every X that satisfies the constraints of (3.2.36) has to lie in the subspace (3.2.35), and it is then immediate that it can be represented as in (3.2.26) for suitable y_r and z_r. These y_r and z_r then also achieve the same value in (3.2.25) as X did in (3.2.36).

It now only remains to show that we can also efficiently calculate all that is required to parameterize this SDP. For the Y_r, this follows from Lemma 3.2.4, and for the Z_r from Corollary 3.2.5. $\operatorname{Tr}(C_r^R)$ can be calculated efficiently by Lemma 3.2.6, $\operatorname{Tr}((C_r^{RB})^\dagger C_r^{RB})$ by Lemma 3.2.7, and the $\gamma_r^{\mathcal{E}}$ and $\gamma_r^{\mathcal{F}}$ by the discussion around (3.2.15). Finally, since we can calculate $\phi_{RB}(C_r^{RB})$ and $\phi_{RB}(C_r^R \otimes \mathbb{1}_{B^n})$, we can calculate a generating set of the subspace (3.2.35), and a basis for its orthogonal complement $\mathcal{A}$ (which is an orthogonal complement within a $\mathcal{O}(\text{poly}(n))$ dimensional space) is then also computable in $\mathcal{O}(\text{poly}(n))$ time.

As a last step, we would like to show that the problem is feasible for an X with bounded operator norm. For the SDP in Lemma 3.2.2, it is easy to see that $\Omega_{R^nB^n} \coloneqq \frac{1}{d_{R^n}} \mathbb{1}_{R^nB^n}$ and $\rho_{R^n} = \frac{1}{d_{R^n}} \mathbb{1}_{R^n}$, satisfy $\operatorname{Tr}\Big(\Omega_{R^nB^n} \Gamma_{R^nB^n}^{\mathcal{E}^{\otimes n}} \Big) = 1$, and so the problem is feasible. Now, since ϕ_{RB} is unital, these choices of $\Omega_{R^nB^n}$ and ρ_{R^n} translate to $X = (\frac{1}{d_{R^n}} \mathbb{1}) \oplus 0 \oplus \varepsilon$ (via the definition of X in (3.2.26)), which satisfies $\|X\|_\infty \leq 1$.

Assuming that we can disregard issues of numerical representation and accuracy, Lemma 2.7.3 then tells us that we can calculate

$$R_n \coloneqq 2^{-D_H^\varepsilon(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n})}, \tag{3.2.37}$$

up to additive error δ in time $\mathcal{O}(\text{poly}(n) \log(1/\delta))$. To show that we can also accurately calculate $D_H^\varepsilon(\mathcal{E}^{\otimes n} \| \mathcal{F}^{\otimes n})$, we apply Lemma 3.2.8, for which we have to show that $-\log(R_n) = \mathcal{O}(n)$ as $n \to \infty$. This follows from

Lemma 2.5.1, which states that for all $\gamma \in (0, 1-\varepsilon)$

$$D_H^{\varepsilon}(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) \leq D_{\max}^{\sqrt{1-\varepsilon-\gamma}}(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) + \log\left(\frac{4(\varepsilon+\gamma)}{\gamma^2}\right) \tag{3.2.38}$$

$$\leq D_{\max}(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) + \log\left(\frac{4(\varepsilon+\gamma)}{\gamma^2}\right) \tag{3.2.39}$$

$$= nD_{\max}(\mathcal{E}\|\mathcal{F}) + \log\left(\frac{4(\varepsilon+\gamma)}{\gamma^2}\right) = \mathcal{O}(n)\,. \tag{3.2.40}$$

Note that also the case $\varepsilon = 0$ (while not directly covered by Lemma 2.5.1) follows immediately by first using that $D_H^0(\mathcal{E}\|\mathcal{F}) \leq D_H^{\varepsilon'}(\mathcal{E}\|\mathcal{F})$ for all $\varepsilon' \in (0,1)$. □

3.3 Parallelizing an n-Shot Adaptive Strategy

After proving the computability of the n-shot parallel exponent in the previous section, our main result in this section is a theorem which links the n-shot parallel and adaptive exponents. We will state this main result in two forms, first in a simple manner that illustrates the main idea and structure of the result, and secondly in a more detailed theorem that gives a tighter bound, and which states in detail what one can choose as a parallel input state.

Remember that there is a trade-off in minimizing the type I and type II errors. In the context of a strategy, for a given type I error α we will write the best achievable type II error as $\beta(\alpha)$. We are especially interested in the exponential decay rate of the type II error with the number of channel uses; i.e., if some strategy involving n uses of the channel has type II error $\beta(\alpha)$, we are interested in the quantity $-\frac{1}{n}\log(\beta(\alpha))$. Our main result compares this error decay rate per channel use between adaptive and parallel strategies:

Corollary 3.3.1 (Main result, simple version). *Let $\mathcal{E}, \mathcal{F} : A \to B$ be two quantum channels such that $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$. Let there be given an adaptive discrimination strategy with n channel uses, and that – for an arbitrary type I error $\alpha_a \in [0,1]$ – achieves type II error $\beta_a(\alpha_a)$, and thus type-II error rate $\frac{-1}{n}\log(\beta_a(\alpha_a))$. Then, for all $\alpha_p \in (0,1]$ there exists a parallel strategy with m channel uses and type II error $\beta_p(\alpha_p)$ such that for all $\alpha_a \in [0,1]$ the type II error rates per channel use obey the following relation:*

$$-\frac{1}{m}\log(\beta_p(\alpha_p)) \geq -\frac{1-\alpha_a}{n}\log(\beta_a(\alpha_a)) - \frac{Cn}{\sqrt{m}}\log\left(\frac{8}{\alpha_p}\right) - \frac{1}{n}. \tag{3.3.1}$$

That is, the type II error rate of the parallel strategy is essentially at least as good as the adaptive one modulo an additional error term, which decays as $m \to \infty$. The constant C is given by

$$C := 7\log\left(2^{D_2(\mathcal{E}\|\mathcal{F})} + 2\right) \leq 7\log\left(2^{D_{\max}(\mathcal{E}\|\mathcal{F})} + 2\right). \tag{3.3.2}$$

Remark 3.3.2. If we take the limit $m \to \infty$ in (3.3.1), then $n \to \infty$, and finally $\alpha_a \to 0$ and $\alpha_p \to 0$, we find that asymptotically there exist parallel strategies with better or at least equal type II error decay than any adaptive one, and hence our result also implies the known result (Fang et al. 2020) that, asymptotically (and with $\alpha \to 0$), adaptive strategies offer no advantage over parallel ones.

Remark 3.3.3. It is known that for small n and m, and suitably chosen α_a and α_p, the type II error rates for adaptive and parallel strategies can be arbitrarily far apart (see Section 3.3.3 below for an example). From the asymptotic equivalence, we know that this difference has to vanish as n and m go to infinity, but the purely asymptotic statement does not tell us how exactly this vanishes, and what the required relationship between n and m is (in principle the required m to reach a similar rate as a given adaptive strategy with n channel uses could grow arbitrarily fast with n). Corollary 3.3.1 now tells us that the difference of type II error rates between an adaptive strategy and the corresponding parallel one will become arbitrarily small if $m = \omega(n^2)$ (i.e., m has to grow faster than n^2). Hence, given a sequence of adaptive strategies with n channel uses, we can convert these into parallel strategies using at most quadratically (or a little bit more than quadratically) as many channel uses each, and will achieve matching rates once n gets large enough. This quadratic relationship is universal in the sense that it holds for all pairs of channels $\mathcal{E}, \mathcal{F}$ with $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$ (where only the prefactor depends on the value of $D_{\max}(\mathcal{E}\|\mathcal{F})$).

Note again that we are not comparing type II errors in Corollary 3.3.1, but rather decay rates of the type II error per channel use. When we say that the rates of a parallel strategy with $m = \omega(n^2)$ channel uses and an adaptive strategy with n channel uses are roughly equal, the parallel strategy will have much smaller type II error because it has many more channel uses.

The following is the more refined version of our result, with a tighter bound and a description of the parallel input state:

Theorem 3.3.4 (Main result, technical version)**.** *Let $\mathcal{E}, \mathcal{F} : A \to B$ be quantum channels such that $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$. Given the original input state $\rho_1 \in \mathcal{D}(R_a \otimes A)$ and the CPTP preparation maps $\{\Lambda_i : R_a \otimes B \to R_a \otimes A\}_{i=2}^n$ of an arbitrary adaptive protocol with n channel uses, we write $\rho_i, \sigma_i \in \mathcal{D}(R_a \otimes A)$, $i \in \{1, \dots, n\}$ for the states that are input into the channel during the adaptive protocol (ρ_i if the channel is $\mathcal{E}$, and σ_i if the channel is $\mathcal{F}$; see Section 3.1 and Figure 3.2 above for a more detailed explanation of this notation, we also write R_a to denote the fixed (and potentially arbitrary large) reference system associated to this adaptive protocol). We define $\ell \in \{1, \dots, n\}$ as the step in the protocol where the distinguishability increases the most, i.e.,*

$$\ell := \operatorname*{arg\,max}_{k \in \{1,\dots,n\}} \left[D(\mathcal{E}(\rho_k)\|\mathcal{F}(\sigma_k)) - D(\rho_k\|\sigma_k) \right]. \tag{3.3.3}$$

Then, for all $\alpha_p \in (0, 1]$, and $m \in \mathbb{N}$, there exists a state $\nu \in \mathcal{D}(R^{\otimes m} \otimes A^{\otimes m})$ – where one has an amount of choice in picking it's reference system R, see below – such that for all $\alpha_a \in [0, 1]$:

$$\begin{aligned} \frac{1}{m} D_H^{\alpha_p}\left(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\right) \geq \frac{1-\alpha_a}{n} D_H^{\alpha_a}\left(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)\right) - \frac{c'_\ell}{\sqrt{m}}\left[\log\left(\frac{4}{\alpha_p}\right) + K\right] \\ - \frac{1}{m}\left[\log\left(\frac{1}{\alpha_p}\right) - \log\left(1 - \frac{\alpha_p}{4}\right)\right] - \frac{h(\alpha_a)}{n}, \end{aligned} \tag{3.3.4}$$

where

$$K := \frac{\ln(2)\log^2(3)}{8} \cosh\left(\frac{\log(3)}{2}\right) \leq 0.29 \tag{3.3.5}$$

and c'_ℓ depends on the pair of channels $\mathcal{E}, \mathcal{F}$ and can be bounded as follows:

$$c'_\ell \coloneqq \frac{4}{\log(3)} \inf_{\gamma_1,\gamma_2\in(0,1]} [c_{\gamma_1}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) + c_{\gamma_2}(\rho_\ell\|\sigma_\ell)] \tag{3.3.6}$$

$$\leq \frac{8\ell}{\log(3)} \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \tag{3.3.7}$$

$$\leq \frac{8n}{\log(3)} \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F})\,, \tag{3.3.8}$$

where

$$c_\gamma(\rho\|\sigma) \coloneqq \frac{1}{\gamma}\log\Big(2^{\gamma D_{1+\gamma}(\rho\|\sigma)} + 2^{-\gamma D_{1-\gamma}(\rho\|\sigma)} + 1\Big), \tag{3.3.9}$$

$$\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \coloneqq \frac{1}{\gamma}\log\Big(2^{\gamma \widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\Big)\,. \tag{3.3.10}$$

Moreover, if

$$m \geq \log\left(\frac{4}{\alpha_p}\right)\left(\frac{4}{\log(3)\sqrt{2\ln(2)}}\right)^2, \tag{3.3.11}$$

then we have the following tighter bound

$$\frac{1}{m} D_H^{\alpha_p}\big(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\big) \geq \frac{1-\alpha_a}{n} D_H^{\alpha_a}\big(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)\big) - \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{4}{\alpha_p}\right)} - \frac{1}{m}\left[\log\left(\frac{1}{\alpha_p}\right) - \log\Big(1-\frac{\alpha_p}{4}\Big)\right] - \frac{h(\alpha_a)}{n}\,, \tag{3.3.12}$$

where c_ℓ is defined in a similar way as c'_ℓ but with different numerical constants K_1 and K_2:

$$K_1 \coloneqq 2\sqrt{2\ln(2)\cosh\left(\frac{\log(3)}{2}\right)} \leq 2.72, \tag{3.3.13}$$

$$K_2 \coloneqq 2\sqrt{2\ln(2)} \leq 2.36, \tag{3.3.14}$$

$$c_\ell \coloneqq \inf_{\gamma_1,\gamma_2\in(0,1]} [K_1 c_{\gamma_1}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) + K_2 c_{\gamma_2}(\rho_\ell\|\sigma_\ell)] \tag{3.3.15}$$

$$\leq \ell(K_1+K_2) \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \tag{3.3.16}$$

$$\leq n(K_1+K_2) \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F})\,. \tag{3.3.17}$$

The parallel input state ν can be chosen as either:

- *An optimizer in the smoothing of the max-divergence $D_{\max}^{\varepsilon}(\rho_\ell^{\otimes m}\|\sigma_\ell^{\otimes m})$, where $\varepsilon = \frac{1}{2}\left(1-\sqrt{1-\alpha_p}\right)$, i.e.,*

$$\nu = \tilde{\nu}_{R_a^m A^m} \coloneqq \operatorname*{arg\,min}_{\tilde{\rho}\in B_\varepsilon^\circ(\rho_\ell^{\otimes m})} D_{\max}(\tilde{\rho}\|\sigma_\ell^{\otimes m})\,. \tag{3.3.18}$$

Note that the reference system R_a depends on the adaptive protocol and can be arbitrarily large, and hence also $\tilde{\nu}$ might have an arbitrarily large reference system.

- *The canonical (or any other) purification of the A^m-marginal $\tilde{\nu}_{A^m} = \mathrm{Tr}_{R_a^m}(\tilde{\nu}_{R_a^m A^m})$. That is, we can choose $\nu = |\psi_{R^m A^m}\rangle\langle\psi_{R^m A^m}|$, with*

$$|\psi_{R^m A^m}\rangle = \mathbb{1}_{R^m} \otimes \sqrt{\tilde{\nu}_{A^m}}\, |\Phi_{R^m A^m}\rangle \tag{3.3.19}$$

where R is isomorphic to A and $|\Phi_{R^m A^m}\rangle$ is an unnormalized maximally entangled state.

Remark 3.3.5. Even though the second choice for the parallel input state ν might seem preferable in most cases (due to the control over the size of the reference system R), it is not necessarily so. This is because even though the reference system of the first choice for ν could be very large, it might not always be. Since the overall state can be mixed, it might actually be smaller than the canonical purification of its marginal. Additionally, with the first choice for ν, one gets a bound that this parallel input state is ε-close to the input state of the adaptive strategy, which might be useful in cases where one wants to show that some property of the adaptive strategy (e.g., a satisfied energy constraint for the input states) is (approximately) satisfied also for the parallel strategy.

Remark 3.3.6. The constraint $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$ is necessary in general, as the example of Harrow et al. (2010) (see also Salek, Hayashi, and Winter (2022)) serves as a counterexample to our statement without this constraint. Specifically, Harrow et al. (2010) construct two channels $\mathcal{E}, \mathcal{F}$ for which its authors then show that there exists an adaptive strategy with only two channel uses that achieves perfect discrimination, i.e., both $\alpha_a = 0$ and $\beta_a = 0$. In our terminology this implies that for all $\alpha_a \in [0, 1]$

$$D_H^{\alpha_a}(\mathcal{E}(\rho_2)\|\mathcal{F}(\sigma_2)) = \infty\,. \tag{3.3.20}$$

On the other hand, Harrow et al. (2010) show for these two channels that with a parallel strategy, even with arbitrarily many channel uses, perfect discrimination can never be achieved (i.e., α_p and β_p cannot both be zero), which in our notation implies that for all m and for all $\nu \in \mathcal{D}\left(R \otimes A^{\otimes m}\right)$

$$D_H^0(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)) < \infty\,. \tag{3.3.21}$$

Since it is well-known (see, e.g., Khatri and Wilde (2020), Prop. 4.66) that for all states ρ, σ

$$\lim_{\alpha_p \to 0} D_H^{\alpha_p}(\rho\|\sigma) = D_H^0(\rho\|\sigma) \tag{3.3.22}$$

this means that for all m, one can find a sufficiently small α_p such that for all $\nu \in \mathcal{D}\left(R \otimes A^{\otimes m}\right)$

$$D_H^{\alpha_p}(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)) < \infty. \tag{3.3.23}$$

and thus a relation like (3.3.12) cannot hold for these two channels. Note that, even in this example, the Stein exponent, i.e., the optimal exponential decay rate of the type II error such that the type I error still goes

to zero, is still identical for the parallel and adaptive strategies, as it is infinite also for the optimal parallel strategy. This follows from

$$D_{\max}(\mathcal{E}\|\mathcal{F}) = \infty \quad \Rightarrow \quad D(\mathcal{E}\|\mathcal{F}) = D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) = \infty\,. \tag{3.3.24}$$

Computability

The usefulness of finite-length bounds often crucially depends on whether the quantities involved can actually be efficiently computed (Hayashi and Watanabe 2020, Watanabe and Hayashi 2017, Hayashi and Watanabe 2016). To demonstrate this for our bound, we consider the following two applications:

1. We are given an adaptive strategy, where we know (potentially only a lower bound on) $D_H^{\alpha_a}(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n))$, and want to employ the theorem to get a lower bound on the performance of the best-possible parallel strategy (i.e., we want to know how much worse could a parallel strategy potentially be).

   Such a lower bound can be computed using Theorem 3.3.4 in time $\mathcal{O}(1)$ (i.e., the computational complexity is independent of the number of channel uses n and m), by using the bound on c_ℓ from (3.3.17). A potentially much tighter bound, obtained by finding ℓ explicitly and then calculating c_ℓ, can be computed in time $\mathcal{O}(n)$.

2. We want to upper bound the performance of the best possible adaptive strategy with n channel uses, by finding the best possible parallel strategy and then using Theorem 3.3.4. In Theorem 3.2.1 we showed that the best possible parallel strategy with m channel uses can actually be calculated in $\mathcal{O}(\mathrm{poly}(m))$ time, and hence by choosing $m = n^{2+\xi}$ for $\xi > 0$ we also obtain asymptotically tight upper bounds on any adaptive strategy with n channel uses in $\mathcal{O}(\mathrm{poly}(n))$ time.

   To our knowledge this is the first such poly-time bound obtained in the literature. Specifically, computing the best parallel strategy and then using our bound is exponentially faster than any currently known way of optimizing over all adaptive strategies. While the authors of Katariya and Wilde (2021) showed that optimizing over all adaptive strategies can be phrased as an SDP, the size of this SDP grows exponentially in n, and there is no obvious symmetry (like the permutation invariance in the parallel case) that allows one to reduce the number of variables.

Proof Overview

The remainder of this section proves Theorem 3.3.4. The general idea of the proof is the following: We will start by moving from the hypothesis testing relative entropy of the adaptive strategy to the Umegaki relative entropy, using Lemma 2.4.1. Then, we will see that by using an amortization argument (equations (3.3.33)–(3.3.36) below), we can bound the performance of the adaptive strategy by the performance of just one of its steps, the step where the distinguishability of the two states (that occur if the channel is either $\mathcal{E}$ or $\mathcal{F}$) increases the most. We then consider m parallel copies of this step and construct a parallel input state using a chain rule for the smoothed max-relative entropy (Lemma 3.3.7); see Figure 3.3 below for an illustration of this step. The smoothed max-relative entropies can be related to the Umegaki relative entropies

we used in the amortization argument by the non-asymptotic bounds presented in Lemma 2.5.2. Finally we connect to the hypothesis testing relative entropy of a parallel protocol using Lemma 2.5.1.

This next subsection presents a key lemma we will use as part of our proof of Theorem 3.3.4.

3.3.1 A Simple One-Shot Version of the Chain Rule

The following is a variant of Fang et al. (2020), Prop. 3.2, where our different convention of smoothing the max-divergence (smoothing only over normalized states) leads to a tighter and simpler bound, and restricting to the case of a single input system also makes things a bit simpler.

Lemma 3.3.7. *Let $\mathcal{E}$ and $\mathcal{F}$ be arbitrary quantum channels from system A to system B, and let $\rho \in \mathcal{D}(A)$, $\sigma \in \mathscr{P}(A)$. Then for all $\varepsilon, \varepsilon' \in [0,1]$ there exists a state $\nu \in B^\circ_\varepsilon(\rho)$ such that*

$$D^{\varepsilon+\varepsilon'}_{\max}(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) \leq D^{\varepsilon}_{\max}(\rho\|\sigma) + D^{\varepsilon'}_{\max}(\mathcal{E}(\nu)\|\mathcal{F}(\nu))\,. \tag{3.3.25}$$

Moreover, ν can be chosen as

$$\nu = \operatorname*{arg\,min}_{\tilde{\rho}\in B^\circ_\varepsilon(\rho)} D_{\max}(\tilde{\rho}\|\sigma)\,. \tag{3.3.26}$$

Proof. Let $\nu \in B^\circ_\varepsilon(\rho)$ be an optimal choice for $D^{\varepsilon}_{\max}(\rho\|\sigma)$; i.e.,

$$\nu \leq 2^{D^{\varepsilon}_{\max}(\rho\|\sigma)}\sigma. \tag{3.3.27}$$

Since $\mathcal{F}$ is a positive map, this implies that

$$\mathcal{F}(\nu) \leq 2^{D^{\varepsilon}_{\max}(\rho\|\sigma)}\mathcal{F}(\sigma)\,. \tag{3.3.28}$$

Furthermore, let $\tau \in B^\circ_{\varepsilon'}(\mathcal{E}(\nu))$ be an optimal choice for $D^{\varepsilon'}_{\max}(\mathcal{E}(\nu)\|\mathcal{F}(\nu))$, so that

$$\tau \leq 2^{D^{\varepsilon'}_{\max}(\mathcal{E}(\nu)\|\mathcal{F}(\nu))}\mathcal{F}(\nu)\,. \tag{3.3.29}$$

Combining the last two inequalities leads to

$$\tau \leq 2^{D^{\varepsilon'}_{\max}(\mathcal{E}(\nu)\|\mathcal{F}(\nu))+D^{\varepsilon}_{\max}(\rho\|\sigma)}\mathcal{F}(\sigma)\,. \tag{3.3.30}$$

It remains to show that $P(\tau, \mathcal{E}(\rho)) \leq \varepsilon + \varepsilon'$ (where P is the purified distance). This follows from

$$P(\tau,\mathcal{E}(\rho)) \leq P(\tau,\mathcal{E}(\nu)) + P(\mathcal{E}(\nu),\mathcal{E}(\rho)) \leq \varepsilon' + P(\nu,\rho) \leq \varepsilon' + \varepsilon\,, \tag{3.3.31}$$

where we used the triangle inequality and the data-processing inequality for the purified distance (see, e.g., Khatri and Wilde 2020). □

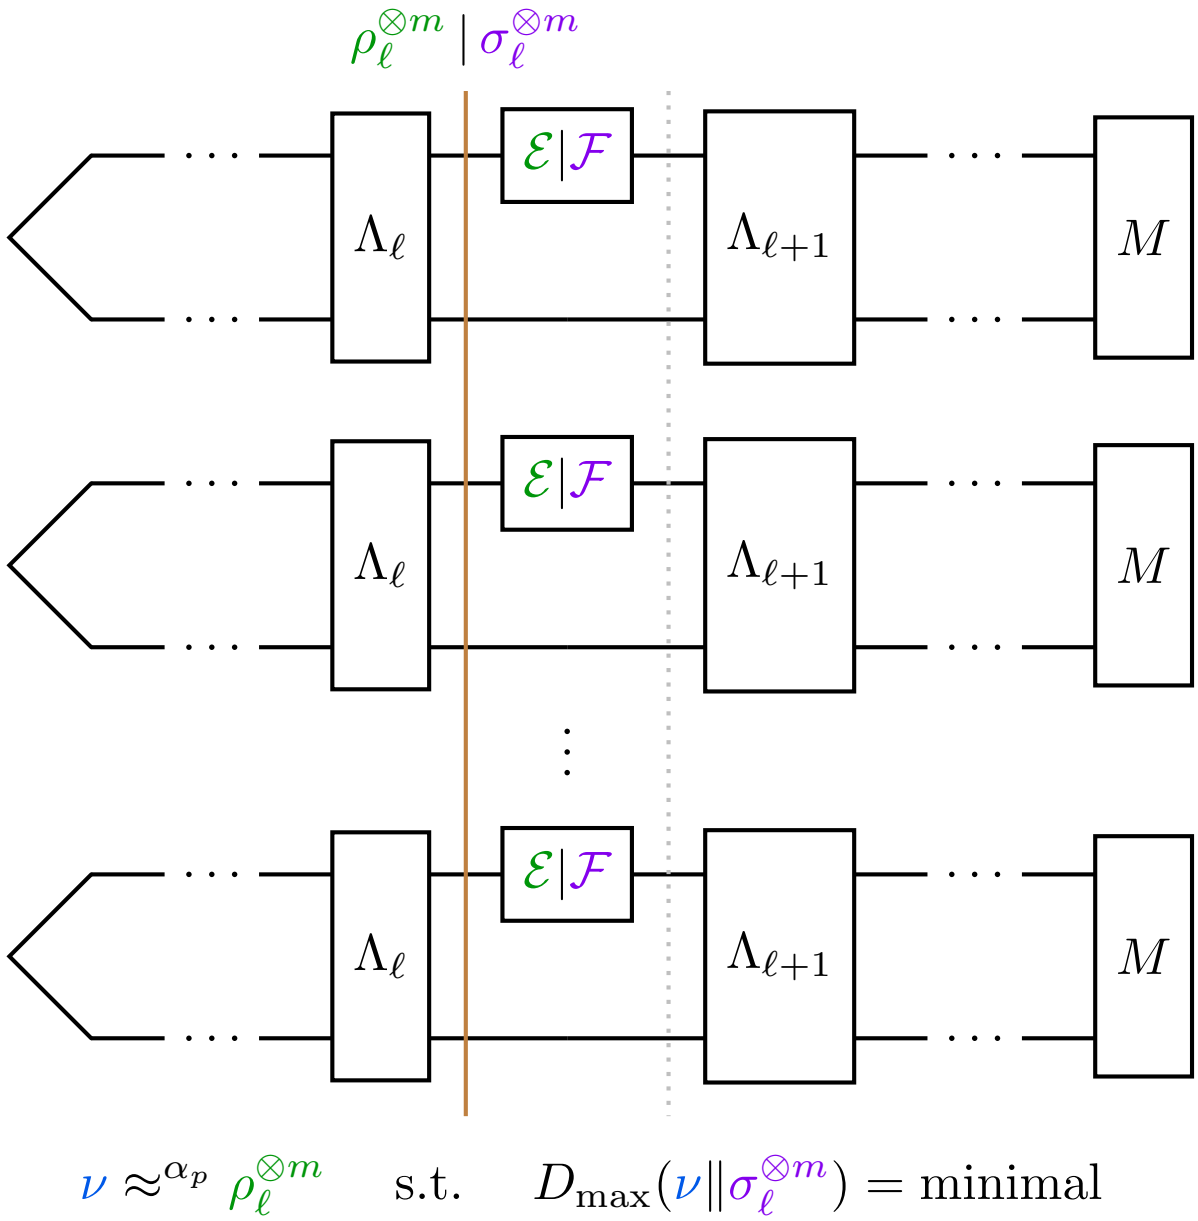

Figure 3.3: Illustration of a key step in our proof, the construction of the parallel input state. We start by picking a single step $\ell \in \{1, \ldots, n\}$ out of the adaptive strategy, where the distinguishability increase $D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell)$ is maximal (this corresponds to the step from the orange to the dotted grey line in the diagram). Now consider m copies of the adaptive strategy in parallel. We construct our parallel input state ν starting from m copies of the input state of the adaptive strategy at this step ℓ if the channel was $\mathcal{E}$ (this is $\rho_\ell^{\otimes m}$). The state ν is then smoothed a bit to reduce its distance to $\sigma_\ell^{\otimes m}$ (which is the input state that we would have if the channel was $\mathcal{F}$). The degree to which we smooth depends on the type I error α_p we want to achieve with the parallel strategy. Having a small type I error means that the state ν is very close to $\rho_\ell^{\otimes m}$, whereas allowing for a larger type I error will move the state closer to $\sigma_\ell^{\otimes m}$.

3.3.2 Proof of Theorem 3.3.4

Proof. We start by applying Lemma 2.4.1 to $D_H^{\alpha_a}(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n))$:

$$\frac{1}{n}D_H^{\alpha_a}(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) \leq \frac{1}{n}\frac{1}{1-\alpha_a}\big(D(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) + h(\alpha_a)\big)\,. \tag{3.3.32}$$

Note that a classical version of this equation in the context of channel discrimination was previously obtained in Hayashi (2009), Eq. (33). Note now that we can write

$$\begin{aligned}
D(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) &= D(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) - D(\rho_n\|\sigma_n) + D(\rho_n\|\sigma_n) && (3.3.33)\\
&= D(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) - D(\rho_n\|\sigma_n) + D(\Lambda_n(\mathcal{E}(\rho_{n-1}))\|\Lambda_n(\mathcal{F}(\sigma_{n-1}))) && (3.3.34)\\
&\leq D(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) - D(\rho_n\|\sigma_n) + D(\mathcal{E}(\rho_{n-1})\|\mathcal{F}(\sigma_{n-1})) && (3.3.35)\\
&\leq \ldots \leq \sum_{k=1}^{n}\Big[D(\mathcal{E}(\rho_k)\|\mathcal{F}(\sigma_k)) - D(\rho_k\|\sigma_k)\Big], && (3.3.36)
\end{aligned}$$

where we used the definition of ρ_k and σ_k, the data-processing inequality, and the fact that $\rho_1 = \sigma_1$. Let us use the index ℓ for the step in the adaptive protocol where this amortized difference is the largest, i.e.

$$\ell \coloneqq \operatorname*{arg\,max}_{k\in\{1,\ldots,n\}}\Big[D(\mathcal{E}(\rho_k)\|\mathcal{F}(\sigma_k)) - D(\rho_k\|\sigma_k)\Big]\,. \tag{3.3.37}$$

Then,

$$\frac{1}{n}D(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) \leq D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell)\,. \tag{3.3.38}$$

We can convert this to smoothed max-relative entropies by using Lemma 2.5.2. We will proceed with the bound requiring a condition on m (or n as it is called in Lemma 2.5.2); the m-independent bound is achieved in complete analogy by just taking the alternative statements from Lemma 2.5.2. We get:

$$\begin{aligned}
&D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell) \leq \frac{1}{m}\Big(D_{\max}^{\varepsilon_1}\big(\mathcal{E}^{\otimes m}(\rho_\ell^{\otimes m})\|\mathcal{F}^{\otimes m}(\sigma_\ell^{\otimes m})\big) - D_{\max}^{\varepsilon_2}\big(\rho_\ell^{\otimes m}\|\sigma_\ell^{\otimes m}\big)\Big)\\
&+\frac{1}{\sqrt{m}}\left[K_1 c_{\gamma_1}\big(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)\big)\sqrt{\log\left(\frac{1}{1-\varepsilon_1}\right)} + K_2 c_{\gamma_2}(\rho_\ell\|\sigma_\ell)\sqrt{\log\left(\frac{1}{\varepsilon_2}\right)}\right] + \frac{1}{m}\log\left(\frac{1}{1-\varepsilon_2^2}\right) \qquad (3.3.39)
\end{aligned}$$

where $\gamma_1, \gamma_2 \in (0,1]$ and $\varepsilon_1, \varepsilon_2 \in (0,1)$ are arbitrary. It will be very convenient to choose $1-\varepsilon_1 = \varepsilon_2 =: \varepsilon$, which is almost optimal as $K_1 \approx K_2$ and $c_\gamma\big(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)\big) \approx c_\gamma(\rho_\ell\|\sigma_\ell)$ for large ℓ (which is the regime we are most interested in). Then,

$$\begin{aligned}
&D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell) \leq \frac{1}{m}\Big(D_{\max}^{1-\varepsilon}\big(\mathcal{E}^{\otimes m}(\rho_\ell^{\otimes m})\|\mathcal{F}^{\otimes m}(\sigma_\ell^{\otimes m})\big) - D_{\max}^{\varepsilon}\big(\rho_\ell^{\otimes m}\|\sigma_\ell^{\otimes m}\big)\Big)\\
&\qquad + \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{1}{\varepsilon}\right)} + \frac{1}{m}\log\left(\frac{1}{1-\varepsilon^2}\right), \qquad (3.3.40)
\end{aligned}$$

where

$$c_\ell := \inf_{\gamma_1,\gamma_2\in(0,1]} \left[K_1 c_{\gamma_1}\left(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)\right) + K_2 c_{\gamma_2}(\rho_\ell\|\sigma_\ell)\right], \tag{3.3.41}$$

and the conditions on m from Lemma 2.5.2 will be satisfied if

$$m \geq \log\left(\frac{1}{\varepsilon}\right)\left(\frac{8}{\log(3)K_2}\right)^2. \tag{3.3.42}$$

Taking $\varepsilon \leq 1/2$, we can apply Lemma 3.3.7 to (3.3.40). We then get a state $\tilde{\nu} \in B^\circ_\varepsilon(\rho_\ell^{\otimes m})$ such that

$$D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell) \leq \frac{1}{m} D_{\max}^{1-2\varepsilon}\left(\mathcal{E}^{\otimes m}(\tilde{\nu})\|\mathcal{F}^{\otimes m}(\tilde{\nu})\right) + \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{1}{\varepsilon}\right)} + \frac{1}{m}\log\left(\frac{1}{1-\varepsilon^2}\right). \tag{3.3.43}$$

Note that $\tilde{\nu} = \tilde{\nu}_{R_a^m A^m} \in \mathcal{D}\left(R_a^m A^m\right)$ is the first of the two possible choices for ν given in Theorem 3.3.4. For the other one, let $\tilde{\nu}_{R' R_a^m A^m}$ be a purification of $\tilde{\nu}_{R_a^m A^m}$, and hence also a purification of $\tilde{\nu}_{A^m}$. As the smoothed max-divergence satisfies the data-processing inequality (see, e.g., Khatri and Wilde 2020, Prop. 4.60) we have

$$D_{\max}^{1-2\varepsilon}\left(\mathcal{E}^{\otimes m}(\tilde{\nu}_{R_a^m A^m})\|\mathcal{F}^{\otimes m}(\tilde{\nu}_{R_a^m A^m})\right) \leq D_{\max}^{1-2\varepsilon}\left(\mathcal{E}^{\otimes m}(\tilde{\nu}_{R' R_a^m A^m})\|\mathcal{F}^{\otimes m}(\tilde{\nu}_{R' R_a^m A^m})\right). \tag{3.3.44}$$

Now, let ν be the canonical purification of $\tilde{\nu}_{A^m}$, as specified in Theorem 3.3.4. Since all purifications are equivalent up to a partial isometry on the purifying system (see, e.g., Khatri and Wilde 2020) on which the channels $\mathcal{E}^{\otimes m}$ and $\mathcal{F}^{\otimes m}$ act as an identity, and since the smoothed max-divergence is invariant under isometries, we have

$$D_{\max}^{1-2\varepsilon}\left(\mathcal{E}^{\otimes m}(\tilde{\nu}_{R' R_a^m A^m})\|\mathcal{F}^{\otimes m}(\tilde{\nu}_{R' R_a^m A^m})\right) = D_{\max}^{1-2\varepsilon}\left(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\right). \tag{3.3.45}$$

We will proceed with ν as the chosen state, but note that everything works analogously by choosing $\tilde{\nu}$. We now further apply the upper bound from Lemma 2.5.1 to find:

$$\begin{aligned} D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell) \leq \frac{1}{m} D_H^{1-(1-2\varepsilon)^2}\left(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\right) + \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{1}{\varepsilon}\right)} \\ + \frac{1}{m}\left[\log\left(\frac{1}{1-\varepsilon^2}\right) + \log\left(\frac{1}{1-(1-2\varepsilon)^2}\right)\right]. \end{aligned} \tag{3.3.46}$$

To get our desired expression we set $\varepsilon := \frac{1}{2}(1-\sqrt{1-\alpha_p})$, which is equivalent to $\alpha_p = 1-(1-2\varepsilon)^2$. Using the estimates

$$\frac{1-\sqrt{1-\alpha_p}}{2} \geq \frac{\alpha_p}{4} \tag{3.3.47}$$

and

$$1-\varepsilon^2 = \frac{1+\sqrt{1-\alpha_p}}{2} + \frac{\alpha_p}{4} \geq 1 - \frac{\alpha_p}{4}, \tag{3.3.48}$$

we arrive at the expression

$$D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell) \leq \frac{1}{m} D_H^{\alpha_p}\left(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\right) + \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{4}{\alpha_p}\right)} + \frac{1}{m}\left[\log\left(\frac{1}{\alpha_p}\right) - \log\left(1-\frac{\alpha_p}{4}\right)\right], \quad (3.3.49)$$

which we can insert into (3.3.38) and then (3.3.32) to get

$$\frac{1}{n} D_H^{\alpha_a}(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) \leq \frac{1}{1-\alpha_a}\left[\frac{1}{m} D_H^{\alpha_p}(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)) + \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{4}{\alpha_p}\right)} + \frac{1}{m}\left(\log\left(\frac{1}{\alpha_p}\right) - \log\left(1-\frac{\alpha_p}{4}\right)\right) + \frac{h(\alpha_a)}{n}\right], \quad (3.3.50)$$

which leads to the desired statement in (3.3.12).

For the bounds on c_ℓ, we start with

$$\begin{aligned} c_\gamma(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) &= \frac{1}{\gamma}\log\left(2^{\gamma D_{1+\gamma}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell))} + 2^{-\gamma D_{1-\gamma}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell))} + 1\right) \\ &\leq \frac{1}{\gamma}\log\left(2^{\gamma \widehat{D}_{1+\gamma}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell))} + 2\right), \end{aligned} \quad (3.3.51)$$

where we used (2.3.85) for the inequality and the fact that $D_{1-\gamma}$ is positive on normalized states. Now, by repeated use of the chain rule for the geometric Rényi divergence (2.3.51) we get

$$\begin{aligned} \widehat{D}_{1+\gamma}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) &\leq \widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F}) + \widehat{D}_{1+\gamma}(\rho_\ell\|\sigma_\ell) && (3.3.52)\\ &= \widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F}) + \widehat{D}_{1+\gamma}(\Lambda_\ell(\mathcal{E}(\rho_{\ell-1}))\|\Lambda_\ell(\mathcal{F}(\sigma_{\ell-1}))) && (3.3.53)\\ &\leq \widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F}) + \widehat{D}_{1+\gamma}(\mathcal{E}(\rho_{\ell-1})\|\mathcal{F}(\sigma_{\ell-1})) && (3.3.54)\\ &\leq \ldots \leq \ell\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F}). && (3.3.55) \end{aligned}$$

Defining the corresponding channel quantity

$$\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) := \frac{1}{\gamma}\log\left(2^{\gamma\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\right), \quad (3.3.56)$$

we find that

$$c_\gamma(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) \leq \frac{1}{\gamma}\log\left(2^{\ell\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\right) \leq \frac{\ell}{\gamma}\log\left(2^{\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\right) = \ell\,\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}). \quad (3.3.57)$$

Using the same argument, we find that also $\widehat{D}_{1+\gamma}(\rho_\ell\|\sigma_\ell) \leq \ell\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})$ and hence also

$$c_\gamma(\rho_\ell\|\sigma_\ell) \leq \ell\,\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}). \quad (3.3.58)$$

Thus,

$$c_\ell \leq \ell(K_1 + K_2) \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \leq n(K_1 + K_2) \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F})\,, \tag{3.3.59}$$

concluding the proof. □

Proof of Corollary 3.3.1

Proof. Corollary 3.3.1 follows from (3.3.4) by making the following estimates

$$h(\alpha_a) \leq 1\,, \tag{3.3.60}$$

$$K \leq 1\,, \tag{3.3.61}$$

$$c'_\ell \leq \frac{8n}{\log(3)} \widehat{c}_1(\mathcal{E}\|\mathcal{F}) \approx 5.05\, n\, \widehat{c}_1(\mathcal{E}\|\mathcal{F}) \leq 6\, n\, \widehat{c}_1(\mathcal{E}\|\mathcal{F})\,, \tag{3.3.62}$$

$$-\log\left(1 - \frac{\alpha_p}{4}\right) \leq 1\,. \tag{3.3.63}$$

Combining all of these we can bound the error term by

$$\frac{6n}{\sqrt{m}} \widehat{c}_1(\mathcal{E}\|\mathcal{F}) \log\left(\frac{8}{\alpha_p}\right) + \frac{1}{m} \log\left(\frac{2}{\alpha_p}\right) + \frac{1}{n} \leq \frac{7n}{\sqrt{m}} \widehat{c}_1(\mathcal{E}\|\mathcal{F}) \log\left(\frac{8}{\alpha_p}\right) + \frac{1}{n}\,, \tag{3.3.64}$$

where we used that $\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \geq 1$ for all $\gamma \in (0, 1]$. Note finally that $\widehat{D}_2(\mathcal{E}\|\mathcal{F}) = D_2(\mathcal{E}\|\mathcal{F}) \leq D_{\max}(\mathcal{E}\|\mathcal{F})$, which leads to the desired expression of C in (3.3.1). □

3.3.3 An Explicit Example

In this section we provide an example that illustrates how adaptive and parallel strategies can differ in the finite-length regime, and how our result bounds this difference. Specifically, we construct an example where the number of channel uses required by a parallel strategy to match a specific adaptive strategy turns out to be arbitrarily large. This demonstrates that the relation between adaptive and parallel strategies in the finite-length regime is in general complex, and one cannot expect a substantially simpler relationship between adaptive and parallel strategies than what we obtain in Corollary 3.3.1 and Theorem 3.3.4.

Our example is inspired by the already mentioned example from Harrow et al. (2010) and Salek, Hayashi, and Winter (2022) that shows a separation between the asymptotic error decay rate of adaptive and parallel strategies in the *symmetric* setting. For a specific pair of channels $(\mathcal{E}, \mathcal{F})$, they construct an adaptive strategy with two channel uses that achieves perfect discrimination (i.e., type I and type II error are equal to zero), whereas for parallel strategies they upper bound the symmetric error exponent by a finite value (i.e., they upper bound the rate at which both errors can simultaneously go to zero); so in particular, perfect discrimination is never possible with parallel strategies. Looking at this example in the asymmetric setting, one finds that there exists an input state $\nu \in \mathcal{D}(A)$ such that for any arbitrary type I error α_p there exists an m such that $D_H^{\alpha_p}((\mathcal{E}(\nu))^{\otimes m}\|(\mathcal{F}(\nu))^{\otimes m}) = \infty$; i.e., already with a parallel strategy with product inputs one can achieve zero type II error with an arbitrary small type I error if one only makes m large enough. Hence, unlike the symmetric setting, in the asymmetric setting there is no asymptotic gap between adaptive and

parallel strategies, but there is still a significant difference for finite m, as the adaptive strategy achieves $\alpha_a = 0, \beta_a = 0$ with only two channel uses, whereas the parallel strategy requires a large m to achieve $\alpha_p \leq \varepsilon, \beta_p = 0$. As mentioned in Remark 3.3.6, the two channels $\mathcal{E}$ and $\mathcal{F}$ used in this example have $D_{\max}(\mathcal{E}\|\mathcal{F}) = \infty$ and hence Theorem 3.3.4 does not immediately apply. What we are going to do though, is use a slightly noisy version of this channel $\mathcal{F}$, which makes $D_{\max}(\mathcal{E}\|\mathcal{F})$ finite.

Define the channels $\mathcal{E}, \mathcal{F} : \mathcal{D}\left(\mathbb{C}^2 \otimes \mathbb{C}^2\right) \to \mathcal{D}\left(\mathbb{C}^2\right)$ for $\kappa \in [0,1]$ as follows:

$$\begin{aligned}\mathcal{E}(\rho \otimes \omega) &:= |0\rangle\langle 0| \langle 0|\omega|0\rangle + |0\rangle\langle 0| \langle 0|\rho|0\rangle \langle 1|\omega|1\rangle + \frac{\mathbb{1}}{2} \langle 11|\rho \otimes \omega|11\rangle \\ \mathcal{F}(\rho \otimes \omega) &:= (1-\kappa)\left[|+\rangle\langle +| \langle 0|\omega|0\rangle + |1\rangle\langle 1| \langle +|\rho|+\rangle \langle 1|\omega|1\rangle + \frac{\mathbb{1}}{2} \langle -1|\rho \otimes \omega\, |-1\rangle \right] + \kappa \frac{\mathbb{1}}{2}.\end{aligned} \tag{3.3.65}$$

It is easy to see that

$$\mathcal{E}\Big(\mathcal{E}\big(|00\rangle\langle 00|\big) \otimes |1\rangle\langle 1|\Big) = |0\rangle\langle 0|, \tag{3.3.66}$$

$$\mathcal{F}\Big(\mathcal{F}\big(|00\rangle\langle 00|\big) \otimes |1\rangle\langle 1|\Big) = (1-\delta(\kappa))\,|1\rangle\langle 1| + \delta(\kappa)\,|0\rangle\langle 0|\,. \tag{3.3.67}$$

where $\delta(\kappa) = (3\kappa - \kappa^2)/4$. For $\kappa = 0$ (which corresponds to the original example in Harrow et al. (2010)) we find $\delta(0) = 0$ and hence the above gives an adaptive strategy that makes the channels perfectly distinguishable with just two channel uses. The exact same strategy will become arbitrarily good if κ is nonzero but small, specifically

$$\frac{1}{2} D_H^0(\mathcal{E}(\rho_2)\|\mathcal{F}(\sigma_2)) = -\frac{1}{2}\log(\delta(\kappa)), \tag{3.3.68}$$

where the adaptive strategy has $\rho_1 = \sigma_1 = |00\rangle\langle 00|$ and $\rho_2 = |01\rangle\langle 01|$, $\sigma_2 = (1-\kappa/2)\,|+1\rangle\langle +1| + \kappa/2\,|-1\rangle\langle -1|$. We also find:

$$\begin{aligned} D(\mathcal{E}(\rho_1)\|\mathcal{F}(\sigma_1)) - D(\rho_1\|\sigma_1) &= D(\mathcal{E}(\rho_1)\|\mathcal{F}(\rho_1)) = D(|0\rangle\langle 0|\,\|(1-\kappa/2)\,|+\rangle\langle +| + \kappa/2\,|-\rangle\langle -|) \\ &= -\frac{1}{2}\left[\log\left(\frac{\kappa}{2}\right) + \log\left(1-\frac{\kappa}{2}\right)\right], \end{aligned} \tag{3.3.69}$$

$$D(\mathcal{E}(\rho_2)\|\mathcal{F}(\sigma_2)) - D(\rho_2\|\sigma_2) = \log(\delta(\kappa)) + \frac{1}{2}\left[\log\left(\frac{\kappa}{2}\right) + \log\left(1-\frac{\kappa}{2}\right)\right], \tag{3.3.70}$$

where interestingly (3.3.69) is always larger than (3.3.70), and so it is the first step that increases the distinguishability the most (when measured in terms of relative entropy). This also implies that this adaptive strategy is not asymptotically optimal, as

$$\frac{1}{2} D(\mathcal{E}(\rho_2)\|\mathcal{F}(\sigma_2)) < D(\mathcal{E}(\rho_1)\|\mathcal{F}(\rho_1))\,, \tag{3.3.71}$$

and hence it is asymptotically better to just use a parallel strategy with tensor copies of ρ_1 as an input state. Nevertheless, we will see that the adaptive strategy is still far superior in a regime where the number of channel uses m is not too large: It is well known (see, e.g., Khatri and Wilde 2020, Prop. 4.66) that for all

states ρ, σ

$$\lim_{\alpha_p \to 0} D_H^{\alpha_p}(\rho\|\sigma) = D_H^0(\rho\|\sigma) = -\log\big(\mathrm{Tr}\big(\rho^0\sigma\big)\big)\,. \tag{3.3.72}$$

In this specific example, for all input states $\nu \in \mathcal{D}\left(\mathcal{H}_R \otimes \mathbb{C}^2 \otimes \mathbb{C}^2\right)$, it can be shown that

$$D_H^0(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) \leq 2\,. \tag{3.3.73}$$

To see this, we start by observing that the second input system of our channels (previously called ω) can be treated as classical, and hence the maximum is attained at either $\omega = |0\rangle\langle 0|$ or $\omega = |1\rangle\langle 1|$ (this follows from joint convexity). If we choose the former, the channel output does not depend on the remaining input state and one finds

$$D_H^0(\mathcal{E}(\rho_{RA}\otimes|0\rangle\langle 0|)\|\mathcal{F}(\rho_{RA}\otimes|0\rangle\langle 0|)) = -\log \mathrm{Tr}\left(|0\rangle\langle 0|\left(|+\rangle\langle +|\,(1-\kappa/2)+|-\rangle\langle -|\,(\kappa/2)\right)\right) = 1\,. \tag{3.3.74}$$

For the second choice of ω one finds that $D_H^0(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = 0$ unless $\nu = \rho_R \otimes |0\rangle\langle 0| \otimes |1\rangle\langle 1|$ and then

$$D_H^0(\mathcal{E}(|0\rangle\langle 0|\otimes|1\rangle\langle 1|)\|\mathcal{F}(|0\rangle\langle 0|\otimes|1\rangle\langle 1|)) = -\log \mathrm{Tr}\left(|0\rangle\langle 0|\,(\mathbb{1}/4(1+\kappa))\right) \leq 2\,. \tag{3.3.75}$$

An analogous argument also immediately yields $D_H^0(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)) \leq 2m$ for any (potentially entangled) ν. Hence, for fixed m, the rate

$$\frac{1}{m} D_H^{\alpha_p}(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)) \tag{3.3.76}$$

of any parallel strategy can be brought down to 2 by making the type I error threshold α_p small enough, whereas the mentioned adaptive strategy achieves zero type I error and a type II error rate that becomes arbitrarily large as $\kappa \to 0$.

To actually calculate the performance of a parallel strategy where the input ρ_1 is used m times, we can use the second-order asymptotics of hypothesis testing relative entropy to relative entropy (this only works here because our input state is a product state). Define $\rho = \mathcal{E}(\rho_1)$, $\sigma = \mathcal{F}(\rho_1)$. Then K. Li (2014), Thm. 5 implies

$$\frac{1}{m} D_H^{\alpha_p}(\rho^{\otimes m}\|\sigma^{\otimes m}) \geq D(\rho\|\sigma) + \sqrt{\frac{V}{m}}\Phi^{-1}\left(\alpha_p - \frac{1}{\sqrt{m}}\frac{CT^3}{\sqrt{V^3}}\right), \tag{3.3.77}$$

$$\frac{1}{m} D_H^{\alpha_p}(\rho^{\otimes m}\|\sigma^{\otimes m}) \leq D(\rho\|\sigma) + \sqrt{\frac{V}{m}}\Phi^{-1}\left(\alpha_p + \frac{1}{\sqrt{m}}\left(\frac{CT^3}{\sqrt{V^3}} + 2\right)\right), \tag{3.3.78}$$

where Φ^{-1} is the inverse of the cumulative distribution function of the standard normal distribution, $C \leq 0.4784$, and

$$V \equiv V(\rho\|\sigma) \coloneqq \mathrm{Tr}\big[\rho(\log\rho - \log\sigma - D(\rho\|\sigma))^2\big], \tag{3.3.79}$$

$$T^3 \equiv T^3(\rho\|\sigma) \coloneqq \sum_{i,j} \lambda_i |\langle x_i|y_j\rangle|^2 |\log(\lambda_i) - \log(\mu_j) - D(\rho\|\sigma)|^3 \tag{3.3.80}$$

for spectral decompositions $\rho = \sum_i \lambda_i |x_i\rangle\langle x_i|$ and $\sigma = \sum_j \mu_j |y_j\rangle\langle y_j|$.

Note that, instead of comparing the type II error decay rate of a parallel strategy with m channel uses to the corresponding rate of an adaptive strategy with two channel uses, it might be more intuitive to compare the parallel strategy to a setup where the adaptive strategy with two channel uses is repeated $m/2$ times in parallel (so that the adaptive and parallel strategies have the same number of channel uses, and comparing rates is the same as comparing type II errors). For this example, since the type I error of the adaptive strategy was chosen to be zero, this is actually equivalent, as

$$\frac{1}{2k} D_H^0((\mathcal{E}(\rho_2))^{\otimes k} \| (\mathcal{F}(\sigma_2))^{\otimes k}) = \frac{1}{2} D_H^0(\mathcal{E}(\rho_2) \| \mathcal{F}(\sigma_2)) = -\frac{1}{2} \log(\delta(\kappa)). \tag{3.3.81}$$

Our theorem (Theorem 3.3.4) gives an upper bound on the extent to which all finite-length parallel strategies can have worse type II errors compared to the adaptive strategy, or equivalently how large m has to be chosen to achieve similar performances. For this example specifically, due to (3.3.69) and (3.3.70) we can choose $\ell = 1$ in Theorem 3.3.4, and hence also the parallel input state ν in Theorem 3.3.4 will just be $\nu = \rho_1^{\otimes m}$.

Figure 3.4 depicts our lower bound (from Theorem 3.3.4) on the performance of a parallel strategy for the given adaptive strategy, together with the actual performance of the parallel strategy choosing $\rho_1^{\otimes m}$ as the input state. For the figure we chose $\kappa = 2^{-50}$, $\alpha_a = 0$, and $\alpha_p = 2^{-5}$. The parameters are chosen to make the following features nicely visible simultaneously in one plot: (i) the range of m where the parallel strategy is worse, (ii) the range of m where it surpasses the adaptive strategy, and (iii) our bound. For c_ℓ in Theorem 3.3.4 we used (3.3.15) with a numerical optimization over γ_1. One finds that there is a range of values for m for which the adaptive strategy is better, and the parallel strategy is lower bounded fairly tightly by our bound. As the given adaptive strategy is not asymptotically optimal it is eventually surpassed by the parallel strategy.

3.4 Outlook

3.4.1 Potential Variants of Our Result and Strong Converse Property

In this section we start by discussing several pathways through which one could hope to improve or extend our results of this chapter, together with a discussion of obstacles we encountered on these pathways, and their relation to different open problems in quantum channel discrimination.

First of all, one might hope to be able to remove the factor of $1 - \alpha_a$ appearing in (3.3.12), perhaps at the cost of an additional error term proportional to $\log(1 - \alpha_a)$. If this additional error term decays in m and n (say, as long as α_a is bounded away from one), this would prove the strong converse property for quantum channel discrimination. To see this, suppose we could show something along the lines of the following inequality:

$$\frac{1}{m} D_H^{\alpha_p}(\mathcal{E}^{\otimes m}(\nu) \| \mathcal{F}^{\otimes m}(\nu)) \overset{?}{\geq} \frac{1}{n} D_H^{\alpha_a}(\mathcal{E}(\rho_n) \| \mathcal{F}(\sigma_n)) - \frac{C_1 n}{\sqrt{m}} \log\left(\frac{8}{\alpha_p}\right) - \frac{C_2}{n} \log\left(\frac{1}{1-\alpha_a}\right). \tag{3.4.1}$$

Now, fixing any $\alpha_a \in (0, 1)$ and taking limits (in this order) $m \to \infty, n \to \infty, \alpha_p \to 0$, we would find that

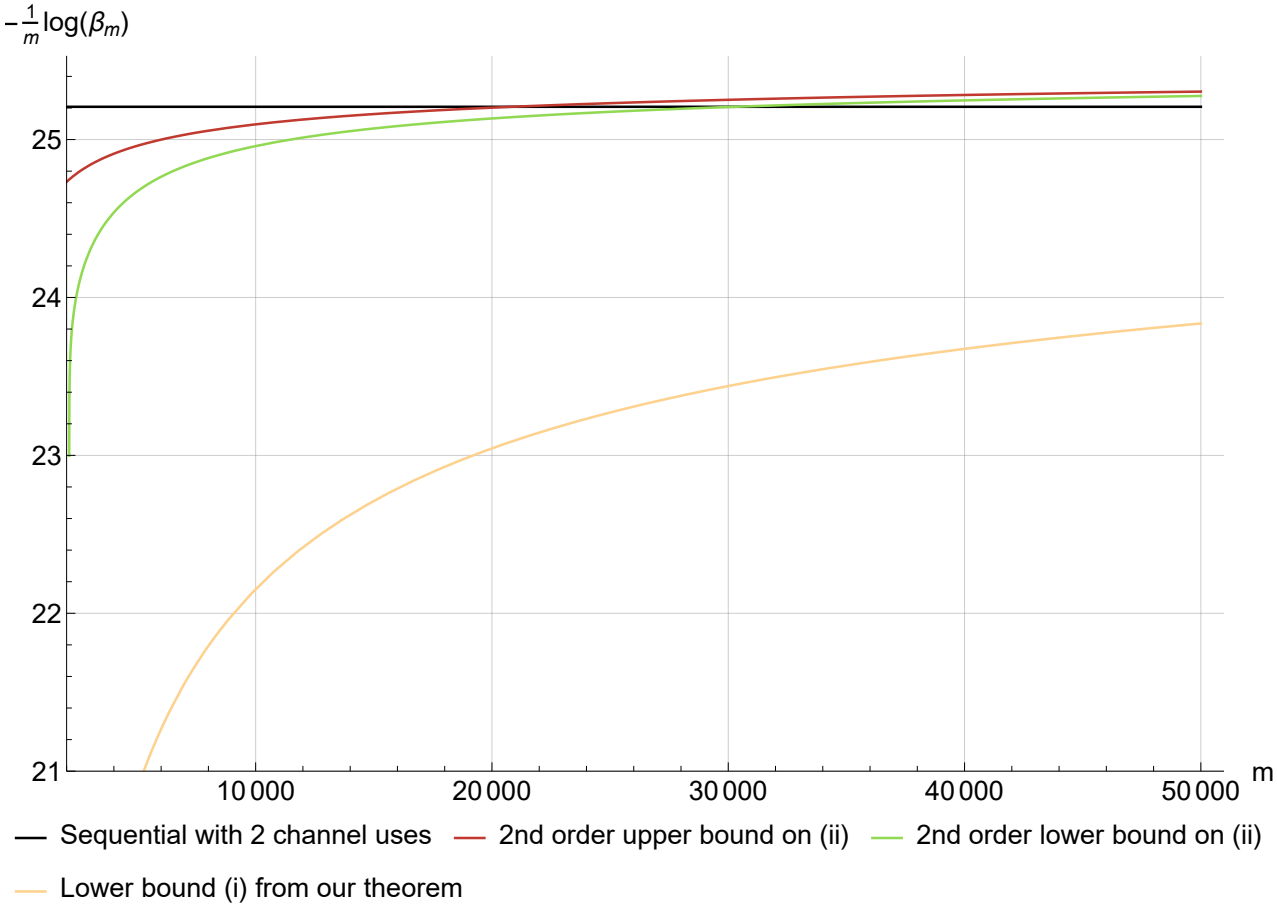

Figure 3.4: Illustration of the type II error decay rate per channel use of a simple adaptive and parallel strategy for a specific pair of channels (see (3.3.65) and (3.3.66) for the definitions of the channels $\mathcal{E}$ and $\mathcal{F}$, respectively, where $\kappa = 2^{-50}$). We compare a fixed adaptive strategy with two channel uses (constant black line) to (i) our lower bound on the performance of a parallel strategy (yellow line) and (ii) the actual performance of a parallel strategy (red and green lines), which are plotted as functions of the number of parallel channel uses m. The black line shows the value of (3.3.68), i.e., the type II error exponent for the given adaptive strategy with two channel uses and type I error $\alpha_a = 0$. This can alternatively be thought of as the rate of repeating the two-step adaptive strategy $m/2$ times in parallel. The yellow line shows the lower bound on the parallel strategy from our theorem (i.e., the right-hand side of (3.3.12)), choosing $\alpha_p = 2^{-5}$. For this specific example we can calculate the parallel input state ν of our theorem, and while we cannot explicitly find the optimal POVM and corresponding type II error (i.e., we cannot explicitly calculate the left-hand side of (3.3.12)), we can bound it from above and below using the second-order asymptotics of the hypothesis testing relative entropy, which is shown in the red and green lines, corresponding to the values of (3.3.78) and (3.3.77). We see that for small m there is a gap between the adaptive and parallel strategies; i.e., the adaptive strategy offers an advantage. This advantage disappears once m gets larger and in this specific example the chosen adaptive strategy even eventually gets surpassed by the parallel strategy, as the adaptive strategy turns out not to be asymptotically optimal.

for an arbitrary $\alpha_a \in (0,1)$

$$D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) = \lim_{\alpha_p \to 0} \lim_{n\to\infty} D_H^{\alpha_p}(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) \overset{?}{\geq} \limsup_{n\to\infty} \frac{1}{n} D_H^{\alpha_a}(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)), \tag{3.4.2}$$

i.e., allowing a finite type I error $\alpha_a \in (0,1)$ does not improve the best achievable type II error rate. This is precisely the strong converse property. Establishing the strong converse property for quantum channel discrimination is an important and very interesting open problem, which is still unsolved despite serious efforts. Significant progress has been made in H. Fawzi and O. Fawzi (2021), and another recent attempt was made in Fang, Gour, and X. Wang (2022); see also Berta et al. (2023).

One might hope to obtain such a variant of our bound without the factor $1-\alpha_a$, by, instead of transitioning from hypothesis testing entropy to relative entropy in (3.3.32), moving to an α-Rényi relative entropy instead and subsequently employing the known relations between Rényi relative entropies and smoothed max-relative entropies. It turns out that this will at some point require bounding the difference between Rényi relative entropies of order $1-\alpha$ and $1+\alpha$. This is possible using Lemma 2.3.1, however at the cost of an error term $(c_\gamma(\rho_n\|\sigma_n))^2$, which will scale quadratically in n. As this does not pick up any dependence in m, such an approach will not lead to anything useful in the asymptotic limit, and is hence unsuccessful.

One further interesting question deals with the relation between ℓ and n in Theorem 3.3.4. One might hope that an optimal adaptive strategy achieves the maximum possible distinguishability gain per channel use, i.e.

$$D^A(\mathcal{E}\|\mathcal{F}) = \sup_{\rho,\sigma\in\mathcal{D}(R\otimes A)} D(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) - D(\rho\|\sigma), \tag{3.4.3}$$

after some finite number of channel uses, or at least comes very close to it. If that was the case, one could bound ℓ by a constant and would remove the dependence on n in (3.3.1). However, note that the supremum in (3.4.3) goes over arbitrarily large reference systems R and hence does not need to be achieved, and we are not aware of any bounds on the system size R to get close to the optimum value. Alternatively, since $D^A(\mathcal{E}\|\mathcal{F}) = D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F})$ also a bound on the speed of convergence of the regularized channel divergence would do, which we are however also not aware of (H. Fawzi and O. Fawzi (2021) prove such a bound for the sandwiched Rényi divergence, which unfortunately cannot be extended straightforwardly to relative entropy). Hence, for now we have to assume that ℓ can in general be n, and that to adequately simulate an adaptive strategy with n channel uses, one might need to employ more than n^2 parallel channel uses to keep the error term in our bound small.

While our result becomes quite powerful in the asymmetric asymptotic setting, Theorem 3.3.4 can in principle be applied for any combination of type I and II errors. Hence, it could be interesting to apply it to scenarios where the type I error decays with m, say as $\alpha_p = 2^{-km}$ for some constant k. In this case, one would want to apply (3.3.12), to have at least a chance of the error term being bounded; however – depending on k – the corresponding condition (3.3.11) might not always be fulfilled. It would be interesting to see whether this condition on m could perhaps be relaxed, to allow for a wider range of k in this specific scenario.

As mentioned already in Remark 2.5.4, our bounds could be significantly tightened if we were able to employ second-order expansions, which however requires controlling the variance of the relative entropy

$V(\rho_n\|\sigma_n)$ in n, which we are unable to do. This seems to be again related to the strong converse problem, as second-order expansions for the hypothesis testing relative entropy imply the strong converse property. Already for the parallel case, if one was able to show

$$V(\mathcal{E}^{\otimes n}(\nu)\|\mathcal{F}^{\otimes n}(\nu)) = \mathcal{O}(n)\,, \tag{3.4.4}$$

where $\nu \in \mathcal{D}\left(R^{\otimes n} \otimes A^{\otimes n}\right)$, and $R \cong A$. is the optimal joint input state, this would be a very significant step towards proving the strong converse for quantum channel discrimination. Conversely, examples where this scales faster than linearly in n are quite likely to lend themselves to counterexamples for the strong converse property. We are not aware of any such examples: we leave this question open for further studies.

3.4.2 The Challenge of Treating Symmetric Channel Discrimination

This chapter has dealt exclusively with asymmetric quantum channel discrimination, however finding asymptotic error rates for the symmetric channel discrimination task is of course also interesting and relevant. While the example from Harrow et al. (2010) and Salek, Hayashi, and Winter (2022) demonstrates that the symmetric asymptotic error exponent cannot be equal for adaptive and parallel strategies, this is also essentially the only thing we know about asymptotic symmetric error exponents (unless we restrict to some special classes of channels, like classical-quantum channels). There currently aren't any kind of "entropic" expressions for one of either the adaptive and parallel symmetric asymptotic error exponents.

The natural conjecture would of course be that the parallel exponent is equal to the regularized Chernoff divergence (the Chernoff divergence is the asymptotic error rate for symmetric quantum state discrimination Audenaert et al. 2007, Nussbaum and Szkoła 2009), and the adaptive exponent equal to the amortized Chernoff divergence, however proving this has not been possible so far. The issue here is the non-i.i.d. structure of the problem once one chooses correlated input states. In the asymmetric setting, the reason we are able to deal with this non-i.i.d. structure is the converse bound from Lemma 2.4.1, which works for all (possibly entangled) states. On the achievability side, we do not have such a strong bound, but since for achievability we only have to show that one strategy exists that achieves the rate, we can choose blocked i.i.d. strategies (with a certain block size, that we can eventually send to infinity) to which we can apply the i.i.d. machinery already established.

In symmetric quantum hypothesis testing, the situation is reversed: The main result of Audenaert et al. (2007) gives very strong one-shot achievability bounds by (essentially) the asymptotic rate (the Chernoff divergence), even for arbitrarily correlated states, while then for the converse proof, one has to rely heavily on the i.i.d. structure to show that one cannot do better (Nussbaum and Szkoła 2009). Unfortunately now, we cannot apply a blocking argument to the converse direction of the proof, since this would only show that we cannot exceed the conjectured asymptotic rate with blocked strategies, but we have to show that we cannot exceed it with any strategy. This is where symmetric quantum channel discrimination is currently stuck, and it looks like it requires a significant jump in understanding of (maybe even classical) converse bounds for non-i.i.d. symmetric hypothesis testing to solve this.

Some progress was made recently, in using the geometric Rényi divergence to obtain a converse upper

bound that is somewhat closer to the conjectured asymptotic rate (Ji et al. 2024), but still known to be larger in general.

3.4.3 Strategies Beyond Adaptive - Non-causal Strategies and the Quantum Switch

In this thesis we consider adaptive strategies as the most general channel discrimination setups, and for practical applications where the channels are given as black-boxes, that is indeed correct. However, in principle one could think of more general strategies, which replace the causal iterative structure of an adaptive strategy with an arbitrary super-channel taking in the n black-boxes as an input, and which in principle do not have to obey to any causal structure (i.e. channel inputs can be influenced by any other channel output), see e.g. Chiribella (2012), Araújo et al. (2015), Chiribella and Ebler (2016), and Wechs et al. (2021). Such non-causal structures can sometimes appear in certain models of quantum gravity, but can also be implemented probabilistically (i.e. with a certain chance of failure) on "normal" causal devices (Chiribella et al. 2013), or sometimes be implemented in a way where the channels are selected (or one could say programmed) coherently (i.e. channels are selected or specified through a quantum register, which may be in a superposistion state, Oreshkov 2019). It is known that such non-causal strategies (assuming they can be implemented perfectly) can lead to smaller error probabilities in channel discrimination tasks with finite number of channel uses (Bavaresco, Murao, and Quintino 2021), but nothing is known yet about the asymptotic error exponent. The natural conjecture would be that for the asymmetric error exponent, also the most general strategies do not lead to any asymptotic advantage, however we have not been able to prove this so far, and leave this question open to further work.

4 Composite Classical and Quantum Channel Discrimination

While the last chapter dealt with binary quantum channel discrimination with simple hypotheses (we discriminate between two single channels), this chapter looks at composite quantum channel discrimination, where we are no longer guaranteed to only encounter two possible channels, but channels which come from two sets, and the task is to determine the set the channel came from. Arguably, this is much more practically relevant than the simple case, as this first of all includes the noisy regime, where the previously exactly specified two channel hypotheses are now turned into small sets of channels due to noise, but also much more general settings, i.e. questions of discriminating big sets with a certain structure (for states this could for example be sets of separable (Brandão et al. 2020) or coherent (Berta, Brandao, and Hirche 2021) states).

We will start with introducing the problem for discriminating two sets of states, and give an overview of previous results in the literature, before moving on to the problem of discriminating two sets of channels. To our knowledge, the task of composite binary quantum *channel* discrimination has not been studied at all thus far (even in the classical literature). Throughout our analysis, we will not restrict ourselves to the composite i.i.d. setting, i.e. we will also allow the provided objects (states or channels) to vary within the sets corresponding to the hypotheses.

What we show is the following: For binary asymmetric composite channel discrimination in this fairly general setting: (i) a characterization of the Stein exponent for parallel channel discrimination strategies (Theorem 4.3.1), and (ii) an upper bound on the Stein exponent for adaptive channel discrimination strategies (Proposition 4.4.1). We further show that already classically this upper bound can sometimes be achieved and be strictly larger than what is possible with parallel strategies (Example 4.4.3), and hence there can be an advantage of adaptive channel discrimination strategies with composite hypotheses. We go on to show that classically this advantage can only exist if the sets of channels corresponding to the hypotheses are non-convex, and additionally assuming this convexity makes parallel strategies asymptotically optimal (Theorem 4.4.4). We leave the question open whether an adaptive advantage can exist in the quantum case when the sets of channels are convex. Table 4.1 gives an overview of what we are able to show regarding composite channel discrimination, and illustrates in which cases an adaptive advantage exists.

As a consequence of our more general treatment which is not limited to the composite i.i.d. setting we also obtain a generalization of the composite state discrimination results of Berta, Brandao, and Hirche (2021) (Theorem 4.1.6). Note, however, that while we do not require provided states or channels to be identical, we still require them to be independent. Hence, our theorems do not directly imply or relate to the (in)famous generalized Stein's lemma (Brandao and Plenio 2010, Berta et al. 2023, Hayashi and Yamasaki 2024, Lami 2024), although one might be able to drop the independence requirement and thus extend our results using the new techniques employed in the very recent proofs of the generalized Stein's lemma (Hayashi and Yamasaki 2024, Lami 2024).

Summary of Main Results

Hypotheses	Asymptotic Parallel Exponent		Asymptotic Adaptive Exponent		Upper Bound	Shown in
Quantum Simple	$D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F})$	$=$	$D_A(\mathcal{E}\|\mathcal{F})$			Sec. 3.1.1
Classical Composite Convex Sets	$\max_{\nu} \min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}(\nu)\|\mathcal{F}(\nu))$	$=$	$\min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}\|\mathcal{F})$	$=$	$\min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}\|\mathcal{F})$	Thm. 4.4.4
Classical Composite Finite Sets	$\max_{\nu} \min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}(\nu)\|\mathcal{F}(\nu))$	$<$ i.g.	?	$\leq$	$\min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}\|\mathcal{F})$	Prop. 4.3.2 Ex. 4.4.3 Prop. 4.4.1
Quantum Composite	$\lim_{n\to\infty} \frac{1}{n} \min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}} D(\mathcal{E}_n\|\mathcal{F}_n)$	$<$ i.g.	?	$<$ i.g.	$\min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D^A(\mathcal{E}\|\mathcal{F})$	Thm. 4.3.1 Ex. 4.4.3 Prop. 4.4.1 Rem. 4.4.2

Table 4.1: Illustration of the relation between adaptive and parallel type II error exponents for various channel discrimination tasks. For the composite problems the task is to discriminate between two sets of channels $\mathcal{S}$ and $\mathcal{T}$ and the table also includes an upper bound based on the worst-case simple i.i.d. problem. "Quantum Simple" refers to the quantum channel discrimination problem with simple hypotheses. With "Classical" we mean that all channels are classical, and "Convex Sets" or "Finite Sets" refers to whether the sets of channels $\mathcal{S}$ and $\mathcal{T}$ are convex or finite. Please see the respective theorems for a general formulation of the results and a precise definition of the quantities involved; $\mathcal{C}$ denotes the convex hull. We write i.g. to denote that these inequalities will be strict in general, although there exist specific examples where equality holds.

4.1 Composite State Discrimination

In simple quantum state discrimination (as discussed in section Section 2.4.2), we are given n identical copies of an unknown state which is promised to be either ρ or σ, and the task is then to decide which of the two options it is. In composite quantum state discrimination however, we are only promised that the states are all from one of two sets S or T, and the task is now to only decide which set they come from (but not to further identify which state exactly was provided). Since there are now multiple states for each hypothesis, there are multiple possible scenarios how the n input states one receives are related: We could still be given n identical copies of a state, or alternatively, we could be given n completely different states but all from the same set S or T, or something in between, where the states are non-identical but still related. We would like to cover all these different scenarios in our analysis, and hence we will describe composite hypotheses as sequences of sets S_n which include all the possible combinations of n states we could get. We will make some small assumptions on these sets:

Definition 4.1.1. For the purpose of this work, a composite quantum state hypothesis (in the asymptotic setting) is a sequence of sets of states

$$\boldsymbol{S} = (S_n \subset \mathcal{D}\left(\mathcal{H}^{\otimes n}\right))_n$$

such that

1. Each set S_n is topologically closed.
2. Each element $\rho_n \in S_n$ is a tensor product of states $\rho_n = \rho^{(1)} \otimes \ldots \otimes \rho^{(n)}$, with each $\rho^{(i)} \in \mathcal{D}\left(\mathcal{H}\right)$ for $i = 1, ..., n$.
3. The sets S_n are closed under tracing out any subsystem, i.e. for any $i = 1, ..., n$ and $\rho_n \in \mathcal{S}_n$ we have that $\mathrm{Tr}_i(\rho_n) \in \mathcal{S}_{n-1}$, where Tr_i denotes the partial trace over the i^{th} subsystem.
4. Each set S_n is closed under permutation of the n subsystems, i.e. for any permutation $\pi \in \mathfrak{S}_n$ and associated canonical unitary representation $P_{\mathcal{H}}(\pi)$, we have for all $\rho_n \in S_n$: $P_{\mathcal{H}}(\pi)\rho_n P_{\mathcal{H}}(\pi)^\dagger \in \mathcal{S}_n$.

Interesting examples of this include:

1. The composite i.i.d. case: We have two sets $S, T \subset \mathcal{D}\left(\mathcal{H}\right)$, and are given n identical copies of an element from S if the null hypothesis is true, and n identical copies of an element from T if the alternate hypothesis is true. This corresponds to:

$$S_n \coloneqq \left\{ \rho^{\otimes n} \;\middle|\; \rho \in S \right\}, \tag{4.1.1}$$
$$T_n \coloneqq \left\{ \sigma^{\otimes n} \;\middle|\; \sigma \in T \right\}. \tag{4.1.2}$$

2. The arbitrarily varying case: This is similar to the composite i.i.d. case, but we are not given n identical copies, but n (potentially different) elements from S or T. This corresponds to:

$$S_n \coloneqq \left\{ \rho_1 \otimes \ldots \otimes \rho_n \mid \rho_1, \ldots, \rho_n \in S \right\}, \tag{4.1.3}$$
$$T_n \coloneqq \left\{ \sigma_1 \otimes \ldots \otimes \sigma_n \mid \sigma_1, \ldots, \sigma_n \in T \right\}. \tag{4.1.4}$$

3. The slightly-varying case: This is an example of a scenario that lies in between the arbitrarily varying case (where there is no correlation between the samples, except for them all being in the same set) and the composite i.i.d. case (where there is maximal correlation between the samples, as they are all identical). For any given $\varepsilon \in [0, 1]$ (which might depend on n) and any distance function $d : \mathcal{D}\left(\mathcal{H}\right) \times \mathcal{D}\left(\mathcal{H}\right) \to [0, 1]$ (e.g. trace distance or purified distance) set

$$S_n \coloneqq \left\{ \rho_1 \otimes \ldots \otimes \rho_n \mid \rho_1, \ldots, \rho_n \in S, \quad d(\rho_i, \rho_j) \leq \varepsilon \; \forall i, j \right\}, \tag{4.1.5}$$
$$T_n \coloneqq \left\{ \sigma_1 \otimes \ldots \otimes, \sigma_n \mid \sigma_1, \ldots, \sigma_n \in T, \quad d(\sigma_i, \sigma_j) \leq \varepsilon \; \forall i, j \right\}. \tag{4.1.6}$$

4. The simple i.i.d. case: The simple i.i.d. case can be seen as a special case of the above where S and T each contain one element.

Lemma 4.1.2. *If $\boldsymbol{S}$ is a quantum state hypothesis, then performing a measurement on any (joint) k subsystems of a state $\rho_n \in S_n$, and conditioning on the measurement result, yields a state ρ_{n-k} on the remaining subsystems that is an element of S_{n-k}. Precisely, let $k \in \{1, ..., n\}$, $\rho_n \in S_n \subset \mathcal{D}\left(\mathcal{H}^{\otimes n}\right)$ and $0 \leq M \leq \mathbb{1} \in \mathcal{B}\left(\mathcal{H}^{\otimes k}\right)$. If*

$$\omega_{n-k} \coloneqq \mathrm{Tr}_{1,...,k}\left[(M \otimes \mathbb{1}_{\mathcal{H}_{k+1}} \otimes ... \otimes \mathbb{1}_{\mathcal{H}_n})\rho_n\right] \tag{4.1.7}$$

then either $\mathrm{Tr}(\omega_{n-k}) = 0$ or $\omega_{n-k}/\mathrm{Tr}(\omega_{n-k}) \in S_{n-k}$.

Proof. This follows immediately from the fact that each element $\rho_n \in S_n$ is a tensor product, and that removing an element in the tensor product gives an element in S_{n-1}. □

Remark 4.1.3. Note that while we state the results below with the assumptions of Definition 4.1.1, we could replace the required tensor product structure (point 2 in Definition 4.1.1) with the statement of Lemma 4.1.2. While this might be more general, we find the assumptions of Definition 4.1.1 to be more natural. Similarly, later on when talking about hypotheses in the context of channel discrimination, we could replace the tensor product structure for channels (point 2 in Definition 4.2.1) with the statement of Lemma 4.1.2 for any tensor product input state.

For our discrimination problem, given an n and an unknown state in $\mathcal{D}\left(\mathcal{H}^{\otimes n}\right)$ we will still want to perform a binary POVM (fully specified by one of its elements, which we write as M) to decide between the two hypotheses. In the end we want to avoid making an error, i.e. claiming that our state comes from S_n when it actually comes from T_n and vice-versa, and these two error probabilities are again what we call type-I and type-II error now in this setting (see below for a formal definition). If we settle on a measurement M, the probability of making an error might still depend on which particular state from either S_n or T_n we actually end up getting. Here, we will be focussing on minimizing the worst case errors, i.e. we want to choose measurements which minimize the error uniformly over all states from S_n and T_n. More formally, we define the type I and type II error probabilities (also again called type I and type II errors) as:

$$\alpha(M, S_n) = \sup_{\rho \in S_n} \mathrm{Tr}((\mathbb{1} - M)\rho) \tag{4.1.8}$$

$$\beta(M, T_n) = \sup_{\sigma \in T_n} \mathrm{Tr}(M\sigma)\,. \tag{4.1.9}$$

Similarly to the rest of this thesis, we will be focussing on the asymmetric setting, where we want to minimize the type II error, β, under the constraint that the type I error α is below a certain threshold, and so the main quantity of interest is again the negative logarithm of this minimal type II error under the type I error constraint, which we also call the hypothesis testing relative entropy of the two sets S_n and T_n:

$$D_H^\varepsilon(S_n \| T_n) \coloneqq -\inf_{\substack{0 \leq M \leq \mathbb{1} \\ \alpha(M, S_n) \leq \varepsilon}} \log \beta(M, T_n)\,. \tag{4.1.10}$$

As the expression in (4.1.8) is linear in ρ, it is easy to see that

$$\alpha(M, \mathcal{C}(S_n)) = \alpha(M, S_n) \tag{4.1.11}$$

(where $\mathcal{C}$ is the convex hull) as the supremum will be achieved at an extremal point, and the same holds also for β.

Hence,

$$D_H^\varepsilon(S_n\|T_n) = D_H^\varepsilon(S_n\|\mathcal{C}(T_n)) = D_H^\varepsilon(\mathcal{C}(S_n)\|T_n) = D_H^\varepsilon(\mathcal{C}(S_n)\|\mathcal{C}(T_n)) \tag{4.1.12}$$

and hence the discrimination task considered is equivalent to discriminating between these convex hulls of the sets S_n and T_n. We are interested in the quantum Stein exponent for this discrimination task, i.e. the optimal exponential decay rate of the type II error in the limit $n \to \infty$ (which corresponds to infinitely many states being provided). Morally, the expression we are interested in is

$$\lim_{\varepsilon\to 0} \text{“}\lim_{n\to\infty}\text{”} \frac{1}{n} D_H^\varepsilon(S_n\|T_n)\,. \tag{4.1.13}$$

However, due to the additional optimization over the elements from the sets S_n and T_n, for any fixed $\varepsilon > 0$, $D_H^\varepsilon(S_n\|T_n)$ need no longer be super-additive in n, and hence we are not aware of any way to show that the limit $n \to \infty$ exists in these expressions. However, in many cases we are able to show that after additionally taking the limit $\varepsilon \to 0$ outside, the final expression does not depend on whether one takes a $\liminf$ or a $\limsup$ inside. Hence, we introduce the following notation

Definition 4.1.4. For any function $f(n,\varepsilon) : \mathbb{N} \times (0,1) \to \mathbb{R}$, and $A \in \mathbb{R} \cup \{-\infty, \infty\}$, we write

$$\lim_{\varepsilon\to 0} \overline{\lim_{n\to\infty}}\, f(n,\varepsilon) = A \tag{4.1.14}$$

if

$$\lim_{\varepsilon\to 0} \liminf_{n\to\infty} f(n,\varepsilon) = A \tag{4.1.15}$$

and

$$\lim_{\varepsilon\to 0} \limsup_{n\to\infty} f(n,\varepsilon) = A\,. \tag{4.1.16}$$

To come back to (4.1.13), in Berta, Brandao, and Hirche (2021) this problem was studied specifically in the composite i.i.d. case, i.e. with $S_n = \{\, \rho^{\otimes n} \mid \rho \in S \,\}$, $T_n = \{\, \sigma^{\otimes n} \mid \sigma \in T \,\}$, where $S, T \subset \mathcal{D}(\mathcal{H})$ were also assumed to be closed and convex, leading to:

Theorem 4.1.5 (Berta, Brandao, and Hirche 2021)**.** *Let S, T be closed and convex, and define for all n: $S_n := \{\, \rho^{\otimes n} \mid \rho \in S \,\}$, $T_n := \{\, \sigma^{\otimes n} \mid \sigma \in T \,\}$. Then*

$$\lim_{\varepsilon\to 0} \overline{\lim_{n\to\infty}}\, \frac{1}{n} D_H^\varepsilon(S_n\|T_n) = \lim_{n\to\infty} \frac{1}{n} \inf_{\substack{\rho_n\in\mathcal{S}_n \\ \sigma_n\in\mathcal{C}(T_n)}} D(\rho_n\|\sigma_n)\,, \tag{4.1.17}$$

where $\overline{\lim}$ can be $\liminf$ or $\limsup$ (see Definition 4.1.4), and one can find cases where this is strictly smaller than

$$\inf_{\substack{\rho\in S \\ \sigma\in T}} D(\rho\|\sigma)\,. \tag{4.1.18}$$

Remember that $\mathcal{C}$ stands for the convex hull, and it is precisely this convex hull in the infimum on the right-hand side of (4.1.17) which prevents the regularization from collapsing, as the elements of S_n and T_n

are tensor products, and the relative entropy is additive. Without this convex hull, the exponent (4.1.17) would be exactly equal to the single-letter expression (4.1.18), which we call the worst-case i.i.d. exponent, as it is equal to the exponent of the worst-case simple i.i.d. problem. Intuitively, one pays for the compositeness by having to include convex combinations in the Stein exponent, and this makes the discrimination problem strictly harder in some cases.

As a consequence of our channel discrimination result further below (Theorem 4.3.1), we will arrive at this following generalization of Theorem 4.1.5:

Theorem 4.1.6. *Let $\boldsymbol{S} = (S_n)_n, \boldsymbol{T} = (T_n)_n$ be two composite quantum state hypotheses. Then*

$$\lim_{\varepsilon\to 0} \overline{\lim_{n\to\infty}} \frac{1}{n} D_H^\varepsilon(S_n\|T_n) = \lim_{n\to\infty} \frac{1}{n} \min_{\substack{\rho_n\in\mathcal{C}(S_n)\\ \sigma_n\in\mathcal{C}(T_n)}} D(\rho_n\|\sigma_n)\,, \tag{4.1.19}$$

where $\overline{\lim}$ can be $\liminf$ or $\limsup$ (see Definition 4.1.4). Furthermore, if each S_n lies in the intersection of $\mathcal{D}(\mathcal{H}^{\otimes n})$ with a linear subspace of $\mathcal{B}(\mathcal{H}^{\otimes n})$ with dimension polynomial in n (this holds for example in the composite i.i.d. case, where each $\rho_n \in S_n$ is permutation invariant), then we can remove the first convex hull to get

$$\lim_{\varepsilon\to 0} \overline{\lim_{n\to\infty}} \frac{1}{n} D_H^\varepsilon(S_n\|T_n) = \lim_{n\to\infty} \frac{1}{n} \min_{\substack{\rho_n\in\mathcal{S}_n\\ \sigma_n\in\mathcal{C}(T_n)}} D(\rho_n\|\sigma_n)\,. \tag{4.1.20}$$

Proof. This follows as a special case of Theorem 4.3.1 below. □

Remark 4.1.7. Our Theorem 4.1.6 generalizes the previous Theorem 4.1.5 in multiple ways: Already in the composite i.i.d. setting it no longer requires the sets $\mathcal{S}$ and $\mathcal{T}$ to be convex. Additionally our theorem also includes all the non-i.i.d. cases such as the arbitrarily (or slightly varying) cases defined above.

4.1.1 Classical Adversarial Hypothesis Testing

Similar to Brandão et al. (2020) and Berta, Brandao, and Hirche (2021), our results are based on a reduction to a classical problem, the one of adversarial hypothesis testing. The following is a brief recapitulation of the treatment of adversarial hypothesis testing in Brandão et al. (2020). Let $P, Q \subset \mathbb{R}^\Omega$ (for a finite domain Ω) be two sets of probability distributions. In the typical composite i.i.d. setting, we are presented with n samples from a distribution in P or Q and have to make a decision which set the distribution comes from. In the adversarial setting, the adversary is allowed to change the distribution within P or Q for each sample, and he can make this change based on the samples we observed previously. Note that while the adversary has access to the previous samples, he can only select a probability distribution $p \in P$ or $q \in Q$ (depending on which hypothesis is true) for the next sample, but he cannot select the sample outcome itself. The adversary is fully specified by two sets of functions $\hat{p}_k : \Omega^{k-1} \to P$ and $\hat{q}_k : \Omega^{k-1} \to Q$, which for each k specify how the adversary picks the next probability distribution based on the previous $k-1$ sample outcomes. The two hypotheses then correspond to whether the adversary uses $\hat{p}_k$, and hence always chooses a probability distribution in P, or $\hat{q}_k$ and always chooses a probability distribution in Q. If the null hypothesis is true (i.e.

the adversary uses $\hat{p}_k$), the probability of a sample string $\mathbf{x} \in \Omega^n$ is then given by

$$\hat{p}(\mathbf{x}) \coloneqq \prod_{k=1}^{n} \hat{p}_k(x_1, \ldots, x_{k-1})(x_k)\,, \tag{4.1.21}$$

and we define $\hat{q}(\mathbf{x})$ in a similar manner. For any decision region $A_n \subset \Omega^n$, the type I and type II errors are then going to be the worst-case errors over all adversarial strategies. We define the corresponding n-shot error exponent as

$$D^{\varepsilon}_{\mathrm{adv},n}(P\|Q) = -\log \inf \left\{ \sup_{\hat{q}} \hat{q}(A_n) \;\middle|\; A_n \subset \Omega^n,\ \sup_{\hat{p}} \hat{p}(A_n^c) \leq \varepsilon \right\} \tag{4.1.22}$$

The key statement of Brandão et al. (2020) is that if the sets P and Q are closed and convex, adversarial hypothesis testing is asymptotically no harder than the worst-case i.i.d. setting, specifically:

Theorem 4.1.8 (Brandão et al. 2020, Theorem 2)**.** *Let Ω be a finite domain and $P, Q \subset \mathbb{R}^{\Omega}$ be two closed, convex sets of probability distributions. Then, for any $\varepsilon \in (0,1)$:*

$$\lim_{n\to\infty} \frac{1}{n} D^{\varepsilon}_{\mathrm{adv},n}(P\|Q) = \min_{p\in P, q\in Q} D(p\|q)\,. \tag{4.1.23}$$

Note that since we are taking the supremum over all adversaries, by picking an adversary that deterministically picks states in a certain sequence, this result implies that also any composite problem is classically asymptotically equally as hard as the worst-case i.i.d. problem (it is also easy to see that the composite problem cannot be simpler than the worst-case i.i.d. problem).

4.2 Composite Channel Discrimination

The task of composite channel discrimination is defined in a very analogous way to composite state discrimination: Given an unknown quantum channel as a black box and the side information that it comes from two sets of possible channels, the task is again to determine the set (but not necessarily the exact identity) of the channel.

The strategies one employs to come to a decision for this task are exactly the same as the ones we described for simple channel discrimination in Section 3.1: One has to find input states for the n channels given, either in a parallel or adaptive way, and then finally apply a binary POVM on the state one is left with. Similarly to the previous section on composite state discrimination, also for channels we will be looking at worst-case errors, and so the key difference to simple channel discrimination is that one has to make sure that the choice of input states and final measurement work well for all possible sequences of channels one could receive. In particular, also for channels we want to allow for the case where the channels we receive can vary from within the set.

Similar to the previous chapter, we will again mostly focus on comparing adaptive and parallel strategies, and in particular how the asymptotic error exponents compare. This turns out to be surprisingly complex,

as we will see that the asymptotic equivalence of adaptive and parallel strategies (when considering the asymmetric error exponent) does no longer always hold when considering composite hypotheses, already classically.

Specifically, in this section we will study the following: 1. We start with a treatment of parallel channel discrimination strategies, where we provide matching achievability and converse bounds for the Stein exponent in terms of a regularized expression (Theorem 4.3.1), in analogy to what has previously been shown (Berta, Brandao, and Hirche 2021) for state discrimination (i.e. Theorem 4.1.5). 2. We prove an upper bound on the Stein exponent for adaptive strategies (Proposition 4.4.1), where we show that this upper bound can sometimes but not always be achieved, and can also be larger than the parallel exponent (Example 4.4.3), hence demonstrating that adaptive strategies can sometimes be advantageous (we show this even classically). 3. We show that classically, under an additional convexity assumption which was not satisfied in the previous example, parallel and adaptive strategies are asymptotically equivalent in the asymmetric composite setting, and the Stein exponent is given by a single-letter entropic formula (Theorem 4.4.4). 4. We further show classically, and in some further restricted setting, that if we replace the convexity assumption with a finiteness assumption, we can still get a single-letter entropic expression for the Stein exponent for parallel strategies (Proposition 4.3.2).

Following the above discussion for composite state discrimination, we want to apply a similar level of generality to discriminating channels, where we want to allow the n black-boxes not be identical. Hence, in analogy with Definition 4.1.1 we will work with general hypotheses satisfying the following conditions:

Definition 4.2.1. For the purpose of this work, a composite quantum channel hypothesis (in the asymptotic setting) is a sequence of sets of channels

$$\boldsymbol{\mathcal{S}} = (\mathcal{S}_n \subset \mathrm{CPTP}(A^n \to B^n))_n$$

such that

1. Each set $\mathcal{S}_n$ is topologically closed.
2. Each element $\mathcal{E}_n \in \mathcal{S}_n$ is a tensor product of channels $\mathcal{E}_n = \mathcal{E}^{(1)} \otimes \ldots \otimes \mathcal{E}^{(n)}$, with $\mathcal{E}^{(i)} \in \mathrm{CPTP}(A \to B)$, for $i = 1, ..., n$.
3. For every $\mathcal{E}_n = \mathcal{E}^{(1)} \otimes \ldots \otimes \mathcal{E}^{(n)} \in \mathcal{S}_n$, removing any element in the tensor product (i.e. discarding one of the n provided channels) yields an element in $\mathcal{S}_{n-1}$.
4. Each set $\mathcal{S}_n$ is closed under permuting the n subsystems of the input and output systems of a channel, i.e. for any permutation $\pi \in \mathfrak{S}_n$ and associated canonical unitary representations $P_A(\pi)$ and $P_B(\pi)$ on A^n and B^n, we have for all $\mathcal{E}_n \in \mathcal{S}_n$ that also the permuted channel $\rho \mapsto P_B(\pi)\mathcal{E}_n(P_A(\pi)\rho P_A(\pi)^\dagger)P_B(\pi)^\dagger$ is an element of $\mathcal{S}_n$.

One can then define the same scenarios, such as the composite i.i.d. setting, the arbitrarily varying setting, and slightly varying settings, as we did for composite state discrimination (below Definition 4.1.1) in a completely analogous way for composite channel discrimination. Note that since the n channels one receives

in the composite hypothesis testing problem need not all be equivalent, one might think that (in particular in an adaptive strategy) one might want to order the channels in a certain way, however it is not hard to see that this does not give any advantage, and so we can restrict to strategies that just take the channels in the order they are given. To see this, note that we assumed the sets of channels to be closed under permutations and we are looking at worst-case error probabilities. Hence, for every reordering one would perform for a given sequence of channels, there exists another sequence in the set that inverts this reordering, and hence in the worst-case one cannot gain anything.

4.2.1 Minimax and Exchange Lemmas

Before we study both parallel and adaptive strategies in detail, we use this section to introduce some of the technical Lemmas we use to prove our results. One key technical result here is Proposition 4.2.4, which allows us to exchange an infimum over sets of channels with a supremum over input states, if the sets of channels are convex.

Lemma 4.2.2. *For any two sets of states S and T, and any $\varepsilon \in [0,1]$,*

$$D_H^\varepsilon(S\|T) \leq \inf_{\substack{\rho\in S\\ \sigma\in T}} D_H^\varepsilon(\rho\|\sigma), \tag{4.2.1}$$

where the intuition is that the left-hand side corresponds to choosing one measurement for all pairings of ρ and σ, and the right-hand side allows for a different measurement for each pairing.

Proof. One finds that

$$2^{-D_H^\varepsilon(S\|T)} = \inf_{\substack{0\leq M\leq \mathbb{1}\\ \sup_{\rho\in S}\operatorname{Tr}\left(\bar{M}\rho\right)\leq\varepsilon}} \sup_{\sigma\in T} \operatorname{Tr}(M\sigma) \geq \sup_{\rho\in S} \inf_{\substack{0\leq M\leq \mathbb{1}\\ \operatorname{Tr}\left(\bar{M}\rho\right)\leq\varepsilon}} \sup_{\sigma\in T} \operatorname{Tr}(M\sigma) \tag{4.2.2}$$

$$\geq \sup_{\rho\in S}\sup_{\sigma\in T} \inf_{\substack{0\leq M\leq \mathbb{1}\\ \operatorname{Tr}\left(\bar{M}\rho\right)\leq\varepsilon}} \operatorname{Tr}(M\sigma) = \sup_{\substack{\rho\in S\\ \sigma\in T}} 2^{-D_H^\varepsilon(\rho\|\sigma)} \tag{4.2.3}$$

where we used the notation $\bar{M} = \mathbb{1} - M$, and the first inequality can be seen as follows: For any $\rho \in S$, if an M is chosen such that $\sup_{\rho'\in S}\operatorname{Tr}\left(\bar{M}\rho'\right) \leq \varepsilon$, then obviously also $\operatorname{Tr}\left(\bar{M}\rho\right) \leq \varepsilon$ and hence the infimum on the right-hand is over a set of M which can only be larger, and hence the expression can only be smaller. The second inequality is just a very basic property of infima and suprema, and the desired statement then follows by taking negative logarithms. □

Lemma 4.2.3 (Generalized minimax theorem Farkas and Revesz 2006, Theorem 5.2)**.** *Let X be a compact and convex subset of a Hausdorff topological vector space and let Y be a convex subset of a linear space. Let $f : X \times Y \to \mathbb{R}\cup\{\infty\}$ be lower semi-continuous on X for fixed $y \in Y$, convex in x and concave in y. Then*

$$\sup_{y\in Y}\inf_{x\in X} f(x,y) = \inf_{y\in Y}\sup_{x\in X} f(x,y)\,. \tag{4.2.4}$$

Remember that we say that a divergence $\mathbf{D}$ satisfies the direct-sum property, if

$$\mathbf{D}\left(\bigoplus_{i=1}^{n} p_i\rho_i \,\middle\|\, \bigoplus_{i=1}^{n} p_i\sigma_i\right) = \sum_{i=1}^{n} p_i \mathbf{D}(\rho_i\|\sigma_i)\,. \tag{4.2.5}$$

whenever $\rho_i, \sigma_i \in \mathcal{H}_i$ are two sets of density matrices and $\{p_i\}_{i=1}^n$ is a probability distribution.

Proposition 4.2.4. *Let $\mathcal{S}, \mathcal{T} \subset \mathrm{CPTP}(A \to B)$ be two closed, convex sets of channels. Let $\mathbf{D}$ be a quantum divergence that satisfies the data-processing inequality, is (jointly) lower semi-continuous, and also satisfies the direct-sum property. Then*

$$\inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} \sup_{\nu\in\mathcal{D}(RA)} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \sup_{\nu\in\mathcal{D}(RA)} \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)). \tag{4.2.6}$$

Proof. This proof is inspired by the proof of the similar minimax result in Brandão et al. (2020), Lemma 13. The $\geq$ direction follows immediately from very basic properties of $\inf$ and $\sup$. For the $\leq$ direction, let μ be a discrete measure on the set of density matrices $\mathcal{D}(RA)$, and consider the function:

$$f((\mathcal{E},\mathcal{F}),\mu) \coloneqq \mathop{\mathbb{E}}_{\nu\sim\mu} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu))\,. \tag{4.2.7}$$

where we write $\mathop{\mathbb{E}}_{\nu\sim\mu}$ for the expectation value with respect to the measure μ (this is the same as integrating ν with respect to the measure μ). It is easy to see (for any divergence) that the data-processing inequality and the direct-sum property together imply joint convexity (see e.g. Khatri and Wilde 2020). Hence the function f is convex in its first argument, and it is clearly also linear (and hence concave) in the second argument. Note also that the set of channels $\mathrm{CPTP}(A \to B)$ is bounded (for example in diamond norm, i.e. the trace distance of the channel outputs evaluated on the same input state, then taking the supremum over input states; this yields a norm on quantum channels which is always upper bounded by 2) and hence compact, and so are the closed subsets $\mathcal{S}$ and $\mathcal{T}$. Then, by Lemma 4.2.3 we have

$$\inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} \sup_{\mu} \mathop{\mathbb{E}}_{\nu\sim\mu} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \sup_{\mu} \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} \mathop{\mathbb{E}}_{\nu\sim\mu} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)). \tag{4.2.8}$$

We can lower bound the left-hand side by restricting the supremum to singular (i.e. Dirac) measures (this is in fact an equality), which recovers the left-hand side of (4.2.6). For the right-hand side, note that by Caratheodory's theorem we can write the expectation value as a convex combination with a finite number of terms. For a given μ we will write

$$\mathop{\mathbb{E}}_{\nu\sim\mu} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \sum_x p_x \mathbf{D}(\mathcal{E}(\nu_x)\|\mathcal{F}(\nu_x)). \tag{4.2.9}$$

Now define the new state

$$\tilde{\nu}_{XRA} \coloneqq \sum_x p_x \left|x\right\rangle\!\left\langle x\right| \otimes \nu_x. \tag{4.2.10}$$

By the direct-sum property of $\mathbf{D}$ (note that $\mathcal{E}$ and $\mathcal{F}$ act with an identity on the new system X)

$$\mathbf{D}(\mathcal{E}(\tilde{\nu})\|\mathcal{F}(\tilde{\nu})) = \sum_x p_x \mathbf{D}(\mathcal{E}(\nu_x)\|\mathcal{F}(\nu_x))\,, \tag{4.2.11}$$

and so the supremum over measures can be replaced by a supremum over states ν_{XRA}. Finally, by Lemma 2.2.4 the supremum can further be restricted to states ν_{RA} where $R \cong A$. □

For the next Lemma, remember our definitions of permutation covariant channels from Section 2.6.

Lemma 4.2.5. *Let $\mathcal{S}_n, \mathcal{T}_n \subset \mathrm{CPTP}(A^n \to B^n)$ be closed convex and also closed under permutations, and let $\mathbf{D}$ be lower semi-continuous and satisfy the data-processing inequality. Then,*

$$\inf_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n}} \sup_{\substack{\nu\in\mathcal{D}(R\otimes A^{\otimes n})\\ R \text{ arbitrary}}} \mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) = \min_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n\\ \mathcal{E}_n,\mathcal{F}_n \text{ perm. covariant.}}} \sup_{\substack{\nu\in\mathcal{D}(R\otimes A^{\otimes n})\\ R \text{ arbitrary}}} \mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) \tag{4.2.12}$$

i.e. the infimum is achieved for permutation covariant elements of $\mathcal{S}_n$ and $\mathcal{T}_n$.

Proof. First, the infimum is achieved since the infimum of a lower semi-continuous function over a compact set is achieved, and the supremum of multiple lower semi-continuous functions is lower semi-continuous. Let $\mathcal{E}_n \in \mathcal{S}_n$ and $\mathcal{F}_n \in \mathcal{T}_n$ be two channels, and consider the permuted versions:

$$\bar{\mathcal{E}}_n(\rho) \coloneqq \frac{1}{n!}\sum_{\pi\in\mathfrak{S}_n} P_B(\pi)\mathcal{E}_n(P_A(\pi)\rho P_A(\pi)^\dagger)P_B(\pi)^\dagger \tag{4.2.13}$$

$$\bar{\mathcal{F}}_n(\rho) \coloneqq \frac{1}{n!}\sum_{\pi\in\mathfrak{S}_n} P_B(\pi)\mathcal{F}_n(P_A(\pi)\rho P_A(\pi)^\dagger)P_B(\pi)^\dagger \tag{4.2.14}$$

These two permuted versions can also be seen as a permutation super-channel having been applied to $\mathcal{E}_n$ and $\mathcal{F}_n$, and it is known that any channel divergence can only decrease under the action of such super-channels (see e.g. Gour 2019), however, we will still show this explicitly again here for the reader's convenience:

Define the channel $\mathcal{A} : A^n \to A^n \otimes R'$

$$\mathcal{A}(\rho) \coloneqq \frac{1}{n!}\sum_{\pi\in\mathfrak{S}_n} (P_A(\pi)\rho P_A(\pi)^\dagger) \otimes |\pi\rangle\langle\pi|_{R'} \tag{4.2.15}$$

where R' is some additional classical register storing the permutation π. Defining $\mathcal{B} : B^n \otimes R' \to B^n$:

$$\mathcal{B}(\rho) \coloneqq \sum_{\pi\in\mathfrak{S}_n} P_B(\pi)\,\langle\pi|_{R'}\,\rho\,|\pi\rangle_{R'}\,P_B(\pi)^\dagger \tag{4.2.16}$$

we find that

$$\bar{\mathcal{E}}_n = \mathcal{B}\circ(\mathcal{E}_n\otimes \mathrm{id}_{R'})\circ\mathcal{A} \tag{4.2.17}$$

$$\bar{\mathcal{F}}_n = \mathcal{B}\circ(\mathcal{F}_n\otimes \mathrm{id}_{R'})\circ\mathcal{A}\,. \tag{4.2.18}$$

Hence

$$\sup_{\substack{\nu\in\mathcal{D}\left(R\otimes A^{\otimes n}\right)\\ R \text{ arbitrary}}} \mathbf{D}(\bar{\mathcal{E}}_n(\nu)\|\bar{\mathcal{F}}_n(\nu)) \leq \sup_{\substack{\nu\in\mathcal{D}\left(R\otimes A^{\otimes n}\right)\\ R \text{ arbitrary}}} \mathbf{D}(\mathcal{E}_n(\mathcal{A}(\nu))\|\mathcal{F}_n(\mathcal{A}(\nu))) \leq \sup_{\substack{\nu\in\mathcal{D}\left(R\otimes A^{\otimes n}\right)\\ R \text{ arbitrary}}} \mathbf{D}(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu)) \tag{4.2.19}$$

where the first inequality is the normal data-processing inequality (we omitted the $\mathrm{id}_{R'}$), and the second inequality uses that all the states $\mathcal{A}(\nu)$ are itself included in the supremum. As it is easy to see that the channels $\bar{\mathcal{E}}_n$ and $\bar{\mathcal{F}}_n$ are permutation covariant and are also included in $\mathcal{S}_n$ and $\mathcal{T}_n$ (as they were assumed to be closed under permutations), the minimum will be achieved for permutation covariant channels. □

4.3 Parallel Discrimination Strategies

Parallel discrimination strategies for composite channel discrimination are essentially the same as in simple channel discrimination (Section 3.1). While the black-box channels need no longer all be identical, the parallel strategy is still fully specified by a joint input state (see Figure 4.1 for an illustration) and a measurement at the end.

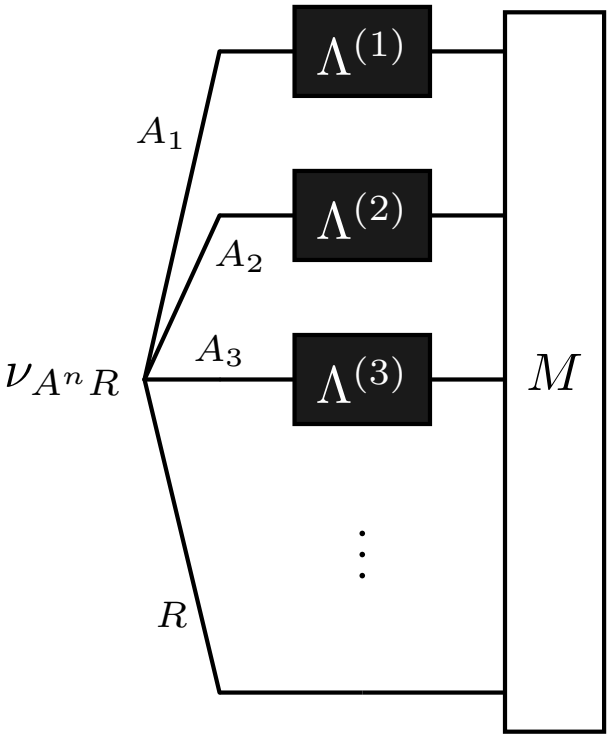

Figure 4.1: Illustration of a parallel protocol with n not necessarily identical black-box channels.

Given a set of channels $\mathcal{A}$ and an input state $\nu \in \mathcal{D}(RA)$ (where R could be any system, possibly also just trivial), we define the set of all output states as

$$\mathcal{A}[\nu] \coloneqq \{ (\mathrm{id}_R \otimes \mathcal{E})(\nu) \mid \mathcal{E} \in \mathcal{A} \} . \tag{4.3.1}$$

Since we want to be looking at worst-case errors again (as introduced for composite state discrimination in Section 4.1), we will be looking for the best input state ν_n and measurement M, such that *for all* $\mathcal{E}_n \in \mathcal{S}_n$ the error of claiming it coming from $\mathcal{T}_n$ (i.e. the type I error) stays below some threshold ε and we otherwise minimize the worst case type II error, i.e. we want to make sure that the probability of claiming an element $\mathcal{F}_n \in \mathcal{T}_n$ to be from $\mathcal{S}_n$ is as low as possible uniformly over all $\mathcal{F}_n \in \mathcal{T}_n$. Given a joint input state ν, the parallel channel discrimination problem turns into a state discrimination problem, and so we define the following type II error exponent for any $\mathcal{S}_n$ and $\mathcal{T}_n$ which satisfy the properties of Definition 4.2.1:

$$D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n) := \sup_{\nu\in\mathcal{D}(RA)} D_H^\varepsilon(\mathcal{S}_n[\nu]\|\mathcal{T}_n[\nu]) = \sup_{\nu\in\mathcal{D}(RA)} \sup_{\substack{0\leq M\leq\mathbb{1}\\ \alpha(M,\mathcal{S}_n[\nu])\leq\varepsilon}} (-\log\beta(M,\mathcal{T}_n[\nu])) \tag{4.3.2}$$

$$e_P(n,\varepsilon,\mathcal{S}_n,\mathcal{T}_n) := \frac{1}{n} D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n)\,. \tag{4.3.3}$$

It is easy to see that $\mathcal{C}(\mathcal{A}[\nu]) = \mathcal{C}(\mathcal{A})[\nu]$, and hence (as above) for any two sets of channels $\mathcal{S},\mathcal{T}$:

$$D_H^\varepsilon(\mathcal{S}\|\mathcal{T}) = D_H^\varepsilon(\mathcal{S}\|\mathcal{C}(\mathcal{T})) = D_H^\varepsilon(\mathcal{C}(\mathcal{S})\|\mathcal{T}) = D_H^\varepsilon(\mathcal{C}(\mathcal{S})\|\mathcal{C}(\mathcal{T}))\,. \tag{4.3.4}$$

Our main theorem of this section is the following:

Theorem 4.3.1. *Let $\boldsymbol{\mathcal{S}} = (\mathcal{S}_n)_n, \boldsymbol{\mathcal{T}} = (\mathcal{T}_n)_n$ be two composite quantum channel hypotheses (as defined in Definition 4.2.1). Then, the quantum Stein exponent of discriminating these two hypotheses with a parallel strategy is given by:*

$$\lim_{\varepsilon\to 0}\overline{\lim_{n\to\infty}}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n) = \lim_{n\to\infty}\frac{1}{n}\min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}}\max_{\nu\in\mathcal{D}(R\otimes A^{\otimes n})} D(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu))\,, \tag{4.3.5}$$

where $\overline{\lim}$ can be $\liminf$ or $\limsup$ (see Definition 4.1.4). Additionally, on the right-hand side the $\min$ and $\max$ can be exchanged, and one can choose the reference system R to be isomorphic to $A^{\otimes n}$ for all n.

Furthermore, if each $\mathcal{S}_n$ lies in the intersection of $\mathrm{CPTP}(A^n\to B^n)$ with a linear subspace, with dimension polynomial in n, of the space of linear maps $A^n\to B^n$ (this is for example the case in the composite i.i.d. setting, where all the $\mathcal{E}_n\in\mathcal{S}_n$ are permutation covariant), we can also remove one convex hull:

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n) = \lim_{n\to\infty}\frac{1}{n}\max_{\nu\in\mathcal{D}(R\otimes A^{\otimes n})}\min_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}} D(\mathcal{E}_n(\nu)\|\mathcal{F}_n(\nu))\,, \tag{4.3.6}$$

where we however cannot say whether $\min$ and $\max$ can be exchanged.

Proof. This proof is very much inspired by the results for composite state discrimination from Berta, Brandao, and Hirche (2021), Theroem 1.1, and Brandão et al. (2020), Theorem 16.

Achievability For the achievability part, let $\varepsilon\in(0,1)$, fix an integer k, and let $\nu_k\in\mathcal{D}\left(RA^k\right)$ be an input state, where R is isomorphic to A^k. Additionally, let $\mathcal{M}_k$ be a POVM measurement on RB^k (where we interpret $\mathcal{M}_k$ as a quantum-classical channel that maps to the probability distribution of measurement

outcomes, as specified in Section 2.1.3). Define the two sets of classical probability distributions $P := \{ \mathcal{M}_k(\mathcal{E}_k(\nu_k)) \mid \mathcal{E}_k \in \mathcal{S}_k \}$ and $Q := \{ \mathcal{M}_k(\mathcal{F}_k(\nu_k)) \mid \mathcal{F}_k \in \mathcal{T}_k \}$. The operational procedure is now to take an unknown channel from either $\mathcal{S}_{nk}$ or $\mathcal{T}_{nk}$, feed it with the input state $\nu_k^{\otimes n}$ and apply the measurement $\mathcal{M}_k^{\otimes n}$ to the outcome. Crucially, due to the assumed structure of the $(\mathcal{S}_n)_n$ and $(\mathcal{T}_n)_n$ (as specified in Definition 4.2.1), the measurement result of each of the individual n POVM measurements will be distributed according to a $p \in P$ or $q \in Q$. Hence, the overall structure of classical outcomes can be seen as an instance of adversarial hypothesis testing with a particular adversary[1]. For this classical problem, by Theorem 4.1.8, the exponent

$$\inf_{p\in P, q\in Q} D(p\|q) \tag{4.3.7}$$

is asymptotically achievable as $n \to \infty$, which just means that

$$\liminf_{n\to\infty} \frac{1}{n} D_H^\varepsilon(\mathcal{S}_{nk}\|\mathcal{T}_{nk}) \geq \inf_{p\in P, q\in Q} D(p\|q) = \inf_{\substack{\rho_k\in\mathcal{S}_k[\nu_k]\\ \sigma_k\in\mathcal{T}_k[\nu_k]}} D(\mathcal{M}_k(\rho_k)\|\mathcal{M}_k(\sigma_k))\,, \tag{4.3.8}$$

where dividing by k yields:

$$\liminf_{n\to\infty} \frac{1}{nk} D_H^\varepsilon(\mathcal{S}_{nk}\|\mathcal{T}_{nk}) \geq \frac{1}{k} \inf_{\substack{\rho_k\in\mathcal{S}_k[\nu_k]\\ \sigma_k\in\mathcal{T}_k[\nu_k]}} D(\mathcal{M}_k(\rho_k)\|\mathcal{M}_k(\sigma_k))\,. \tag{4.3.9}$$

Now, to obtain a procedure for discriminating m channels where m is not a multiple of k, we can just ignore at most $k-1$ channels so that we are left with a multiple of k channels and then do the above. This yields a strategy to distinguish $\mathcal{S}_m$ and $\mathcal{T}_m$ for any m and asymptotically the $k-1$ discarded channels do not matter, so we get:

$$\liminf_{m\to\infty} \frac{1}{m} D_H^\varepsilon(\mathcal{S}_m\|\mathcal{T}_m) \geq \frac{1}{k} \inf_{\substack{\rho_k\in\mathcal{S}_k[\nu_k]\\ \sigma_k\in\mathcal{T}_k[\nu_k]}} D(\mathcal{M}_k(\rho_k)\|\mathcal{M}_k(\sigma_k)) \geq \inf_{\substack{\rho_k\in\mathcal{C}(\mathcal{S}_k[\nu_k])\\ \sigma_k\in\mathcal{C}(\mathcal{T}_k[\nu_k])}} \frac{1}{k} D(\mathcal{M}_k(\rho_k)\|\mathcal{M}_k(\sigma_k))\,, \tag{4.3.10}$$

where we added convex hulls on the right-hand side (this just makes the infimum smaller). We can now take the supremum over all measurements $\mathcal{M}_k$ on the right-hand side, and by Brandão et al. (2020), Lemma 13, we can exchange this supremum with the already present infimum, to find

$$\liminf_{m\to\infty} \frac{1}{m} D_H^\varepsilon(\mathcal{S}_m\|\mathcal{T}_m) \geq \inf_{\substack{\rho_k\in\mathcal{C}(\mathcal{S}_k[\nu_k])\\ \sigma_k\in\mathcal{C}(T_k[\nu_k])}} \frac{1}{k} D_M(\rho_k\|\sigma_k)\,. \tag{4.3.11}$$

Note that Brandão et al. (2020), Lemma 13, requires the infimum to be over a convex set, which is why we

[1] In fact, this problem can also be seen to be at most as hard as a composite hypothesis testing task in the arbitrarily varying case, and a similar statement as Theorem 4.1.8 for this composite arbitrarily varying task would be sufficient for our purposes.

introduced convex hulls in the previous step. Additionally, we now take the supremum over ν_k to find

$$\liminf_{m\to\infty}\frac{1}{m}D_H^\varepsilon(\mathcal{S}_m\|\mathcal{T}_m) \geq \sup_{\nu_k\in\mathcal{D}(RA^k)}\inf_{\substack{\rho_k\in\mathcal{C}(\mathcal{S}_k[\nu_k])\\ \sigma_k\in\mathcal{C}(T_k[\nu_k])}}\frac{1}{k}D_M(\rho_k\|\sigma_k) \tag{4.3.12}$$

$$= \sup_{\nu_k\in\mathcal{D}(RA^k)}\inf_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(T_k)}}\frac{1}{k}D_M(\mathcal{E}_k(\nu_k)\|\mathcal{F}_k(\nu_k)) \tag{4.3.13}$$

$$= \inf_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(T_k)}}\sup_{\nu_k\in\mathcal{D}(RA^k)}\frac{1}{k}D_M(\mathcal{E}_k(\nu_k)\|\mathcal{F}_k(\nu_k)) \tag{4.3.14}$$

where the first equality is just a rewriting, and for the second equality we used that by Proposition 4.2.4 (since the infimum is over convex sets, and D_M satisfies the direct sum property) we can exchange infimum and supremum. We take the lim sup over k to get

$$\liminf_{m\to\infty}\frac{1}{m}D_H^\varepsilon(\mathcal{S}_m\|\mathcal{T}_m) \geq \limsup_{k\to\infty}\inf_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(T_k)}}\sup_{\nu_k\in\mathcal{D}(RA^k)}\frac{1}{k}D_M(\mathcal{E}_k(\nu_k)\|\mathcal{F}_k(\nu_k))\,. \tag{4.3.15}$$

Now, by Lemma 4.2.5 the infimum is achieved for permutation covariant channels $\mathcal{E}_k$, $\mathcal{F}_k$, and by Lemma 2.6.1 the supremum is achieved for a permutation invariant state (note that the channels $\mathcal{E}_k$ and $\mathcal{F}_k$ are of course also permutation covariant with regards to permutations within R, as they act with the identity on the reference system). Hence the state $\mathcal{F}_k(\nu_k)$ is permutation invariant, and thus by Lemma 2.6.2 we get

$$\liminf_{m\to\infty}\frac{1}{m}D_H^\varepsilon(\mathcal{S}_m\|\mathcal{T}_m) \geq \limsup_{k\to\infty}\min_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(T_k)}}\max_{\nu_k\in\mathcal{D}(R\otimes A^k)}\frac{1}{k}D(\mathcal{E}_k(\nu_k)\|\mathcal{F}_k(\nu_k)) \tag{4.3.16}$$

$$= \limsup_{k\to\infty}\min_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(T_k)}}\frac{1}{k}D(\mathcal{E}_k\|\mathcal{F}_k)\,. \tag{4.3.17}$$

Converse For the converse part, let $R\cong A$, and then note that by Lemma 4.2.2:

$$D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n) = \sup_{\nu_n\in\mathcal{D}(R^nA^n)}D_H^\varepsilon(\mathcal{S}_n[\nu_n]\|T_n[\nu_n]) \leq \sup_{\nu_n\in\mathcal{D}(R^nA^n)}\inf_{\substack{\rho_n\in\mathcal{S}_n[\nu_n]\\ \sigma_n\in\mathcal{T}_n[\nu_n]}}D_H^\varepsilon(\rho_n\|\sigma_n)\,. \tag{4.3.18}$$

By Lemma 2.4.1 we have that for any two states ρ,σ:

$$D_H^\varepsilon(\rho\|\sigma) \leq \frac{1}{1-\varepsilon}(D(\rho\|\sigma)+h(\varepsilon))\,. \tag{4.3.19}$$

Thus,

$$\lim_{\varepsilon\to 0}\liminf_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n)\overset{(4.3.4)}{=}\lim_{\varepsilon\to 0}\liminf_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{C}(\mathcal{S}_n)\|\mathcal{C}(\mathcal{T}_n)) \tag{4.3.20}$$

$$=\lim_{\varepsilon\to 0}\liminf_{n\to\infty}\frac{1}{n}\sup_{\nu_n\in\mathcal{D}(R^nA^n)}D_H^\varepsilon(\mathcal{C}(\mathcal{S}_n)[\nu_n]\|\mathcal{C}(\mathcal{T}_n)[\nu_n]) \tag{4.3.21}$$

$$\overset{(4.3.18)}{\leq}\lim_{\varepsilon\to 0}\liminf_{n\to\infty}\sup_{\nu_n\in\mathcal{D}(R^nA^n)}\min_{\substack{\rho_n\in\mathcal{C}(\mathcal{S}_n)[\nu_n]\\ \sigma_n\in\mathcal{C}(\mathcal{T}_n)[\nu_n]}}\frac{1}{n}D_H^\varepsilon(\rho_n\|\sigma_n) \tag{4.3.22}$$

$$=\liminf_{n\to\infty}\sup_{\nu_n\in\mathcal{D}(R^nA^n)}\min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}}\frac{1}{n}D(\mathcal{E}_n(\nu_n)\|\mathcal{F}_n(\nu_n)) \tag{4.3.23}$$

$$\overset{(4.2.4)}{=}\liminf_{n\to\infty}\min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}}\frac{1}{n}D(\mathcal{E}_n\|\mathcal{F}_n) \tag{4.3.24}$$

where the optimizations are achieved by the same argument as above. Equivalently, one finds the same with $\liminf$ replaced with $\limsup$:

$$\lim_{\varepsilon\to 0}\limsup_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n)\leq\limsup_{n\to\infty}\min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}}\frac{1}{n}D(\mathcal{E}_n\|\mathcal{F}_n)\,. \tag{4.3.25}$$

Combining (4.3.24) with the achievability result (4.3.17), we find

$$\liminf_{n\to\infty}\min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}}\frac{1}{n}D(\mathcal{E}_n\|\mathcal{F}_n)\geq\lim_{\varepsilon\to 0}\liminf_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n)\geq\limsup_{k\to\infty}\min_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(\mathcal{T}_k)}}\frac{1}{k}D(\mathcal{E}_k\|\mathcal{F}_k) \tag{4.3.26}$$

and hence both inequalities in this line are in fact equalities. Also, combining this again with (4.3.24) and (4.3.25) we find

$$\lim_{k\to\infty}\min_{\substack{\mathcal{E}_k\in\mathcal{C}(\mathcal{S}_k)\\ \mathcal{F}_k\in\mathcal{C}(\mathcal{T}_k)}}\frac{1}{k}D(\mathcal{E}_k\|\mathcal{F}_k)\leq\lim_{\varepsilon\to 0}\liminf_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n) \tag{4.3.27}$$

$$\leq\lim_{\varepsilon\to 0}\limsup_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{S}_n\|\mathcal{T}_n)\leq\lim_{n\to\infty}\min_{\substack{\mathcal{E}_n\in\mathcal{C}(\mathcal{S}_n)\\ \mathcal{F}_n\in\mathcal{C}(\mathcal{T}_n)}}\frac{1}{n}D(\mathcal{E}_n\|\mathcal{F}_n)\,, \tag{4.3.28}$$

and hence all of these expressions coincide, and we get the desired statement with $\overline{\lim}$.

Finally, the second part of the theorem that applies if $\mathcal{S}_n$ also lies in a linear space with dimension polynomial in n, can be seen as an immediate consequence of the first part of the theorem after using Proposition 4.2.4 and Lemma 2.6.3. Note that after the application of Lemma 2.6.3 we do no longer satisfy the convexity assumption necessary for another application of Proposition 4.2.4, and hence we cannot conclude that the min and max can be exchanged again at this point. □

4.3.1 Classical Parallel Exponent for Finite Sets in the Composite IID Setting

The characterizations of the asymptotic error exponent in Theorem 4.3.1 are generally hard to calculate, and we would like to find scenarios where one can instead find single-letter formulas. One case in which one can do so (and which will be useful for upcoming examples) is in the *classical* composite i.i.d. case when the two sets $\mathcal{S}$ and $\mathcal{T}$ are also assumed to be finite.

Proposition 4.3.2. *Let $\mathcal{S}, \mathcal{T} \subset \mathrm{CPTP}(\mathcal{X} \to \mathcal{Y})$ be two finite sets of classical channels. In the composite i.i.d setting, i.e. with*

$$\mathcal{S}_n \coloneqq \left\{ \mathcal{E}^{\otimes n} \;\middle|\; \mathcal{E} \in \mathcal{S} \right\} \tag{4.3.29}$$

$$\mathcal{T}_n \coloneqq \left\{ \mathcal{F}^{\otimes n} \;\middle|\; \mathcal{F} \in \mathcal{T} \right\} \tag{4.3.30}$$

the Stein exponent of distinguishing these two hypotheses with a parallel strategy is given by

$$\lim_{\varepsilon \to 0} \overline{\lim_{n \to \infty}}\, e_P(n, \varepsilon, \mathcal{S}_n, \mathcal{T}_n) = \max_{\nu \in \mathcal{D}(\mathcal{X}'\mathcal{X})} \min_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) \,, \tag{4.3.31}$$

where $\mathcal{X}'$ is another classical system with the size of $\mathcal{X}$, and $\overline{\lim}$ can be $\liminf$ *or* $\limsup$ *(see Definition 4.1.4),*

Proof. We split the proof into the achievability and converse part.

Achievability Picking any classical input state $\nu \in \mathcal{D}(\mathcal{X}'\mathcal{X})$ and feeding identical copies of it into the n classical channels, turns this problem into the classical composite i.i.d. hypothesis testing problem of distinguishing the sets $P = \mathcal{S}[\nu]$, $Q = \mathcal{T}[\nu]$. Since they are both finite, we can apply Mosonyi, Szilágyi, and Weiner (2022), Theorem III.2, which states that the optimal exponent of this composite state discrimination problem is given by

$$\min_{p \in P, q \in Q} D(P \| Q) = \min_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) \,. \tag{4.3.32}$$

Taking the supremum over all input states ν yields the desired achievability result, and the supremum is achieved by an argument similar to Lemma 2.2.5, as the minimum over a finite number of elements does not affect any of the required continuity properties.

Converse It follows immediately from the definition (4.3.2), Lemma 4.2.2 and Lemma 2.4.1 that

$$\frac{1}{n} D_H^\varepsilon(\mathcal{S}_n \| \mathcal{T}_n) = \frac{1}{n} \sup_{\nu_n \in \mathcal{D}(R\mathcal{X}^n)} D_H^\varepsilon(\mathcal{S}_n[\nu_n] \| \mathcal{T}_n[\nu_n]) \tag{4.3.33}$$

$$\leq \frac{1}{n(1-\varepsilon)} \sup_{\nu_n \in \mathcal{D}(R\mathcal{X}^n)} \inf_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} D(\mathcal{E}^{\otimes n}(\nu_n) \| \mathcal{F}^{\otimes n}(\nu_n)) + o(1) \,. \tag{4.3.34}$$

where strictly speaking R could be a quantum system, and hence ν_n a quantum-classical state of the form

$$\nu_n = \sum_{i=1}^{d^n} p_i \rho_R^{(i)} \otimes |i\rangle\langle i|_{\mathcal{X}^n} \tag{4.3.35}$$

where $\{p_i\}_{i=1}^{d^n}$ is a probability distribution and the $\rho_R^{(i)}$ are density matrices on R.

By using the joint convexity of relative entropy and additivity under tensor products, it is easy to see though that for any such state ν_n there exists a probability distribution $\{q_i\}_{i=1}^{d}$ such that for all channels $\mathcal{E}$ and $\mathcal{F}$:

$$\frac{1}{n} D(\mathcal{E}^{\otimes n}(\nu_n) \| \mathcal{F}^{\otimes n}(\nu_n)) \leq \sum_{i=1}^{d} q_i D(\mathcal{E}(|i\rangle\langle i|) \| \mathcal{F}(|i\rangle\langle i|)) = D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) \,. \tag{4.3.36}$$

where $\nu = \sum_i q_i \, |i\rangle\langle i|_{\mathcal{X}'} \otimes |i\rangle\langle i|_{\mathcal{X}}$. Hence, we can upper bound the last part of (4.3.34) with the single-letter expression

$$\frac{1}{(1-\varepsilon)} \sup_{\nu \in \mathcal{D}(\mathcal{X}'\mathcal{X})} \min_{\substack{\mathcal{E}\in\mathcal{S} \\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) + o(1) \tag{4.3.37}$$

and the statement follows in the limit $n \to \infty$, $\varepsilon \to 0$.

□

4.4 Adaptive Discrimination Strategies

As for simple channel discrimination, the most general setup of the channels will allow for channel inputs to depend on previous channel outputs, and thus be adaptive. Let n be fixed and let $\Lambda_n = \Lambda^{(1)} \otimes \ldots \otimes \Lambda^{(n)}$ be n black-box channels given to us, where the task is to determine whether they come from $\mathcal{S}_n$ or $\mathcal{T}_n$, where $\mathcal{S}_n$ and $\mathcal{T}_n$ are part of quantum channel hypotheses as specified in Definition 4.2.1. We write just the first i channels as $\Lambda_i \coloneqq \Lambda^{(1)} \otimes \ldots \otimes \Lambda^{(i)}$ for $i = 1, ..., n$. A general adaptive channel discrimination protocol for these Λ_n, can now be fully specified by an initial state $\omega_0 \in \mathcal{D}(R \otimes A)$, a set of $n-1$ CPTP maps $\mathcal{N}_i : R \otimes B \to R \otimes A$, that transform the state before it is fed into the next black-box channel, and a final binary POVM $\{M, \mathbb{1} - M\}$ on $R \otimes B$. We will assume the size of reference system R to be fixed and identical throughout the protocol (this is without loss of generality). The protocol consists of alternating applications of a black-box channel and the preparation CPTP maps $\mathcal{N}_i$ (see Figure 4.2). We define:

$$\omega_i(\Lambda_i) \coloneqq \Lambda^{(i)}(\mathcal{N}_i(\omega_{i-1}(\Lambda_{i-1}))), \qquad \text{for } i \in \{2, \ldots, n\}, \tag{4.4.1}$$

where we do not make identities on reference systems explicit (as previously), and $\omega_1(\Lambda_1) \coloneqq \Lambda^{(1)}(\omega_0)$. With our notation, the final state before the action of the POVM will be $\omega_n(\Lambda_n)$. Note that since the sets $\mathcal{S}_n$ and $\mathcal{T}_n$ were assumed to be permutation invariant, there is no advantage to be gained from reordering the black-box channels and so this is indeed the most general setup.

For a set $\mathcal{S}_n$ corresponding to a hypothesis, we write $\omega_n(\mathcal{S}_n) \coloneqq \{\, \omega_n(\mathcal{E}_n) \mid \mathcal{E}_n \in \mathcal{S}_n \,\}$. Given an ω_n, the problem then reduces to the composite state-discrimination problem $\omega_n(\mathcal{S}_n)$ vs. $\omega_n(\mathcal{T}_n)$. Note that $\omega_n(\mathcal{S}_n) \subset \mathcal{D}(R \otimes B)$, so this state discrimination problem will not be an instance of a many-copy discrimination problem as studied above, the n just indicates how many channel black-boxes were used in obtaining the

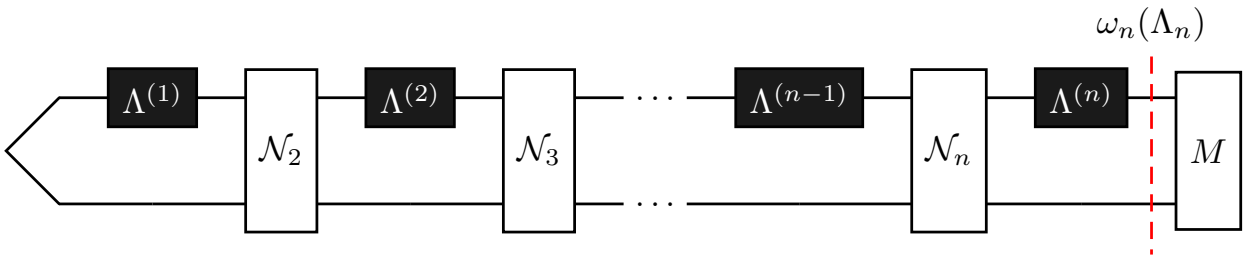

Figure 4.2: Illustration of a general adaptive protocol with n not necessarily identical black-box channels. The top row makes use of the given black-boxes while the bottom row depicts the memory system R.

states in the set. We can again define the corresponding worst-case type II error exponent as

$$e_A(n,\varepsilon,\mathcal{S}_n,\mathcal{T}_n) \coloneqq \frac{1}{n}\sup_{\omega_n} D_H^\varepsilon(\omega_n(\mathcal{S}_n)\|\omega_n(\mathcal{T}_n)) = \frac{1}{n}\sup_{\omega_n}\sup_{\substack{0\leq M\leq \mathbb{1}\\ \alpha(M,\omega_n(\mathcal{S}_n))\leq\varepsilon}} [-\log\beta(M,\omega_n(\mathcal{T}_n))] \tag{4.4.2}$$

where the supremum over ω_n goes over all adaptive strategies, i.e. all initial states ω_0 and all preparation maps $\mathcal{N}_i$, $i=2,\ldots,n$.

4.4.1 An Upper Bound for Adaptive Strategies

We can prove the following upper bound on the Stein exponent for discriminating two composite channel hypotheses with adaptive strategies. This captures the intuition that if the sets $\mathcal{S}_n$ and $\mathcal{T}_n$ are such that they include the i.i.d. problem, then the error exponent has to be less than the worst-case i.i.d. error exponent (for a similar statement for composite state discrimination, see e.g. Mosonyi, Szilágyi, and Weiner 2022).

Proposition 4.4.1. *Let $\boldsymbol{\mathcal{S}} = (\mathcal{S}_n)_n, \boldsymbol{\mathcal{T}} = (\mathcal{T}_n)_n$ be two quantum composite channel hypotheses (as in Definition 4.2.1). Then, for all n and $\varepsilon\in[0,1]$, it holds that*

$$e_A(n,\varepsilon,\mathcal{S}_n,\mathcal{T}_n) \leq \inf_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n}} e_A(n,\varepsilon,\mathcal{E}_n,\mathcal{F}_n)\,. \tag{4.4.3}$$

Furthermore, let $\mathcal{S}\coloneqq\mathcal{S}_1$ and $\mathcal{T}\coloneqq\mathcal{T}_1$. If the hypotheses are such that for all n

$$\mathcal{E}^{\otimes n}\in\mathcal{S}_n \quad \forall\,\mathcal{E}\in\mathcal{S} \tag{4.4.4}$$

$$\mathcal{F}^{\otimes n}\in\mathcal{T}_n \quad \forall\,\mathcal{F}\in\mathcal{T} \tag{4.4.5}$$

then the Stein exponent for distinguishing these two composite hypotheses by an adaptive strategy is upper bounded by

$$\lim_{\varepsilon\to0}\limsup_{n\to\infty} e_A(n,\varepsilon,\mathcal{S}_n,\mathcal{T}_n) \leq \min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D^A(\mathcal{E}\|\mathcal{F}) = \min_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F})\,. \tag{4.4.6}$$

Proof. As mentioned above, we write $\omega_n(\Lambda_n)$ for the state at the end of an adaptive strategy ω_n with n

channel uses, when the n black-box channels are given by Λ_n. With Lemma 4.2.2 we get:

$$e_A(n,\varepsilon,\mathcal{S}_n,\mathcal{T}_n) = \frac{1}{n}\sup_{\omega_n} D_H^\varepsilon(\omega_n(\mathcal{S}_n)\|\omega_n(\mathcal{T}_n)) \leq \frac{1}{n}\sup_{\omega_n}\inf_{\substack{\rho_n\in\omega_n(\mathcal{S}_n)\\ \sigma_n\in\omega_n(\mathcal{T}_n)}} D_H^\varepsilon(\rho_n\|\sigma_n) \tag{4.4.7}$$

$$= \frac{1}{n}\sup_{\omega_n}\inf_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n}} D_H^\varepsilon(\omega_n(\mathcal{E}_n)\|\omega_n(\mathcal{F}_n)) \leq \frac{1}{n}\inf_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n}}\sup_{\omega_n} D_H^\varepsilon(\omega_n(\mathcal{E}_n)\|\omega_n(\mathcal{F}_n)) \tag{4.4.8}$$

$$= \inf_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n}} e_A(n,\varepsilon,\mathcal{E}_n,\mathcal{F}_n) \tag{4.4.9}$$

which proves the first claim. For the second claim, if the requirements are satisfied, we get

$$\inf_{\substack{\mathcal{E}_n\in\mathcal{S}_n\\ \mathcal{F}_n\in\mathcal{T}_n}} e_A(n,\varepsilon,\mathcal{E}_n,\mathcal{F}_n) \leq \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} e_A(n,\varepsilon,\mathcal{E}^{\otimes n},\mathcal{F}^{\otimes n})\,. \tag{4.4.10}$$

The known characterization of the adaptive asymptotic error exponent of i.i.d. channels (see Section 3.1.1) thus implies

$$\lim_{\varepsilon\to 0}\limsup_{n\to\infty} e_A(n,\varepsilon,\mathcal{S}_n,\mathcal{T}_n) \leq \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D^A(\mathcal{E}\|\mathcal{F}) = \inf_{\substack{\mathcal{E}\in\mathcal{S}\\ \mathcal{F}\in\mathcal{T}}} D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F})\,. \tag{4.4.11}$$

Finally, the sequence $D(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n})$ is superadditive in n, and hence is monotonically increasing in n, and hence by Fekete's Lemma we can replace the limit $n\to\infty$ in the regularized divergence with a supremum over n. Thus, the regularized divergence is lower semi-continuous (as the supremum of lower semi-continuous functions is lower semi-continuous), and hence the infimum is achieved (and the same also for the infimum in D^A, since $D^A = D^{\mathrm{reg}}$). □

We will give a classical example in the next section where this upper bound is achieved and is strictly larger than the achievable exponent of parallel strategies. Hence this demonstrates an advantage of adaptive strategies for composite channel discrimination even if everything is classical.

Remark 4.4.2. While the upcoming example demonstrates that this upper bound can sometimes be achieved, it cannot always be achieved. Hence, it is not a candidate for the optimal asymptotic exponent of adaptive strategies. This can be seen by taking all channels to be replacer channels[2]. In this case the task of channel discrimination reduces to that of state discrimination, for which adaptive and parallel strategies are equivalent. In the composite i.i.d. setting (i.e. when $S_n = \{\,\rho^{\otimes n} \mid \rho\in S\,\}$ and $T_n = \{\,\sigma^{\otimes n} \mid \sigma\in T\,\}$) it has been shown that there exist sets S and T such that

$$\lim_{\varepsilon\to 0}\overline{\lim_{n\to\infty}} D_H^\varepsilon(S_n\|T_n) = \lim_{n\to\infty}\frac{1}{n}\inf_{\substack{\rho_n\in\mathcal{C}(S_n)\\ \sigma_n\in\mathcal{C}(T_n)}} D(\rho_n\|\sigma_n) < \inf_{\substack{\rho\in S\\ \sigma\in T}} D(\rho\|\sigma), \tag{4.4.12}$$

and different examples exist where S and T are either convex (Berta, Brandao, and Hirche 2021, Section 4.2) or discrete (Mosonyi, Szilágyi, and Weiner 2022, Section IV.A).

[2] A replacer channel is a quantum channel which outputs a fixed quantum state regardless of the input.

4.4.2 A Classical Example of an Adaptive Advantage

In the following, we give a fully classical example that demonstrates how adaptive strategies can be (also asymptotically) beneficial with composite hypotheses in the composite i.i.d. setting.

Example 4.4.3. There exist classical composite channel hypotheses $\mathcal{S} = \{\mathcal{E}_1, \mathcal{E}_2\}$ and $\mathcal{T} = \{\mathcal{F}_1, \mathcal{F}_2\}$, such that the adaptive error exponent in the composite i.i.d. setting is strictly larger than the parallel one. Specifically, we show that

$$\lim_{\varepsilon\to 0} \overline{\lim_{n\to\infty}} \frac{1}{n} e_A(n, \varepsilon, \mathcal{S}_n, \mathcal{T}_n) = \min_{i,j\in\{1,2\}} D(\mathcal{E}_i \| \mathcal{F}_j) = 2 \lim_{\varepsilon\to 0} \overline{\lim_{n\to\infty}} \frac{1}{n} e_P(n, \varepsilon, \mathcal{S}_n, \mathcal{T}_n) \tag{4.4.13}$$

where $\mathcal{S}_n = \left\{ \mathcal{E}_i^{\otimes n} \;\middle|\; i = 1, 2 \right\}$, $\mathcal{T}_n = \left\{ \mathcal{F}_i^{\otimes n} \;\middle|\; i = 1, 2 \right\}$.

When defining the channels, we will use quantum notation for convenience, but everything should be seen as classical, i.e. all states are diagonal in the computational basis.

The channels used in our example are then:

$$\mathcal{E}_1(\rho) = \tau \otimes |0\rangle\langle 0| \tag{4.4.14}$$

$$\mathcal{E}_2(\rho) = \tau \otimes |1\rangle\langle 1| \tag{4.4.15}$$

$$\mathcal{F}_1(\rho) = \frac{1}{2} \left[\tau + \langle 0|\rho|0\rangle\, |0\rangle\langle 0| + \langle 1|\rho|1\rangle\, \tau\right] \otimes |0\rangle\langle 0| \tag{4.4.16}$$

$$\mathcal{F}_2(\rho) = \frac{1}{2} \left[\tau + \langle 0|\rho|0\rangle\, \tau + \langle 1|\rho|1\rangle\, |0\rangle\langle 0|\right] \otimes |1\rangle\langle 1| \tag{4.4.17}$$

Where $\tau = \mathbb{1}_2/2$ is the maximally mixed state. For notational simplicity, we denote $\mathcal{E}(0) \coloneqq \mathcal{E}(|0\rangle\langle 0|)$, $\mathcal{E}(1) \coloneqq \mathcal{E}(|1\rangle\langle 1|)$.

The adaptive strategy The channels are constructed to allow for the following adaptive strategy: Given a black-box channel, we first use it with an arbitrary input state. Depending on the second output bit we will be able to determine with certainty the "index" of the channel, i.e. we will know that the channel is either $\mathcal{E}_1$ or $\mathcal{F}_1$ if the second bit is zero, or alternatively if the second bit is one we will know that the channel is either $\mathcal{E}_2$ or $\mathcal{F}_2$. It is easy to see that the optimal input state to discriminate $\mathcal{E}_1$ from $\mathcal{F}_1$ (optimal in the sense of asymptotic type II error decay rate) is $|0\rangle\langle 0|$, whereas the optimal input state to discriminate $\mathcal{E}_2$ from $\mathcal{F}_2$ is $|1\rangle\langle 1|$. Hence, in our adaptive strategy, for all subsequent channel uses, we input the value of the second bit we received out of the first channel use. This will lead to the following exponent:

$$\min_{i\in\{1,2\}} \max_{\rho\in\mathcal{D}(\mathcal{X})} D(\mathcal{E}_i(\rho)\|\mathcal{F}_i(\rho)) = D\big(\mathcal{E}_1(0)\|\mathcal{F}_1(0)\big) = D\big(\mathcal{E}_2(1)\|\mathcal{F}_2(1)\big) = \log_2(4/3)/2\,. \tag{4.4.18}$$

It is easy to see that this is also equal to

$$\min_{i,j\in\{1,2\}} \max_{\rho\in\mathcal{D}(\mathcal{X})} D(\mathcal{E}_i(\rho)\|\mathcal{F}_j(\rho)) \tag{4.4.19}$$

since this minimum is always achieved for $i = j$, as otherwise the second output bit allows for the two channels to be distinguished with certainty, which makes the relative entropy infinite. Since this is equal to the upper bound from Proposition 4.4.1 (for classical channels the regularized channel divergence collapses to the single-letter channel divergence), this is an asymptotically optimal adaptive strategy.

The best parallel strategy By Proposition 4.3.2, the optimal parallel exponent is given by

$$\max_{\nu \in \mathcal{D}(\mathcal{X}'\mathcal{X})} \min_{\substack{\mathcal{E}\in\mathcal{S} \\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}(\nu)\|\mathcal{F}(\nu))\,. \tag{4.4.20}$$

Similarly to the argument used in the proof of Proposition 4.3.2, by using the joint convexity of the relative entropy we find that for any state $\nu \in \mathcal{D}(\mathcal{X}'\mathcal{X})$ there exists a $p \in [0,1]$ such that for any i, j:

$$D(\mathcal{E}_i(\nu)\|\mathcal{F}_j(\nu)) \leq pD(\mathcal{E}_i(0)\|\mathcal{F}_j(0)) + (1-p)D(\mathcal{E}_i(1)\|F_j(1))\,, \tag{4.4.21}$$

and picking $\nu = p\,|00\rangle\langle 00|_{\mathcal{X}'\mathcal{X}} + (1-p)\,|11\rangle\langle 11|_{\mathcal{X}'\mathcal{X}}$ achieves the right-hand side. Hence, we can write:

$$\max_{\nu \in \mathcal{D}(\mathcal{X}'\mathcal{X})} \min_{\substack{\mathcal{E}\in\mathcal{S} \\ \mathcal{F}\in\mathcal{T}}} D(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) = \max_{0\leq p\leq 1} \min_{i,j\in\{1,2\}} \big(pD(\mathcal{E}_i(0)\|\mathcal{F}_j(0)) + (1-p)D(\mathcal{E}_i(1)\|F_j(1))\big)\,. \tag{4.4.22}$$

Similarly to above, the minimum will be achieved at $i = j$, and it is easy to see by explicit computation that the optimum value of p is $1/2$. Since $D(\mathcal{E}_2(0)\|\mathcal{F}_2(0)) = D(\mathcal{E}_1(1)\|F_1(1)) = 0$ the parallel exponent is thus

$$\frac{1}{2}D\big(\mathcal{E}_1(0)\|\mathcal{F}_1(0)\big) \tag{4.4.23}$$

which is half the exponent we were able to achieve with the adaptive strategy. It is also easy to see that a way to asymptotically achieve this parallel exponent is just to alternate the two input states 0 and 1. This captures the intuition that since we do not know the "index" of the channel in advance, we have to balance between the two optimal input states, and half of the time we will have chosen the wrong one, which means that half the channel outputs will be useless, and hence we can only achieve half the rate.

4.4.3 Classical Equality under Convexity

Looking back at the previous example, one finds that the advantage of the adaptive strategy can be seen as coming from the fact that the order of the maximum over input states and minimum over channels (for example in (4.4.18)) matters: The parallel strategy has to find a good input state for all channels (this corresponds to taking the maximum over states outside), whereas the adaptive strategy can reduce the problem to a simple discrimination problem between just two channels and then tailor the input state to these two channels (this corresponds to taking the infimum over channels outside). Indeed, one also finds that an application of our exchange result Proposition 4.2.4 (or similar minimax theorems) is not permitted in this example, as the sets of channels $\mathcal{S}$ and $\mathcal{T}$ are not convex. We show subsequently that, in the classical case, convexity of these sets is indeed sufficient for there not to be an advantage of adaptive strategies.

Theorem 4.4.4. *Let $\boldsymbol{\mathcal{S}} = (\mathcal{S}_n \subset \mathrm{CPTP}(\mathcal{X}^n \to \mathcal{Y}^n))_n, \boldsymbol{\mathcal{T}} = (\mathcal{T}_n \subset \mathrm{CPTP}(\mathcal{X}^n \to \mathcal{Y}^n))_n$ be two composite classical channel hypotheses (still satisfying the properties of Definition 4.2.1). If $\mathcal{S} := \mathcal{S}_1$ and $\mathcal{T} := \mathcal{T}_1$ are convex, and additionally for all n*

$$\mathcal{E}^{\otimes n} \in \mathcal{S}_n \quad \forall \mathcal{E} \in \mathcal{S} \tag{4.4.24}$$

$$\mathcal{F}^{\otimes n} \in \mathcal{T}_n \quad \forall \mathcal{F} \in \mathcal{T} \tag{4.4.25}$$

then the Stein exponent of distinguishing these two composite hypotheses (with a possible adaptive strategy) is given by

$$\lim_{\varepsilon \to 0} \overline{\lim_{n \to \infty}} \, e(n, \varepsilon, \mathcal{S}_n, \mathcal{T}_n) = \min_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} \max_{\nu \in \mathcal{D}(\mathcal{X})} D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) \tag{4.4.26}$$

where $\overline{\lim}$ can be $\liminf$ or $\limsup$ (see Definition 4.1.4), and this optimal exponent can be achieved with a parallel strategy.

Proof. We split the proof in to the achievability and converse parts.

Achievability Picking any classical input state $\nu \in \mathcal{D}(\mathcal{X})$ and feeding identical copies of it into the n classical channels, turns this problem into the classical composite hypothesis testing problem which is at most as hard as distinguishing the sets $P = \mathcal{S}[\nu]$, $Q = \mathcal{T}[\nu]$ in an adversarial setting (this follows from the properties of a composite channel hypothesis as specified in Definition 4.2.1). Then, by Theorem 4.1.8, the exponent

$$\min_{p \in P, q \in Q} D(p \| q) \tag{4.4.27}$$

is achievable, and hence, by optimizing over ν, also the exponent

$$\sup_{\nu \in \mathcal{D}(\mathcal{X})} \min_{\substack{p \in \mathcal{S}[\nu] \\ q \in \mathcal{T}[\nu]}} D(p \| q) = \sup_{\nu \in \mathcal{D}(\mathcal{X})} \min_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) \tag{4.4.28}$$

is achievable. Now, since $\mathcal{S}$ and $\mathcal{T}$ are convex, we can apply Proposition 4.2.4 and exchange the minimum and the supremum (where the supremum is also achieved, e.g. by Lemma 2.2.5).

Converse From Proposition 4.4.1 we get:

$$\lim_{\varepsilon \to 0} \lim_{n \to \infty} e(n, \varepsilon, \mathcal{S}_n, \mathcal{T}_n) \leq \min_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} D^{\mathrm{reg}}(\mathcal{E} \| \mathcal{F}) \,. \tag{4.4.29}$$

If all channels $\mathcal{E}$ and $\mathcal{F}$ are classical, the regularization is not necessary (Hayashi 2009). This can be easily seen as follows: Since the relative entropy is jointly convex, the optimization over the input state is achieved at an extreme point of the convex set of input states, and classically all extreme points are product distributions, which makes the regularization collapse and also eliminates the need for any reference system. Hence

$$\lim_{\varepsilon \to 0} \limsup_{n \to \infty} e(n, \varepsilon, \mathcal{S}_n, \mathcal{T}_n) \leq \min_{\substack{\mathcal{E} \in \mathcal{S} \\ \mathcal{F} \in \mathcal{T}}} \max_{\nu \in \mathcal{D}(\mathcal{X})} D(\mathcal{E}(\nu) \| \mathcal{F}(\nu)) \tag{4.4.30}$$

which is what we wanted to prove.

□

4.5 Outlook

We have been able to provide new insight into the relation between adaptive and parallel channel discrimination strategies, by studying such strategies for composite channel hypotheses and demonstrating that there is a gap in the asymptotic setting. However, there are still many open questions regarding composite channel discrimination, as can be seen by the number of cells in Table 4.1 for which we cannot give a definitive answer. Here, we want to briefly describe some of these problems and elaborate on possible solutions.

First of all, for classical composite hypotheses which are non-convex, we currently do not have an entropic expression for the optimal achievable rate of adaptive strategies, so far we have not even been able to prove that the worst-case i.i.d. upper bound cannot always be achieved[3]. Intuitively though, we consider it to be unlikely that this bound is always achieved, and we are also not particularly hopeful that there will be a simple entropic formula for the adaptive exponent. This comes from imagining generalizations of Example 4.4.3: In our example, determining the index of the channel within the two sets was possible perfectly after only one use, and hence one was able to use the optimal input state for all subsequent channel uses. One could, however, think about examples where determining this index is not perfectly possible, and hence one is expected to have to pay a certain (asymptotically non-vanishing) number of channel uses to distinguish the individual elements of the sets and then prepare the best input state, which should make the upper bound of Proposition 4.4.1 not achievable in this case. This procedure of determining which channels in the set we seem to be provided with also becomes significantly more complex once one stops having the symmetry between the sets $\mathcal{S}$ and $\mathcal{T}$ which we have in Example 4.4.3, and in the general case it is not obvious at all how one could capture in a simple entropic expression the intricacies of gaining knowledge about which elements in this set one might be given.

Additionally, we would like to see if there is an advantage for adaptive strategies in the quantum composite i.i.d. case when the sets of channels $\mathcal{S}_1$ and $\mathcal{T}_1$ are convex (recall that we showed that this is not possible classically). Given that the regularization is necessary in general in the quantum case, and the sets $\mathcal{S}_n$ and $\mathcal{T}_n$ will not be convex even if $\mathcal{S}_1$ and $\mathcal{T}_1$ are, we consider it not unlikely that there will again be an asymptotic gap between adaptive and parallel strategies.

[3] Mosonyi, Szilágyi, and Weiner (2022) provide an example where discriminating states is not possible with this upper bound. This, however, requires continuous probability distributions (i.e. the analogue of infinite-dimensional Hilbert spaces), which we do not consider here.

5 Infinite-Dimensional Quantum Channel Discrimination

5.1 Introduction – Why Infinite Dimensions?

This chapter deals with channel discrimination in infinite dimensions, extending many of the previous results to the infinite-dimensional setting, and establishing the theory of error exponents for asymmetric quantum channel discrimination also in this setting. While the prototypical system in quantum information theory is often considered finite-dimensional, infinite-dimensional, or continuous variable (CV), quantum systems are still of huge technological and experimental relevance. This is because quantum optics, which utilizes continuous quadrature amplitudes of the quantized electromagnetic field, has been shown to be a promising platform for efficient implementation of the essential steps in quantum communication protocols, namely, preparation, (unitary) manipulation, and measurement of (entangled) quantum states, (see e.g. Braunstein and van Loock 2005 for a review). Examples of such systems include collections of electromagnetic modes traveling along an optical fibre, and massive harmonic oscillators. Infinite-dimensional quantum channels which are of particular relevance include bosonic Gaussian channels (see e.g. Guha (2008) and references therein) and bosonic dephasing channels (see e.g. Lami and Wilde (2023)). The study of quantum information in infinite dimensions has found applications in quantum communication (Holevo and Werner 2001, Wolf, Pérez-García, and Giedke 2007, Takeoka, Guha, and Wilde 2014, Pirandola et al. 2017, Wilde, Tomamichel, and Berta 2017, Rosati, Mari, and Giovannetti 2018), the most noteworthy of them being the experimental realization of quantum teleportation for optical fields. Other applications include quantum computing and quantum error correction (Gottesman, Kitaev, and Preskill 2001, Guillaud and Mirrahimi 2019, Ofek et al. 2016, Michael et al. 2016, Guillaud and Mirrahimi 2019), quantum simulations (Flurin et al. 2017) and quantum sensing (Aasi et al. 2013, Zhang et al. 2018, Meyer et al. 2001, McCormick et al. 2019).

5.2 Quantum Information Theory in Infinite-Dimensional Hilbert Spaces

5.2.1 Operators on Infinite-Dimensional Hilbert Spaces

This section will very briefly introduce some of the fundamental concepts of infinite-dimensional Hilbert spaces. For a more thorough introduction, including proofs of all the statements, we refer to the literature, e.g. Pedersen (1989) and Reed and Simon (1980). Throughout this section we assume $\mathcal{H}$ to be a separable (in general infinite-dimensional) Hilbert space, that means there exists a countable orthonormal basis of $\mathcal{H}$, in the sense that every vector in $\mathcal{H}$ can be approximated arbitrarily well with finite linear combinations of basis elements. For any operator $A \in \mathrm{Lin}(\mathcal{H} \to \mathcal{H})$, we write its *operator norm* as $\|A\| := \|A\|_\infty := \sup_{v\in\mathcal{H},\|v\|=1} \|Av\|$, and write $\mathcal{B}(\mathcal{H}) := \{ A \in \mathrm{Lin}(\mathcal{H} \to \mathcal{H}) \mid \|A\| < \infty \}$ for the set of *bounded operators*. For any $A \in \mathcal{B}(\mathcal{H})$, let $A^\dagger \in \mathcal{B}(\mathcal{H})$ be the *adjoint* of A in the inner product of the vector space (i.e. for

$|\phi\rangle, |\psi\rangle \in \mathcal{H}$, $\langle\phi|A^\dagger\psi\rangle = \langle A\phi|\psi\rangle$). We call a bounded operator $A \in \mathcal{B}(\mathcal{H})$ *self-adjoint* if $A^\dagger = A$, i.e. for all $|\phi\rangle, |\psi\rangle \in \mathcal{H}$, $\langle\phi|A\psi\rangle = \langle A\phi|\psi\rangle$. We write $\mathcal{B}_{\mathrm{sa}}(\mathcal{H})$ for all bounded self-adjoint operators, and similar to the finite-dimensional notation, for $A, B \in \mathcal{B}_{\mathrm{sa}}(\mathcal{H})$ we say $A \leq B$ if $\langle\psi|A\psi\rangle \leq \langle\psi|B\psi\rangle$ for all $|\psi\rangle \in \mathcal{H}$. For a self-adjoint operator $A \in \mathcal{B}_{\mathrm{sa}}(\mathcal{H})$ that is also positive (i.e. $A \geq 0$), we can define its trace as

$$\mathrm{Tr}(A) = \sum_{i=1}^{\infty} \langle e_i| A |e_i\rangle \ , \tag{5.2.1}$$

where $\{|e_i\rangle\}_{i=1}^{\infty}$ is an orthonormal basis of $\mathcal{H}$ (which exists since $\mathcal{H}$ is separable). Note that the positivity of A implies that this infinite sum always exists (although it can be ∞), and is independent of the chosen orthonormal basis. Furthermore, for any $A \in \mathcal{B}(\mathcal{H})$, let $|A| \coloneqq \sqrt{A^\dagger A}$, where the square root can be defined e.g. by acting on the spectral decomposition of $A^\dagger A$. Then, we say that an operator $A \in \mathrm{Lin}(\mathcal{H} \to \mathcal{H})$ is *trace-class*, and write $A \in \mathcal{B}_1(\mathcal{H})$, if $\mathrm{Tr}(|A|) < \infty$. Furthermore, if A is trace-class, we can again define $\mathrm{Tr}(A)$ as in (5.2.1), and one can show that this limit is always finite and also does not depend on the chosen orthonormal basis. More generally, for $p \in [1, \infty)$ we define the *Schatten p-norm* for $A \in \mathcal{B}(\mathcal{H})$ as $\|A\|_p \coloneqq [\mathrm{Tr}(|A|^p)]^{1/p}$, and then write $\mathcal{B}_p(\mathcal{H}) \coloneqq \left\{ A \in \mathrm{Lin}(\mathcal{H} \to \mathcal{H}) \;\middle|\; \|A\|_p < \infty \right\}$. One can show that $\lim_{p\to\infty} \|A\|_p = \|A\|_\infty$ is the operator norm defined above. Finally, write $\mathscr{P}(\mathcal{H}) \coloneqq \{ A \in \mathcal{B}_1(\mathcal{H}) \mid A \geq 0 \}$ for all *positive* trace-class operators, and the set of all *density matrices* is then given by $\mathcal{D}(\mathcal{H}) \coloneqq \{ \rho \in \mathscr{P}(\mathcal{H}) \mid \mathrm{Tr}(\rho) = 1 \}$.

Relation to Concepts in General von Neumann Algebras

When defining and using properties of many quantum divergences, we will cite heavily from the works Hiai (2021), 2018, 2019 where many of these objects are introduced in the even more general setting of von Neumann algebras. Let us briefly discuss how the definitions and statements there correspond to the objects we will use in our simpler case of operators on a separable Hilbert space (note that this is not meant as an introduction into the theory of von Neumann algebras, for more details see e.g. Hiai (2021), Appendix A). For a general von Neumann algebra M, thought of as an "observable algebra" (the analogue of our $\mathcal{B}(\mathcal{H})$), the corresponding state-space then lies within its pre-dual M_* (the analogue of our $\mathcal{B}_1(\mathcal{H})$), in particular in the positive cone M_*^+ of the pre-dual (the analogue of our $\mathscr{P}(\mathcal{H})$), and the normalization constraint (analogous to our definition of $\mathcal{D}(\mathcal{H})$) corresponds to $\omega(\mathbb{1}) = 1$ for $\omega \in M_*^+$.

For a given state $\omega \in M_*^+$, one then often considers a representation of the von Neumann algebra M on a Hilbert space $\mathcal{H}_\omega$, given by $\pi_\omega : M \to \mathcal{B}(\mathcal{H}_\omega)$, such that there exists a vector $|\Omega\rangle \in \mathcal{H}_\omega$ which satisfies that $\omega(x) = \langle\Omega| \pi_\omega(x) |\Omega\rangle$ for all $x \in M$ (such a representation is for example given by the GNS construction). If $M = \mathcal{B}(\mathcal{H})$, then we can construct such a representation where $\mathcal{H}_\omega$ and π_ω are actually independent of ω, in particular we can choose $\mathcal{H}_w \coloneqq \mathcal{B}_2(\mathcal{H})$ the space of Hilbert-Schmidt operators (note that $\mathcal{B}_2(\mathcal{H})$ forms a Hilbert space with the inner product $\langle A, B\rangle = \mathrm{Tr}\big(A^\dagger B\big)$), and the corresponding vector $|\Omega\rangle$ for any $\omega \in \mathcal{D}(\mathcal{H}) \subset \mathcal{B}_1(\mathcal{H})$ is then given by $|\sqrt{\omega}\rangle \equiv \sqrt{\omega} \in \mathcal{B}_2(\mathcal{H})$. The representation π_ω then acts on a state $|\sigma\rangle \in \mathcal{B}_2(\mathcal{H})$ via operator (left-)multiplication, i.e. $\pi_\omega(x) |\sigma\rangle = |x\sigma\rangle$, where $x \in \mathcal{B}(\mathcal{H})$ guarantees that also $x\sigma \in \mathcal{B}_2(\mathcal{H})$. Furthermore, given such a representation (one can in fact show that the final construction is independent of the chosen representation), one can construct the Haagerup L^p spaces $L^p(M)$, which in our

case will correspond to the Schatten spaces $\mathcal{B}_p(\mathcal{H})$. One can then show that there exists a bijection between the pre-dual M_* and $L^1(M)$, written as $\psi \in \mathcal{M}_* \mapsto h_\psi \in L^1(M)$, which in our case is just the identity on $\mathcal{B}_1(\mathcal{H})$. While we write $A^\dagger$ for the adjoint of an element $A \in \mathcal{B}(\mathcal{H})$, we will usually write B^* for the adjoint of an operator $B \in \mathcal{B}(\mathcal{H}_\omega) = \mathcal{B}(\mathcal{B}_2(\mathcal{H}))$ on the representation space.

Unbounded Operators

We will occasionally also have to deal with unbounded operators, for example the relative modular operator defined below is such an unbounded operator on the representation space. Such an operator in general need not (and cannot) be defined on the entire space $\mathcal{H}$, so we write $A : \mathcal{D}(A) \to \mathcal{H}$, where $\mathcal{D}(A) \subset \mathcal{H}$ is a linear subspace called the *domain* of A. We say A is *densely defined*, if $\mathcal{D}(A)$ is dense in $\mathcal{H}$. For unbounded operators, one also has to take slightly more care when defining the adjoint of A. Let A be a densely defined operator, then a vector $|\psi\rangle \in \mathcal{H}$ will be in the domain of the *adjoint* A^* if $|\phi\rangle \in \mathcal{D}(A) \mapsto \langle\psi|A\phi\rangle$ is a continuous linear functional on the domain of A. For any such $|\psi\rangle \in \mathcal{D}(A^*)$, there then exists an extension of this linear functional to the entire space $\mathcal{H}$ (and this extension is unique since A is densely defined), and by the Riesz representation theorem this linear functional can be written as $|\phi\rangle \mapsto \langle\omega|\phi\rangle$, for some $\omega \in \mathcal{H}$, and one then defines $A^*|\psi\rangle := |\omega\rangle$. We call a densely defined operator A *self-adjoint* if $\mathcal{D}(A) = \mathcal{D}(A^*)$ and $A = A^*$ on their (now equal) domain. Note that the condition on the domains to be equal is necessary, the fact that A is equal to its adjoint A^* on its domain $\mathcal{D}(A)$ does not imply that the domains are equal.

We call an operator A *closed* if its graph $\Gamma(A) = \{(|\psi\rangle, A|\psi\rangle) \mid |\psi\rangle \in \mathcal{H}\} \subset \mathcal{H} \oplus \mathcal{H}$ is closed, or equivalently if for all sequences $(|\psi_n\rangle)_n \subset \mathcal{D}(A)$ such that $|\psi_n\rangle \to |\psi\rangle \in \mathcal{H}$ and $A|\psi_n\rangle \to |\phi\rangle \in \mathcal{H}$ it holds that $|\psi\rangle \in \mathcal{D}(A)$ and $A|\psi\rangle = |\phi\rangle$ (note that if $\mathcal{D}(A)$ is closed this is equivalent to A being continuous on its domain, otherwise the notions of continuity and closedness are different). From the definition of the adjoint above, it is not hard to see that the adjoint of any operator A is always closed, and hence also any self-adjoint operator is always closed.

We call an operator A *closable* if A is the restriction of a closed operator from a larger domain onto $\mathcal{D}(A)$. One can show that a densely defined operator A is closable if and only if its adjoint A^* is densely defined, in which case the closure of A is equal to $(A^*)^*$. Furthermore, if A is densely defined and closed then A^*A is densely defined and self-adjoint.

5.2.2 Quantum Channels

Quantum channels can in general be defined analogously to the finite-dimensional setting, for details and proofs of the following statements, see e.g. Holevo (2001) and references therein. For any C^* subalgebra (that is, any norm-closed linear sub-algebra, that is also closed under taking adjoints) $U \subset \mathcal{B}(\mathcal{H}_A)$ we call a linear map $\mathcal{E} : U \to \mathcal{B}(\mathcal{H}_B)$ *positive* if $\mathcal{E}(A) \geq 0$ whenever $A \geq 0$. With $\mathcal{M}_n$ the set of $n \times n$ matrices, we say that $\mathcal{E}$ is *completely positive* if $\mathcal{E} \otimes \mathrm{id}_{\mathcal{M}_n} : U \otimes \mathcal{M}_n \to \mathcal{B}(\mathcal{H}_B) \otimes \mathcal{M}_n$ is positive for all $n \in \mathbb{N}$. If $U = \mathcal{B}_1(\mathcal{H}_A)$, we say that $\mathcal{E}$ is *trace-preserving* if $\mathcal{E} : U \to \mathcal{B}_1(\mathcal{H}_B)$ and $\mathrm{Tr}(\mathcal{E}(\rho)) = \mathrm{Tr}(\rho)$ for all $\rho \in \mathcal{B}_1(\mathcal{H}_A)$. We call $\mathcal{E}$ a *quantum channel* if it is completely positive and trace-preserving. One can show (Davies 1976) that any quantum channel is bounded as an operator between $\mathcal{B}_1(\mathcal{H}_A)$ and $\mathcal{B}_1(\mathcal{H}_B)$ and hence norm-continuous.

5.2.3 Distance Metrics

Trace Distance

The trace distance in infinite dimensions is defined analogously as

$$T(\rho,\sigma) \coloneqq \frac{1}{2}\left\|\rho-\sigma\right\|_1 . \tag{5.2.2}$$

The trace distance satisfies the data-processing inequality since the trace norm is contractive under quantum channels. The latter can be seen by a standard argument which also works in infinite dimensions (throughout this section we write again $\langle A,B\rangle = \mathrm{Tr}\big(A^\dagger B\big)$ for the Hilbert-Schmidt inner product): It is well-known that the dual space of $\mathcal{B}_1(\mathcal{H})$ is $\mathcal{B}(\mathcal{H})$, i.e. any continuous linear functional $\mathcal{B}_1(\mathcal{H})\to\mathbb{C}$ can be written as $A\in\mathcal{B}_1(\mathcal{H})\mapsto\langle B,A\rangle$, where $B\in\mathcal{B}(\mathcal{H})$. For any channel $\Lambda:\mathcal{B}_1(\mathcal{H}_A)\to\mathcal{B}_1(\mathcal{H}_B)$, and any $B\in\mathcal{B}(\mathcal{H}_B)$, $A\mapsto\langle B,\Lambda(A)\rangle$ is a continuous linear functional on $\mathcal{B}_1(\mathcal{H}_A)$, which can hence be written as $A\mapsto\langle C,A\rangle$, for some $C\in\mathcal{B}(\mathcal{H}_A)$. We then set $\Lambda^*(B)\coloneqq C$, which defines a linear map $\Lambda^*:\mathcal{B}(\mathcal{H}_B)\to\mathcal{B}(\mathcal{H}_A)$. It is easy to see from the definition that Λ^* is completely positive iff Λ is, and Λ^* is unital (i.e. $\Lambda^*(\mathbb{1}_B)=\mathbb{1}_A$) iff Λ is trace-preserving (and vice-versa). In particular, if Λ is completely positive and trace-preserving, then $-\mathbb{1}_B\leq B\leq\mathbb{1}_B$ implies $-\mathbb{1}_A\leq\Lambda^*(B)\leq\mathbb{1}_A$, and hence Λ^* is contractive in $\left\|\cdot\right\|_\infty$-norm. We can then use the dual representation of the trace-norm (for $X\in\mathcal{B}_1(\mathcal{H}_B)$):

$$\left\|X\right\|_1 = \sup_{\substack{B\in\mathcal{B}(\mathcal{H}_B)\\ \left\|B\right\|_\infty=1}} |\langle B,X\rangle| \tag{5.2.3}$$

to conclude that (for $A\in\mathcal{B}_1(\mathcal{H}_A)$)

$$\left\|\Lambda(A)\right\|_1 = \sup_{\substack{B\in\mathcal{B}(\mathcal{H}_B)\\ \left\|B\right\|_\infty=1}} |\langle\Lambda^*(B),A\rangle| \leq \sup_{\substack{A'\in\mathcal{B}(\mathcal{H}_A)\\ \left\|A'\right\|_\infty=1}} \left|\langle A',A\rangle\right| = \left\|A\right\|_1 \tag{5.2.4}$$

and hence any quantum channel Λ is contractive in $\left\|\cdot\right\|_1$-norm.

Fidelity and Purified Distance

We can also define the fidelity in infinite dimensions completely analogously to the finite-dimensional case via (for $\rho,\sigma\in\mathcal{D}(\mathcal{H})$)

$$F(\rho,\sigma) \coloneqq \left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1 . \tag{5.2.5}$$

A key property of the fidelity is Uhlmann's Theorem, which in the infinite-dimensional context can be expressed as follows: For a state $\rho\in\mathcal{D}(\mathcal{H})$ we call an element $|\psi\rangle\in\mathcal{B}_2(\mathcal{H})$ a *purification* of ρ, if for all $A\in\mathcal{B}(\mathcal{H})$

$$\rho(A) = \mathrm{Tr}(\rho A) = \langle\psi|\,A\,|\psi\rangle . \tag{5.2.6}$$

Intuitively, one can easily see that this will be satisfied for $|\psi\rangle=\big|\sqrt{\rho}U\big\rangle$, for an arbitrary unitary U on $\mathcal{H}$, and in fact from the polar decomposition it follows that it is satisfied for only such $|\psi\rangle$. This corresponds

to the statement in the finite-dimensional setting that all purifications can be written as $(U_R \otimes \sqrt{\rho}) |\Gamma\rangle_{RA}$, where $|\Gamma\rangle_{RA}$ is an unnormalized maximally entangled state.

Uhlmann showed that (Uhlmann 1976)

$$F(\rho,\sigma) = \sup_{\substack{|\psi\rangle \text{ purification of } \rho \\ |\phi\rangle \text{ purification of } \sigma}} |\langle\phi|\psi\rangle| \,. \tag{5.2.7}$$

While expressing the purifications in terms of Hilbert-Schmidt vectors is theoretically appealing, we can also express things more in the spirit of the finite-dimensional approach including an additional reference Hilbert space. Let $\{|e_i\rangle\}_i$ be an orthonormal basis of $\mathcal{H}$, and for $\sqrt{\rho}U \in \mathcal{B}_2(\mathcal{H})$ define the state

$$|\Psi_\rho\rangle := \sum_i (\sqrt{\rho}U |e_i\rangle) \otimes |e_i\rangle \in \mathcal{H} \otimes \mathcal{H} \tag{5.2.8}$$

and it is easy to see that this is indeed a well-defined state with norm 1 on $\mathcal{H} \otimes \mathcal{H}$. We find that

$$\mathrm{Tr}_2(|\Psi_\rho\rangle\langle\Psi_\rho|) = \sum_i \sqrt{\rho}U |e_i\rangle\langle e_i| U^\dagger \sqrt{\rho} = \rho \tag{5.2.9}$$

and in this sense, then also $|\Psi_\rho\rangle$ is a purification of ρ. Furthermore, if $\left|\sqrt{\rho}U\right\rangle, |\sqrt{\sigma}V\rangle \in \mathcal{B}_2(\mathcal{H})$ are two purifications of ρ, σ in the original sense, and $|\Psi_\rho\rangle, |\Phi_\sigma\rangle \in \mathcal{H} \otimes \mathcal{H}$ are the corresponding constructions in the latter sense, then

$$\left\langle \sqrt{\sigma}V \middle| \sqrt{\rho}U \right\rangle_{\mathcal{B}_2(\mathcal{H})} = \mathrm{Tr}\left(V^\dagger \sqrt{\sigma}\sqrt{\rho}U\right) = \langle\Phi_\sigma|\Psi_\rho\rangle_{\mathcal{H}\otimes\mathcal{H}} \tag{5.2.10}$$

and so the two notions are equivalent. One can show that the fidelity satisfies the data-processing inequality also in infinite dimensions (Alberti and Uhlmann 1983).

The purified distance of two states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ is again defined as

$$P(\rho,\sigma) := \sqrt{1 - F(\rho,\sigma)^2} \,. \tag{5.2.11}$$

A Fuchs-van de Graaf Inequality

For completeness, we show here that one of the Fuchs-van de Graaf inequalities we will subsequently be using also holds in infinite dimensions, by a very straightforward adaption of the finite-dimensional proof (see e.g. Khatri and Wilde 2020). If $|\Psi_\rho\rangle, |\Phi_\sigma\rangle \in \mathcal{H} \otimes \mathcal{H}$ are two purifications of $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, then we have

$$T(\rho,\sigma) = \frac{1}{2}\|\rho - \sigma\|_1 \leq \frac{1}{2}\||\Psi_\rho\rangle\langle\Psi_\rho| - |\Phi_\sigma\rangle\langle\Phi_\sigma|\|_1 = \sqrt{1 - |\langle\Psi_\rho|\Phi_\sigma\rangle|^2} \tag{5.2.12}$$

where we used the data-processing inequality for the trace-distance, and for the second equality note that the problem is essentially two-dimensional at this point, and so we can evaluate the trace norm using

finite-dimensional methods. Now taking the infimum over all such purifications, we get

$$T(\rho,\sigma) \leq \sqrt{1-F(\rho,\sigma)^2} = P(\rho,\sigma) \tag{5.2.13}$$

5.3 Quantum Divergences in Infinite Dimensions

While already in finite dimensions we had to occasionally be careful when defining quantum (and classical) divergences, the situation is even more tricky in infinite dimensions, where the naive expressions are often not well-defined. In particular also, the condition $\rho \ll \sigma$ (again defined by the support of ρ being contained in the support of σ) will no longer be sufficient for the divergences in this thesis to be finite (see Section 5.6.3 for an exhaustive discussion of this).

Fundamentally though, we want a quantum divergence $\mathbf{D}$ to have similar properties as in finite dimensions: For all separable Hilbert spaces $\mathcal{H}$, it should be defined for pairs of positive trace-class operators $\mathbf{D} : \mathscr{P}(\mathcal{H}) \times \mathscr{P}(\mathcal{H}) \to \mathbb{R} \cup \{\infty\}$, such that it is again positive on states (i.e. for $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, $\mathbf{D}(\rho\|\sigma) \geq 0$), and satisfies the data-processing inequality (i.e. for $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, and a quantum channel $\mathcal{E} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{K})$ it holds that $\mathbf{D}(\mathcal{E}(\rho)\|\mathcal{E}(\sigma)) \leq \mathbf{D}(\rho\|\sigma)$).

Generally also, for any divergence $\mathbf{D}$ on states, and any two channels $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}_A) \to \mathscr{P}(\mathcal{H}_B)$ we will define the associated channel divergence

$$\mathbf{D}(\mathcal{E}\|\mathcal{F}) := \sup_{\nu \in \mathcal{D}(\mathcal{H}_{A'} \otimes \mathcal{H}_A)} \mathbf{D}(\mathcal{E}(\nu)\|\mathcal{F}(\nu)), \tag{5.3.1}$$

where $\mathcal{H}_{A'}$ is isomorphic to $\mathcal{H}_A$, and we use the same notation regarding implicit identities as in finite dimensions (see Section 2.1.2).

5.3.1 Standard f-Divergences and the Relative Modular Operator

To see why the definition of some of the previously encountered divergences can be tricky in infinite dimensions, remember the finite-dimensional definition of the Petz–Rényi divergence:

$$D_\alpha(\rho\|\sigma) = \frac{1}{\alpha-1} \log \operatorname{Tr}\left(\rho^\alpha \sigma^{1-\alpha}\right). \tag{5.3.2}$$

We see that for $\alpha > 1$ the definition will include σ with a negative exponent, which, (even if we restrict it to the support of σ) will in infinite dimensions not be a bounded operator, and hence troublesome to deal with. One way to deal with the associated problems, and give a well-defined definition of such divergences is to define them through what is called the relative modular operator. One advantage of this treatment is that the relative modular operator can be defined in general von Neumann algebras (see e.g. Hiai (2021) and Araki (1975)), although we will be introducing it here in a slightly more explicit way for the infinite-dimensional Hilbert space setting (this section is largely taken from and/or inspired by Androulakis and John (2023), Hiai (2021), and Araki (1977)).

Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be two given states, and let Π_σ be the projection onto the support of σ, and $\Pi_\sigma^\perp$ be the projection onto the orthogonal complement (i.e. the kernel of σ). The relative modular operator will be an

operator acting on "vectorized operators" (or in the von Neumann algebra setting, on the representation space), i.e. it is a linear operator on $\mathcal{B}_2(\mathcal{H})$. Throughout this section all kets (like $|\sqrt{\sigma}\rangle$) will refer to such elements of the Hilbert space $\mathcal{B}_2(\mathcal{H})$. Consider the subset

$$\mathcal{D}(S) \coloneqq \left\{ \left| X\sqrt{\sigma} + Y\Pi_\sigma^\perp \right\rangle \;\middle|\; X \in \mathcal{B}(\mathcal{H}), Y \in \mathcal{B}_2(\mathcal{H}) \right\} \subset \mathcal{B}_2(\mathcal{H}), \tag{5.3.3}$$

which is dense, since every element of $\mathcal{B}_2(\mathcal{H})$ can be approximated by an operator supported on a finite-dimensional subspace, and all such operators are included in $\mathcal{D}(S)$. Now, define $S : D(S) \to \mathcal{B}_2(\mathcal{H})$ as (this S depends on ρ and σ but we will still write S instead of $S_{\rho,\sigma}$ to not overload the notation too much):

$$S\left| X\sqrt{\sigma} + Y\Pi_\sigma^\perp \right\rangle = \left| \Pi_\sigma X^\dagger \sqrt{\rho} \right\rangle . \tag{5.3.4}$$

Note that this is not a linear, but an anti-linear operator. Additionally, consider the set

$$\mathcal{D}(S^*) \coloneqq \left\{ \left| \sqrt{\sigma}X + \Pi_\sigma^\perp Y \right\rangle \;\middle|\; X \in \mathcal{B}(\mathcal{H}), Y \in \mathcal{B}_2(\mathcal{H}) \right\} \subset \mathcal{B}_2(\mathcal{H}), \tag{5.3.5}$$

and the anti-linear operator

$$S^*\left| \sqrt{\sigma}X + \Pi_\sigma^\perp Y \right\rangle \coloneqq \left| \sqrt{\rho}X^\dagger \Pi_\sigma \right\rangle, \tag{5.3.6}$$

which is again densely defined.

Note that for an anti-linear operator A, the defining equation for the adjoint is that $\langle A^*\phi|\psi\rangle \overset{!}{=} \overline{\langle \phi|A\psi\rangle} = \langle A\psi|\phi\rangle$, where the overline denotes complex conjugation. This can be seen by following the argument for the definition of the adjoint for linear operators in Section 5.2.1, and noting that if A is anti-linear, then $|\psi\rangle \mapsto \overline{\langle\phi|A\psi\rangle}$ will be a continuous linear functional, and hence can be written as $\langle\omega|\psi\rangle =: \langle A^*\phi|\psi\rangle$, and A^* will be then again an anti-linear operator. One then also obtains analogous results for the theory of anti-linear operators, in particular any adjoint is automatically closed, and a densily defined anti-linear operator is closable if and only if its adjoint is densily defined. We can verify that S^* is indeed (possibly a restriction of) the adjoint of S:

$$\begin{aligned}
\left\langle \sqrt{\sigma}X_2 + \Pi_\sigma^\perp Y_2 \right| S \left| X_1\sqrt{\sigma} + Y_1\Pi_\sigma^\perp \right\rangle &= \left\langle \sqrt{\sigma}X_2 + \Pi_\sigma^\perp Y_2 \middle| \Pi_\sigma X_1^\dagger \sqrt{\rho} \right\rangle = \mathrm{Tr}\left(X_2^\dagger \sqrt{\sigma} X_1^\dagger \sqrt{\rho} \right) && (5.3.7)\\
&= \mathrm{Tr}\left(\sqrt{\sigma}X_1^\dagger \sqrt{\rho} X_2^\dagger \right) && (5.3.8)\\
&= \left\langle X_1\sqrt{\sigma} + Y_1\Pi_\sigma^\perp \middle| \sqrt{\rho}X_2^\dagger \Pi_\sigma \right\rangle && (5.3.9)\\
&= \left\langle X_1\sqrt{\sigma} \middle| S^* \middle| \sqrt{\sigma}X_2 + \Pi_\sigma^\perp Y_2 \right\rangle . && (5.3.10)
\end{aligned}$$

Hence, since S^* is densely defined, S is closable, and thus we can define the relative modular operator as

$$\Delta_{\rho,\sigma} \coloneqq S^*\overline{S}, \tag{5.3.11}$$

where $\overline{S}$ is the closure of S, and $\Delta_{\rho,\sigma}$ is thus a self-adjoint operator. One then finds that the domain of $\Delta_{\rho,\sigma}$

at least contains the set

$$\left\{ \left| X\sigma + Y\Pi_\sigma^\perp \right\rangle \;\middle|\; X \in \mathcal{B}(\mathcal{H}), Y \in \mathcal{B}_2(\mathcal{H}) \right\} \subset \mathcal{B}_2(\mathcal{H}) \tag{5.3.12}$$

and for $A \in \mathcal{B}(\mathcal{H})$ it holds that

$$\Delta_{\rho,\sigma} |A\sigma\rangle = S^*\overline{S} \left|A\sqrt{\sigma}\sqrt{\sigma}\right\rangle = S^*S \left|A\sqrt{\sigma}\sqrt{\sigma}\right\rangle = S^* \left|\sqrt{\sigma}A^\dagger\sqrt{\rho}\right\rangle = |\rho A \Pi_\sigma\rangle \tag{5.3.13}$$

where we used that the state $|A\sigma\rangle$ already lies in the domain of S. In that sense, this generalizes the finite-dimensional definition $\Delta_{\rho,\sigma}(X) \coloneqq \rho X \sigma^{-1}$. In fact, (5.3.13) allows us to conclude that if $\rho = \sum_i a_i P_i$, and $\sigma = \sum_i b_i Q_i$ are two spectral decompositions (where we assume without loss of generality the a_i and b_i to be non-zero), then

$$\Delta_{\rho,\sigma} |\psi\rangle = \sum_{i,j} a_i b_j^{-1} |P_i \psi Q_j\rangle \;, \tag{5.3.14}$$

and this is a (potentially degenerate) spectral decomposition of $\Delta_{\rho,\sigma}$, since the map $|\psi\rangle \mapsto |P_i\psi Q_j\rangle$ is a self-adjoint projection on $\mathcal{B}_2(\mathcal{H})$ (it is a projection since the P_i and Q_j are, and the self-adjointness follows from $\langle\phi|P_i\psi Q_j\rangle = \mathrm{Tr}\big(\phi^\dagger P_i \psi Q_j\big) = \langle P_i \phi Q_j | \psi\rangle$).

Subsequently, we will denote the spectral measure of $\Delta_{\rho,\sigma}$ as $\xi^{\Delta_{\rho,\sigma}}$, i.e. $\Delta_{\rho,\sigma} = \int_{[0,\infty)} \lambda d\xi^{\Delta_{\rho,\sigma}}(\lambda)$. From the explicit expression (5.3.14) it follows that

$$\left\langle\sqrt{\sigma}\right| d\xi^{\Delta_{\rho,\sigma}}(\lambda) \left|\sqrt{\sigma}\right\rangle = \sum_{i,j} \delta(a_i b_j^{-1} - \lambda) \,\mathrm{Tr}\big(P_i\sqrt{\sigma}Q_j\sqrt{\sigma}\big) = \sum_{i,j} \delta(a_i b_j^{-1} - \lambda) b_j \,\mathrm{Tr}(P_i Q_j) \tag{5.3.15}$$

$$= \sum_{i,j} \delta(a_i b_j^{-1} - \lambda) \frac{1}{\lambda} \,\mathrm{Tr}(\sqrt{\rho}P_i\sqrt{\rho}Q_j) = \frac{1}{\lambda} \left\langle\sqrt{\rho}\right| d\xi^{\Delta_{\rho,\sigma}}(\lambda) \left|\sqrt{\rho}\right\rangle \;, \tag{5.3.16}$$

where δ is the Dirac measure[1]. This also implies that if $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, the measure

$$\lambda \left\langle\sqrt{\sigma}\right| d\xi^{\Delta_{\rho,\sigma}}(\lambda) \left|\sqrt{\sigma}\right\rangle = \left\langle\sqrt{\rho}\right| d\xi^{\Delta_{\rho,\sigma}}(\lambda) \left|\sqrt{\rho}\right\rangle \tag{5.3.17}$$

is a probability measure.

For a convex (or concave) function $f : (0,\infty) \to \mathbb{R}$ its *standard f-divergence* is defined as (Hiai 2018, 2021)

$$D_f(\rho\|\sigma) \coloneqq f(0)\,\mathrm{Tr}[\sigma\{\rho = 0\}] + f'(\infty)\,\mathrm{Tr}[\rho\{\sigma = 0\}] + \int_0^\infty f(\lambda) \langle\sqrt{\sigma}| d\xi^{\Delta_{\rho,\sigma}}(\lambda) |\sqrt{\sigma}\rangle \,, \tag{5.3.18}$$

$$= f(0)\,\mathrm{Tr}[\sigma\{\rho = 0\}] + f'(\infty)\,\mathrm{Tr}[\rho\{\sigma = 0\}] + \int_0^\infty \frac{f(\lambda)}{\lambda} \langle\sqrt{\rho}| d\xi^{\Delta_{\rho,\sigma}}(\lambda) |\sqrt{\rho}\rangle \,. \tag{5.3.19}$$

where $\{\sigma = 0\}$ is the projector onto the kernel of σ, $f'(\infty) \coloneqq \lim_{t\to\infty} \frac{f(t)}{t}$, and $f(0) \coloneqq \lim_{t\to 0} f(t)$. In this thesis, we will consider functions f for which always $f(0) = 0$, but often $f'(\infty) = \infty$. In that case, $\rho \ll \sigma$ (i.e. the support of ρ being contained in the support of σ) is a necessary condition for $D_f(\rho\|\sigma) = \infty$.

[1] Strictly speaking these statements should be understood in terms of Radon-Nikodym derivatives, i.e. if we introduce the two real-valued (and positive) measures $d\mu(\lambda) = \left\langle\sqrt{\sigma}\right| d\xi^{\Delta_{\rho,\sigma}}(\lambda) \left|\sqrt{\sigma}\right\rangle$, $d\nu(\lambda) = \left\langle\sqrt{\rho}\right| d\xi^{\Delta_{\rho,\sigma}}(\lambda) \left|\sqrt{\rho}\right\rangle$, then $\frac{d\nu}{d\mu}(\lambda) = \lambda$.

5.3.2 The Quantum Relative Entropy

The Umegaki Quantum Relative Entropy in infinite dimensions is the standard f-divergence associated to the function $f(\lambda) = \lambda \log \lambda$. This is a convex (even operator convex) function with $f(0) = 0$ and $f'(\infty) = \infty$, thus we can write

$$D(\rho\|\sigma) := D_f(\rho\|\sigma) = \begin{cases} \langle\sqrt{\rho}| \log \Delta_{\rho,\sigma} |\sqrt{\rho}\rangle & \text{if } \rho \ll \sigma \\ \infty & \text{otherwise} \end{cases} \tag{5.3.20}$$

Note that the expression $\langle\sqrt{\rho}| \log \Delta_{\rho,\sigma} |\sqrt{\rho}\rangle$ need not be finite, and is hence somewhat formal, and should really be understood as the corresponding integral

$$\int_0^\infty \log(\lambda) \, \langle\sqrt{\rho}| \, d\xi^{\Delta_{\rho,\sigma}}(\lambda) \, |\sqrt{\rho}\rangle \tag{5.3.21}$$

which converges (possibly to ∞) due to (5.3.17) and the convexity of $f(\lambda) = \lambda \log(\lambda)$. We will use similar expressions below also for other f-divergences which should also strictly speaking be understood as their corresponding integrals.

As a standard f-divergence associated to an operator convex function f, the quantum relative entropy satisfies the data-processing inequality (Hiai 2019), more direct proofs of this statement can also be found in Lindblad (1975) and Müller-Hermes and Reeb (2017). We can relate it to the finite-dimensional quantum relative entropy as follows:

Lemma 5.3.1 (Martingale convergence, Hiai 2018, Theorem 4.1(v) and Theorem 4.5)**.** *Let $\{M_a\}_a$ be an increasing net of unital von Neumann subalgebras of $\mathcal{B}(\mathcal{H})$ such that $(\bigcup_a M_a)'' = \mathcal{B}(\mathcal{H})$. Then, for every $\rho, \sigma \in \mathscr{P}(\mathcal{H})$*

$$\lim_a D(\rho|_{M_a}\|\sigma|_{M_a}) \nearrow D(\rho\|\sigma) \, . \tag{5.3.22}$$

In particular, for every increasing sequence of projections $(e_n)_n$ in $\mathcal{B}(\mathcal{H})$ such that $e_n \xrightarrow{n\to\infty} \mathbb{1}$, one can construct such a sequence of subalgebras as $M_n = e_n \mathcal{B}(\mathcal{H}) \, e_n \oplus \mathbb{C}(\mathbb{1} - e_n)$.

In this sense we can approximate the quantum relative entropy by taking projections onto larger and larger finite-dimensional subspaces, and adding one additional matrix element for the normalization. This allows for fairly straightforward generalizations of some finite-dimensional results, such as the following, which is known as "almost-concavity".

Lemma 5.3.2. *Let $\sigma \in \mathcal{D}(\mathcal{H})$ be a density matrix, $\{\lambda_i\}$, $i = 1, \ldots, n$ be a normalized probability distribution, and $\rho_i \in \mathcal{D}(\mathcal{H})$, $i = 1, \ldots, n$ be a set of density matrices. Then,*

$$D\left(\sum_{i=1}^n \lambda_i \rho_i \middle\| \sigma\right) + H(\lambda) \geq \sum_{i=1}^n \lambda_i D(\rho_i\|\sigma) \tag{5.3.23}$$

where $H(\lambda)$ is the Shannon entropy of λ.

Proof. We thank Milán Mosonyi for pointing out that this can be seen as a straightforward consequence of the above martingale convergence property. Let $(e_n)_n$ be an increasing sequence of finite-dimensional

projections on $\mathcal{B}(\mathcal{H})$ such that $\lim_{n\to\infty} e_n = \mathbb{1}$. Then, with M_n constructed as in Lemma 5.3.1, $\rho_i^{(n)} := \rho_i|_{M_n} = e_n\rho_i e_n \oplus \operatorname{Tr}(\rho_i(\mathbb{1}-e_n))$, and $\sigma^{(n)} = \sigma|_{M_n} = e_n\sigma e_n \oplus \operatorname{Tr}(\sigma(\mathbb{1}-e_n))$ are finite-dimensional, and so the finite-dimensional version of "almost-concavity" (see e.g. Khatri and Wilde 2020) implies

$$D\left(\sum_{i=1}^n \lambda_i\rho_i^{(n)}\middle\|\sigma^{(n)}\right) + H(\lambda) \geq \sum_{i=1}^n \lambda_i D(\rho_i^{(n)}\|\sigma^{(n)}) \tag{5.3.24}$$

and the desired statement follows by taking limits $n\to\infty$ on both sides. □

5.3.3 The Petz–Rényi Relative Entropy

The function $f_\alpha(\lambda) = \lambda^\alpha$ is concave for $\alpha\in[0,1]$ and convex for $\alpha\in[1,\infty)$, with $f_\alpha(0) = 0$ for all α. Futhermore, $f'_\alpha(\infty) = 0$ if $\alpha\in[0,1)$, and $f'_\alpha(\infty) = \infty$ for $\alpha\in(1,\infty)$. Hence, we get (Hiai 2021)

$$D_{f_\alpha}(\rho\|\sigma) = \begin{cases} \left\|\Delta_{\rho,\sigma}^{\alpha/2}|\sqrt{\sigma}\rangle\right\|^2 & \text{if } \alpha\in[0,1], \text{ or } \left[\alpha\in(1,\infty) \text{ and } \rho\ll\sigma \text{ and } |\sqrt{\sigma}\rangle\in\mathcal{D}(\Delta_{\rho,\sigma}^{\alpha/2})\right] \\ \infty & \text{otherwise.}\end{cases} \tag{5.3.25}$$

where the domain condition corresponds to the finiteness of the integral in (5.3.18), and due to the concavity (and positivity) of f_α for $\alpha\in[0,1)$ it is not hard to see that the integral will always be finite in that case.

We can then define the *α-Petz-Rényi* divergence for $\alpha\in[0,1)\cup(1,\infty)$ as

$$D_\alpha(\rho\|\sigma) := \frac{1}{\alpha-1}\log D_{f_\alpha}(\rho\|\sigma)\,. \tag{5.3.26}$$

We also have the following further characterization, which can be seen as analogous to the finite-dimensional case

Lemma 5.3.3 (Hiai 2021, Theorem 3.6)**.** *If $0\leq\alpha<1$, then*

$$D_\alpha(\rho\|\sigma) = \frac{1}{\alpha-1}\log\operatorname{Tr}\left[\rho^\alpha\sigma^{1-\alpha}\right]. \tag{5.3.27}$$

If $\alpha>1$ and $\rho\ll\sigma$, then $\sqrt{\rho}\in\mathcal{D}(\Delta_{\rho,\sigma}^{\alpha/2})$ (and hence $D_\alpha(\rho\|\sigma)<\infty$) if and only if there exists a unique $\eta\in\mathcal{B}_2(\mathcal{H})$ with support and image in the support of σ, such that $\sqrt{\rho^\alpha} = \eta\sqrt{\sigma^{\alpha-1}}$, and then

$$D_\alpha(\rho\|\sigma) = \frac{1}{\alpha-1}\log\|\eta\|_2^2 = \frac{1}{\alpha-1}\log\operatorname{Tr}\left(\eta^\dagger\eta\right). \tag{5.3.28}$$

In particular, this also implies that the Petz–Rényi divergence is additive on tensor products. As shown in Hiai (2018), the Petz–Rényi divergence is monotonically increasing in α on $(0,1)\cup(1,\infty)$, and satisfies $\lim_{\alpha\uparrow1} D_\alpha(\rho\|\sigma) = D(\rho\|\sigma)$, and $\lim_{\alpha\downarrow1} D_\alpha(\rho\|\sigma) = D(\rho\|\sigma)$ if there exists an $\alpha_0>1$ such that $D_{\alpha_0}(\rho\|\sigma)<\infty$. Similar to the finite-dimensional case, D_α satisfies the data-processing inequality for (and only for) $\alpha\in[0,2]$, which corresponds to the range where f_α is either operator convex or operator concave (see e.g. Hiai 2018).

Continuity of the Petz-Rényi divergence in α

Since our finite-dimensional proof of a one-shot relation between adaptive and parallel strategies relied on the continuity estimate Lemma 2.3.1 of the Petz–Rényi divergence, we will also need a similar estimate in infinite dimensions.

Lemma 5.3.4. *Lemma 2.3.1 also holds in infinite dimensions. In particular, let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be any two states on a separable Hilbert space $\mathcal{H}$ and $\gamma \in (0, 1]$. Define*

$$c_\gamma(\rho\|\sigma) := \frac{1}{\gamma} \log \left(2^{\gamma D_{1+\gamma}(\rho\|\sigma)} + 2^{-\gamma D_{1-\gamma}(\rho\|\sigma)} + 1 \right) . \tag{5.3.29}$$

Then, for all $\gamma \in (0, 1]$ and $\delta \in (0, \frac{\gamma}{2}]$:

$$\begin{aligned} D_{1+\delta}(\rho\|\sigma) &\leq D(\rho\|\sigma) + \ln(2)\delta(c_\gamma(\rho\|\sigma))^2 && (5.3.30)\\ &\leq D(\rho\|\sigma) + \delta(c_\gamma(\rho\|\sigma))^2. && (5.3.31) \end{aligned}$$

Furthermore, if $D(\rho\|\sigma) < \infty$, then for all $\gamma \in (0, 1]$ and $\delta \in (0, \frac{\gamma}{2}]$

$$D_{1-\delta}(\rho\|\sigma) \geq D(\rho\|\sigma) - \ln(2)\delta(c_\gamma(\rho\|\sigma))^2 \cosh(\ln(2)\delta c_\gamma(\rho\|\sigma)) . \tag{5.3.32}$$

and for all $\delta \in (0, \frac{\log 3}{2c_\gamma(\rho\|\sigma)}]$:

$$\begin{aligned} D_{1-\delta}(\rho\|\sigma) &\geq D(\rho\|\sigma) - \ln(2) \cosh(\log(3)/2)\delta(c_\gamma(\rho\|\sigma))^2 && (5.3.33)\\ &\geq D(\rho\|\sigma) - \delta(c_\gamma(\rho\|\sigma))^2 . && (5.3.34) \end{aligned}$$

Proof. This generalization of the finite-dimensional proof Lemma 2.3.1 was found in joint discussions with Jan Kochanowski. Recall the definitions of the Umegaki relative entropy and the Petz-Rényi relative entropies via the relative modular operator as

$$\begin{aligned} D(\rho\|\sigma) &= \int_{(0,\infty)} \log \lambda \langle\sqrt{\rho}| d\xi^{\Delta_{\rho,\sigma}}(\lambda)|\sqrt{\rho}\rangle = \langle\sqrt{\rho}| \log \Delta_{\rho,\sigma}|\sqrt{\rho}\rangle, && (5.3.35)\\ D_\alpha(\rho\|\sigma) &= \frac{1}{\alpha - 1} \log \int_{(0,\infty)} \lambda^{\alpha-1} \langle\sqrt{\rho}| d\xi^{\Delta_{\rho,\sigma}}(\lambda)|\sqrt{\rho}\rangle = \frac{1}{\alpha - 1} \log\langle\sqrt{\rho}|\Delta_{\rho,\sigma}^{\alpha-1}|\sqrt{\rho}\rangle, && (5.3.36)\\ D_{1+\delta}(\rho\|\sigma) &= \frac{1}{\delta} \log \int_{(0,\infty)} \lambda^{\delta} \langle\sqrt{\rho}| d\xi^{\Delta_{\rho,\sigma}}(\lambda)|\sqrt{\rho}\rangle = \frac{1}{\delta} \log\langle\sqrt{\rho}|\Delta_{\rho,\sigma}^{\delta}|\sqrt{\rho}\rangle. && (5.3.37) \end{aligned}$$

To prove the lemma, we follow the proof of Lemma 2.3.1, essentially only replacing $X = \rho \otimes (\sigma^{-1})^T$ by $\Delta_{\rho,\sigma}$ and $|\phi\rangle = \sum_i \sqrt{\rho}|i\rangle \otimes |i\rangle \in \mathcal{H} \otimes \mathcal{H}$ with $|\sqrt{\rho}\rangle \in \mathcal{B}_2(\mathcal{H})$.

We write $t^\delta = 1 + \delta \ln(t) + r_\delta(t)$, where the first two summands are the Taylor-coefficients of t^δ when

expanding in δ around $\delta = 0$. Set $r_\delta(t) := t^\delta - \delta \ln(t) - 1$. Using $e^x \geq 1 + x$ we can upper bound it by

$$r_\delta(t) = t^\delta - \delta \ln(t) - 1 \leq t^\delta + e^{-\delta \ln(t)} - 2 = e^{\delta \ln(t)} + e^{-\delta \ln(t)} - 2 \tag{5.3.38}$$

$$= 2(\cosh(\delta \ln(t)) - 1) =: s_\delta(t). \tag{5.3.39}$$

It is easy to check that $s_\delta(t)$ is monotonically increasing in t on $[1, \infty)$, concave in t on $[3, \infty)$ if $\delta \leq \frac{1}{2}$, and satisfies $s_{-\delta}(t) = s_\delta(t) = s_{\gamma\delta}(t^{\frac{1}{\gamma}})$, see the proof of Lemma 2.3.1 for more detailed explanations of this. We get for $t > 0$ and $\gamma \in (0, 1]$

$$s_\delta(t) = s_{\frac{\delta}{\gamma}}(t^\gamma) \leq s_{\frac{\delta}{\gamma}}(t^\gamma + t^{-\gamma}) \leq s_{\frac{\delta}{\gamma}}(t^\gamma + t^{-\gamma} + 1). \tag{5.3.40}$$

It is easy to see that $t^\gamma + t^{-\gamma} + 1 \geq 3$ for all $t \in (0, \infty)$. Thus, we can use Jensens inequality to get

$$\langle \sqrt{\rho} | s_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle = \int_0^\infty s_\delta(\lambda) \, \mathrm{d}\mu_\rho(\lambda) \leq \int_0^\infty s_{\frac{\delta}{\gamma}}(\lambda^\gamma + \lambda^{-\gamma} + 1) \, \mathrm{d}\mu_\rho(\lambda) \tag{5.3.41}$$

$$\overset{\text{Jensen}}{\leq} s_{\frac{\delta}{\gamma}}\left(\int_0^\infty (\lambda^\gamma + \lambda^{-\gamma} + 1) \, \mathrm{d}\mu_\rho(\lambda) \right) = s_{\frac{\delta}{\gamma}}\left(2^{\gamma c_\gamma(\rho\|\sigma)} \right). \tag{5.3.42}$$

where we wrote $\mathrm{d}\mu_\rho(\lambda) \coloneqq \langle \sqrt{\rho} | \, \mathrm{d}\xi^{\Delta_{\rho,\sigma}}(\lambda) \, | \sqrt{\rho} \rangle$, where $\mathrm{d}\xi^{\Delta_{\rho,\sigma}}(\lambda)$ is the spectral measure of the relative modular operator $\Delta_{\rho,\sigma}$. Now using Taylor's theorem with the Lagrange remainder we can bound

$$s_\delta(t) = s_0(t) + \frac{d}{d\delta} s_\delta(t)|_{\delta=0} \delta + \frac{1}{2} \frac{d^2}{d\delta^2} s_\delta(t)|_{\delta=\xi} \delta^2 \tag{5.3.43}$$

$$= \delta^2 (\ln(t))^2 \cosh(\xi \ln(t)) \leq \delta^2 (\ln(t))^2 \cosh(\delta \ln(t)) \tag{5.3.44}$$

for all $t > 0$ and some $\xi \in (0, \delta)$, where we used that $s_0(t) = \frac{d}{d\delta} s_\delta(t)|_{\delta=0} = 0$. Hence,

$$\langle \sqrt{\rho} | s_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle \leq s_{\frac{\delta}{\gamma}}\left(2^{\gamma c_\gamma(\rho\|\sigma)} \right) \leq (\delta \ln(2) c_\gamma(\rho\|\sigma))^2 \cosh(\delta \ln(2) c_\gamma(\rho\|\sigma)). \tag{5.3.45}$$

We can now apply this to derive the upper bound in the Lemma. For $\delta > 0$ have

$$D_{1+\delta}(\rho\|\sigma) = \frac{1}{\delta} \log \langle \sqrt{\rho} | \Delta_{\rho,\sigma}^\delta | \sqrt{\rho} \rangle = \frac{1}{\delta} \log \langle \sqrt{\rho} | 1 + \delta \ln(2) \log \Delta_{\rho,\sigma} + r_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle \tag{5.3.46}$$

$$= \frac{1}{\delta} \log(1 + \delta \ln(2) D(\rho\|\sigma) + \langle \sqrt{\rho} | r_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle) \tag{5.3.47}$$

$$\leq \frac{1}{\delta} \log(1 + \delta \ln(2) D(\rho\|\sigma) + \langle \sqrt{\rho} | s_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle) \tag{5.3.48}$$

$$= \frac{1}{\delta} \log(1 + \delta \ln(2) D(\rho\|\sigma)) + \frac{1}{\delta} \log\left(1 + \frac{\langle \sqrt{\rho} | s_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle}{1 + \delta \ln(2) D(\rho\|\sigma)} \right) \tag{5.3.49}$$

$$\leq D(\rho\|\sigma) + \frac{1}{\delta} \log\left(1 + \langle \sqrt{\rho} | s_\delta(\Delta_{\rho,\sigma}) | \sqrt{\rho} \rangle \right) \tag{5.3.50}$$

$$\leq D(\rho\|\sigma) + \frac{1}{\delta} \log\left(1 + (\delta \ln(2) c_\gamma(\rho\|\sigma))^2 \cosh(\delta \ln(2) c_\gamma(\rho\|\sigma)) \right) \tag{5.3.51}$$

$$\leq D(\rho\|\sigma) + \delta \ln(2) (c_\gamma(\rho\|\sigma))^2, \tag{5.3.52}$$

where in the fifth line it was used that $\log(1+x) \leq \frac{x}{\ln(2)}$ and $\delta > 0$. The final inequality follows from the fact that $k \mapsto k^2 - \ln\big(1 + k^2 \cosh(k)\big)$ is monotonically increasing and hence positive (see again the proof of Lemma 2.3.1 for a proof of this monotonicity). For the lower bound in the Lemma, i.e. the case $\delta < 0$, a slightly different argument has to be applied.

$$D_{1+\delta}(\rho\|\sigma) = \frac{1}{\delta}\log\langle\sqrt{\rho}|\Delta_{\rho,\sigma}^{\delta}|\sqrt{\rho}\rangle = \frac{1}{\delta}\log\langle\sqrt{\rho}|1 + \delta\ln(2)\log\Delta_{\rho,\sigma} + r_\delta(\Delta_{\rho,\sigma})|\sqrt{\rho}\rangle \tag{5.3.53}$$

$$\geq \frac{1}{\delta}\log(1 + \delta\ln(2)D(\rho\|\sigma) + \langle\sqrt{\rho}|s_\delta(\Delta_{\rho,\sigma})|\sqrt{\rho}\rangle) \tag{5.3.54}$$

$$\geq D(\rho\|\sigma) + \frac{1}{\delta\ln(2)}\langle\sqrt{\rho}|s_\delta(\Delta_{\rho,\sigma})|\sqrt{\rho}\rangle \tag{5.3.55}$$

$$\geq D(\rho\|\sigma) + \delta\ln(2)(c_\gamma(\rho\|\sigma))^2\cosh(\delta\ln(2)c_\gamma(\rho\|\sigma)), \tag{5.3.56}$$

where in the second inequality again $\log(1+x) \leq \frac{x}{\ln(2)}$ was used. If now $|\delta| \leq \frac{\log(3)}{2c_\gamma(\rho\|\sigma)} \leq \frac{\gamma}{2}$, then $\ln(2)\cosh(\ln(3)/2) < 1$ and thus

$$D_{1+\delta}(\rho\|\sigma) \geq D(\rho\|\sigma) + \delta(c_\gamma(\rho\|\sigma))^2. \tag{5.3.57}$$

□

5.3.4 The Geometric Rényi Divergence

The most natural way to define the geometric Rényi divergence in infinite dimensions is by generalizing the reverse-test characterization (see Section 2.3.3). Similar to the finite-dimensional case, we call the divergence obtained from the axiomatic characterization the maximal Rényi divergence, and call it the geometric Rényi divergence only for $\alpha \in (0, 2]$. The infinite-dimensional definition of the maximal Rényi divergence then goes as follows:

For $\rho, \sigma \in \mathscr{P}(\mathcal{H})$, we say that (Γ, g, h, μ) is a reverse test of (ρ, σ), if (X, μ) is a σ-finite measure space, $g, h \in L^1(X, \mu)$ with $g, h \geq 0$, and Γ is a positive trace-preserving map $\Gamma : L^1(X, \mu) \to \mathcal{B}_1(\mathcal{H})$ such that $\Gamma(g) = \rho$, $\Gamma(h) = \sigma$. By Γ being positive, we mean that $\Gamma(f) \geq 0$ for any $f \geq 0$, and by Γ being trace-preserving we mean that $\mathrm{Tr}(\Gamma(f)) = \int f d\mu$. Throughout this section we will encounter some expressions of the form $h(\lambda)f(g(\lambda)/h(\lambda))$, where $h(\lambda)$ can potentially be zero, for which we will use the following conventions

$$0f\left(\frac{0}{0}\right) = 0, \qquad 0f\left(\frac{g(\lambda)}{0}\right) = g(\lambda)f'(\infty) = g(\lambda)\lim_{x\to\infty} x^{-1}f(x)\,. \tag{5.3.58}$$

Definition 5.3.5 (for equivalence to other definitions see Hiai (2019), Theorem 6.3)**.** In infinite dimensions, for $\rho, \sigma \in \mathscr{P}(\mathcal{H})$ and any $\alpha \in (1, \infty)$, the *maximal Rényi trace function* can be defined via the following optimization problem

$$\widehat{S}_\alpha(\rho\|\sigma) \coloneqq \min_{\Gamma,g,h,\mu}\{S_\alpha^\mu(g\|h) \mid (\Gamma, g, h, \mu) \text{ is a reverse test of } (\rho, \sigma)\} \tag{5.3.59}$$

where

$$S_\alpha^\mu(g\|h) = \int h \left(\frac{g}{h}\right)^\alpha d\mu = \int h(\lambda) \left(\frac{g(\lambda)}{h(\lambda)}\right)^\alpha d\mu(\lambda)\,. \tag{5.3.60}$$

For $\alpha \in (0,1)$, the minimization is replaced with a maximization

$$\widehat{S}_\alpha(\rho\|\sigma) \coloneqq \max_{\Gamma,g,h,\mu} \{S_\alpha^\mu(g\|h) \mid (\Gamma, g, h, \mu) \text{ is a reverse test of } (\rho,\sigma)\}\,. \tag{5.3.61}$$

The *maximal Rényi divergence* is then defined as

$$\widehat{D}_\alpha(\rho\|\sigma) \coloneqq \frac{1}{\alpha-1} \log \widehat{S}_\alpha(\rho\|\sigma)\,. \tag{5.3.62}$$

This can be seen as a special case of what is called a maximal f-divergence for $f(\lambda) = \lambda^\alpha$. A lot is known about these divergences when f is operator convex (e.g. if $\alpha \in [1,2]$, or by considering $-f$, if $\alpha \in (0,1]$), in which case the solution to the optimization problem can be explicitly characterized (Hiai 2019), and as stated above, we refer to the divergence as the geometric Rényi divergence in this range of $\alpha \in (0,2]$.

Our main theorem of this section is the following:

Theorem 5.3.6. *For all $\alpha \in (0,1) \cup (1,2]$, the geometric Rényi divergence satisfies the chain rule, i.e. for all states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ and any two channels $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{K})$ we have that*

$$\widehat{D}_\alpha(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) \leq \widehat{D}_\alpha(\rho\|\sigma) + \widehat{D}_\alpha(\mathcal{E}\|\mathcal{F})\,. \tag{5.3.63}$$

This also directly implies the additivity of the geometric Rényi channel divergence, i.e. for all channels $\mathcal{E}_1, \mathcal{F}_1 : \mathscr{P}(\mathcal{H}_1) \to \mathscr{P}(\mathcal{H})(\mathcal{K}_1)$, $\mathcal{E}_2, \mathcal{F}_2 : \mathscr{P}(\mathcal{H}_2) \to \mathscr{P}(\mathcal{K}_2)$ it holds that

$$\widehat{D}_\alpha(\mathcal{E}_1 \otimes \mathcal{E}_2\|\mathcal{F}_1 \otimes \mathcal{F}_2) = \widehat{D}_\alpha(\mathcal{E}_1\|\mathcal{F}_1) + \widehat{D}_\alpha(\mathcal{E}_2\|\mathcal{F}_2). \tag{5.3.64}$$

We show this by proving that at least for some restricted set of states, the optimization in the reverse tests can be restricted to probability distributions on a countable set, rather than L^1 functions (although the optimum might not necessarily be achieved anymore in this case), which then allows us to adapt the finite-dimensional chain rule proof of Berta and Tomamichel (2022) to show the desired result. To simplify the argument also in the case where $\alpha \in (0,1)$, we show this first step in the slightly more general setting of maximal f-divergences for an operator convex function $f : [0,\infty) \to \mathbb{R}$, where such maximal f-divergences are defined as follows:

$$\widehat{S}_f(\rho\|\sigma) \coloneqq \min_{\Gamma,g,h,\mu} \{S_f^\mu(g\|h) \mid (\Gamma, g, h, \mu) \text{ is a reverse test of } (\rho,\sigma)\} \tag{5.3.65}$$

where

$$S_f^\mu(g\|h) = \int h f\left(\frac{g}{h}\right) d\mu = \int h(\lambda) f\left(\frac{g(\lambda)}{h(\lambda)}\right) d\mu(\lambda)\,. \tag{5.3.66}$$

Furthermore, for $\rho, \sigma \in \mathscr{P}(\mathcal{H})$, we say (Γ, p, q) is a *discrete* reverse test for (ρ,σ), if $p, q \in \ell_1$, $p, q \geq 0$, and $\Gamma : \ell_1 \to \mathcal{B}(\mathcal{H})$ is a linear, positive, and trace-preserving (in the sense of $\operatorname{Tr}(\Gamma(r)) = \sum_i r_i$ for all $r \in \ell_1$)

map such that $\Gamma(p) = \rho, \Gamma(q) = \sigma$.

Lemma 5.3.7. *Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be such that $\exists c \geq 0$ s.t. $\rho \leq c\sigma \leq c^2\rho$. Let f be a continuous operator convex function on $[0, \infty)$. Then, the maximal f divergence can expressed as the following optimization problem:*

$$\widehat{S}_f(\rho\|\sigma) = \inf_{\Gamma,p,q} \{S_f(p\|q) \mid (\Gamma, p, q) \text{ is a discrete reverse test of } (\rho, \sigma)\} \tag{5.3.67}$$

where

$$S_f(p\|q) \coloneqq \sum_i q_i f\left(\frac{p_i}{q_i}\right). \tag{5.3.68}$$

Proof. For any measure μ on $[0, 1]$, we say (Γ, g, h, μ) is a piecewise reverse test of (ρ, σ) if (Γ, g, h, μ) is a reverse test, and additionally there exists a countable partition of $[0, 1]$ into disjoint μ-measurable sets $\{A_i\}$ of non-zero measure, such that

$$g = \sum_{i=1}^{\infty} g^{(i)} 1_{A_i} \qquad h = \sum_{i=1}^{\infty} h^{(i)} 1_{A_i}. \tag{5.3.69}$$

where the $g^{(i)}$ and $h^{(i)}$ are constants. We first show the following statement:

$$\widehat{S}_f(\rho\|\sigma) = \inf_{\Gamma,g,h,\mu} \{S_f^{\mu}(g\|h) \mid (\Gamma, g, h, \mu) \text{ is a piecewise reverse test of } (\rho, \sigma)\} \tag{5.3.70}$$

Since any piecewise reverse test is also a reverse test, we only have to show that there exists a sequence of piecewise reverse tests that converges to the optimum value. Additionally, we can restrict to $\widehat{S}_f(\rho\|\sigma) < \infty$, since otherwise we know that no reverse test can achieve a finite value, and we are done. By Hiai (2019), Theorem 6.3, since the function f is operator convex, we can further restrict to the case where the measure space X is [0,1], and furthermore the optimum reverse test can be chosen as $g(t) = t$, $h(t) = 1 - t$, $t \in [0, 1]$, for some suitable Γ and μ (see Hiai (2019) for the exact expression of Γ and μ).

For $n \geq 2$, consider the following piecewise aproximation g_n of $g(t) = t$:

$$g_n(t) = \sum_{k=1}^{\infty} \frac{1}{n(k+1)} 1_{(\frac{1}{n(k+1)}, \frac{1}{nk}]}(t) + \sum_{i=1}^{n-1} \frac{i}{n} 1_{(i/n,(i+1)/n]}(t) \tag{5.3.71}$$

This satisfies the following properties:

1. $|g_n(t) - g(t)| \leq \frac{1}{n} \qquad \forall t \in [0, 1]$
2. $\frac{1}{2} g(t) \leq g_n(t) \leq g(t) \qquad \forall t \in [0, 1]$.

As $h(t) = g(1 - t)$, we then set $h_n(t) \coloneqq g_n(1 - t)$. See Figure 5.1 for an illustration of $g_n(t)$.

The motivation for this approximation is that we want g_n and h_n to approximate g and h from below, but we also would like g_n/h_n to be well-defined everywhere except for $t = 1$ (where also g/h is infinite), hence the need for an infinite series at least in the definition of h_n. Note that the measurability of these functions follows from the measurability of g and h, and (since $n \geq 2$) we also have $\frac{g_n}{h_n} \leq 2\frac{g}{h}$. These approximations

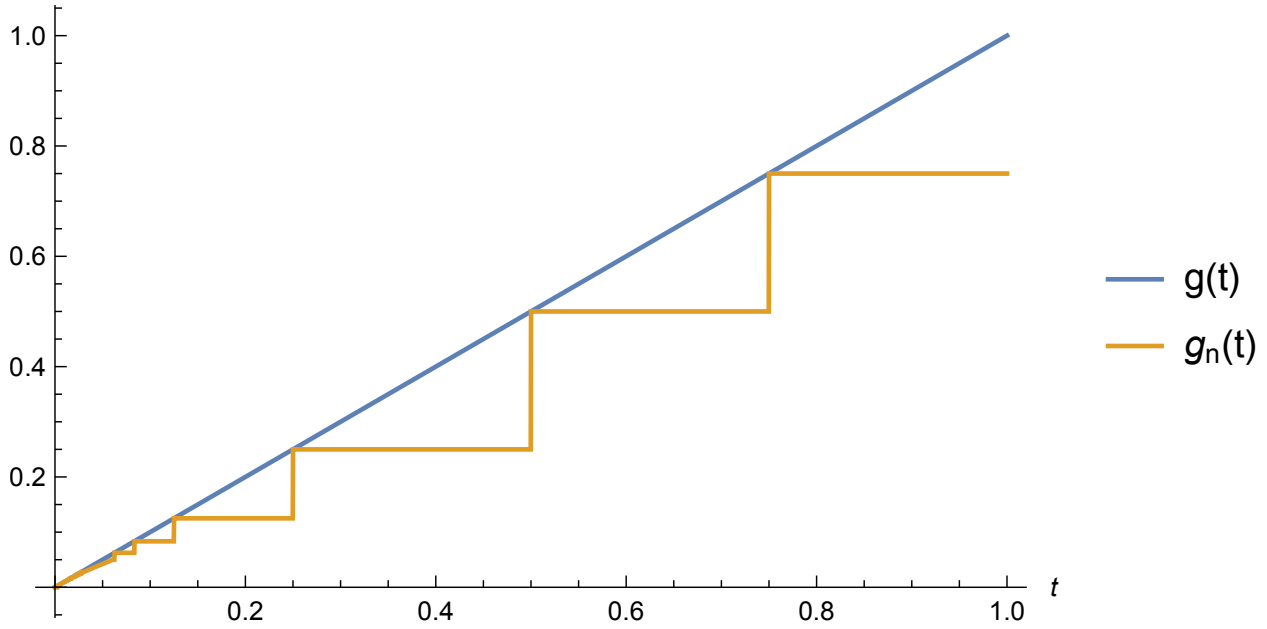

Figure 5.1: Illustration of the piecewise approximation $g_n(t)$ of $g(t) = t$.

will no longer form a reverse test, since for example in general $\Gamma(g_n) < \rho$, but the goal is now to modify these further to construct a $(\bar{\Gamma}_n, \bar{g}_n, \bar{h}_n)$ that does form a piecewise reverse test. Define

$$\Delta_g^{(n)} := \Gamma(g) - \Gamma(g_n) \tag{5.3.72}$$

$$\Delta_h^{(n)} := \Gamma(h) - \Gamma(h_n)\,. \tag{5.3.73}$$

By the linearity and positivity of Γ, we have that these are positive operators. Moreover,

$$\Delta_g^{(n)} = \Gamma(g - g_n) \leq \Gamma\left(\frac{1}{n} 1_{[0,1]}\right) = \frac{1}{n}\Gamma(g+h) = \frac{1}{n}(\rho+\sigma) \leq \frac{1+c}{n}\sigma\,. \tag{5.3.74}$$

and similarly for $\Delta_h^{(n)}$ (remember that $g(t) = t$, $h(t) = 1 - t$). Subsequently we will want to ensure that $\Delta_g^{(n)} \leq \Delta_h^{(n)}$. This will generally not be satisfied, but we can achieve this by replacing h_n with $\tilde{h}_n = \left(1 - \frac{1}{\sqrt{n}}\right) h_n$. This is still a piecewise approximiation of h that converges pointwise in the limit $n \to \infty$, but also

$$\Delta_{\tilde{h}}^{(n)} := \Gamma(h - \tilde{h}_n) = \left(1 - \frac{1}{\sqrt{n}}\right)\Gamma(h - h_n) + \frac{1}{\sqrt{n}}\Gamma(h) = \left(1 - \frac{1}{\sqrt{n}}\right)\Delta_h^{(n)} + \frac{1}{\sqrt{n}}\sigma \geq \frac{1}{\sqrt{n}}\sigma \geq \frac{\sqrt{n}}{1+c}\Delta_g^{(n)}\,. \tag{5.3.75}$$

For n large enough, $\frac{\sqrt{n}}{1+c} \geq 1$ and then $\Delta_g^{(n)} \leq \Delta_{\tilde{h}}^{(n)}$. We will construct our piecewise reverse test by extending g_n and h_n to functions on $[0, 2]$ and setting suitable values on $(1, 2]$. As a measure on $[0, 2]$, we choose $\nu := \mu \oplus \lambda$, where λ is the Lebesgue measure on $(1, 2]$, and we mean by the notation that when

integrating functions w.r.t. ν, we integrate with μ over $[0,1]$ and with λ over $(1,2]$, i.e.

$$\int_0^2 f d\nu = \int_0^1 f d\mu + \int_1^2 f d\lambda \,. \tag{5.3.76}$$

Let

$$x \coloneqq \frac{\operatorname{Tr}\left(\Delta_g^{(n)}\right)}{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)} \in [0,1] \tag{5.3.77}$$

and if $x < 1$, define the state ω by:

$$(1-x)\omega \coloneqq \frac{\Delta_{\tilde{h}}^{(n)}}{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)} - x\frac{\Delta_g^{(n)}}{\operatorname{Tr}\left(\Delta_g^{(n)}\right)} = \frac{1}{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)}(\Delta_{\tilde{h}}^{(n)} - \Delta_g^{(n)}) \geq 0 \tag{5.3.78}$$

which satisfies $\operatorname{Tr}(\omega) = 1$. We employ the convention $\frac{0}{0} = 0$, i.e. if $\Delta_g^{(n)} = 0$, the term $(\Delta_g^{(n)} / \operatorname{Tr}\left(\Delta_g^{(n)}\right))$ is zero. If $x = 1$, ω turns out to be irrelevant, so we can just pick any $\omega \in \mathcal{D}(\mathcal{H})$. Define further the following two normalized functions on $[0,2]$:

$$\bar{g}_n = g_n 1_{[0,1]} + \frac{\operatorname{Tr}\left(\Delta_g^{(n)}\right)}{x} 1_{(1,1+x]} = g_n 1_{[0,1]} + \operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) 1_{(1,1+x]} \tag{5.3.79}$$

$$\bar{h}_n = \tilde{h}_n 1_{[0,1]} + \operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) 1_{(1,2]} \tag{5.3.80}$$

and the following map from L^1 functions on $[0,2]$ to bounded operators

$$\bar{\Gamma}_n(f) \coloneqq \Gamma(f 1_{[0,1]}) + \frac{\Delta_g^{(n)}}{\operatorname{Tr}\left(\Delta_g^{(n)}\right)} \int_1^{1+x} f d\lambda + \omega \int_{1+x}^2 f d\lambda \,. \tag{5.3.81}$$

This is positive, in the sense that $\bar{\Gamma}_n(f) \geq 0$ if $f \geq 0$, and normalization-preserving in the sense that

$$\operatorname{Tr}\big(\bar{\Gamma}_n(f)\big) = \int_0^2 f d\nu \,. \tag{5.3.82}$$

It also satisfies

$$\bar{\Gamma}_n(\bar{g}_n) = \Gamma(g_n) + \Delta_g^{(n)} x \frac{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)}{\operatorname{Tr}\left(\Delta_g^{(n)}\right)} = \Gamma(g_n) + \Delta_g^{(n)} = \rho \tag{5.3.83}$$

$$\bar{\Gamma}_n(\bar{h}_n) = \Gamma(\tilde{h}_n) + \Delta_g^{(n)} x \frac{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)}{\operatorname{Tr}\left(\Delta_g^{(n)}\right)} + (1-x)\omega \operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) = \Gamma(\tilde{h}_n) + \operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) \frac{\Delta_{\tilde{h}}^{(n)}}{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)} = \sigma \tag{5.3.84}$$

and hence $(\bar{\Gamma}_n, \bar{g}_n, \bar{h}_n, \nu)$ is a piecewise reverse test of (ρ, σ). It remains to show that

$$S_f^{\nu}(\bar{g}_n \| \bar{h}_n) \xrightarrow{n\to\infty} S_f^{\mu}(g\|h)\,. \tag{5.3.85}$$

On $(1, 2]$, we have

$$\int_1^2 \left|\bar{h}_n f\left(\frac{\bar{g}_n}{\bar{h}_n}\right)\right| d\lambda = \int_1^2 \operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) \left| f\left(\frac{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) 1_{(1,1+x]}}{\operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right)}\right)\right| d\lambda \tag{5.3.86}$$

$$\leq C \operatorname{Tr}\left(\Delta_{\tilde{h}}^{(n)}\right) = C \int_0^1 (h - \tilde{h}_n) d\mu \to 0 \tag{5.3.87}$$

by monotone convergence, where $C = \max\{|f(0)|\,, |f(1)|\} < \infty$.

For the bound on $[0, 1]$, we will now first show some properties of f. For this, note that it is quite immediate from the definition of $\widehat{S}_f$ (see also e.g. Hiai (2019), Example 2.14) that for $s(x) = a + bx$

$$\widehat{S}_{f+s}(\rho\|\sigma) = \widehat{S}_f(\rho\|\sigma) + a \operatorname{Tr}(\sigma) + b \operatorname{Tr}(\rho) = \widehat{S}_f(\rho\|\sigma) + a + b \tag{5.3.88}$$

and similarly for S_f. Hence, one easily sees that showing our claim (5.3.67) for any such shifted $f + s$ is sufficient, since it directly implies the claim for the original f. Since f was assumed to be continuous and convex, in particular its epigraph is a closed convex set, and so by the supporting hyperplane theorem we can pick an s such that $(f + s)(x) \geq 0$ for all $x \in [0, \infty)$. In particular, for any point $x_0 \in (0, \infty)$ there exists such an s that additionally achieves $(f + s)(x_0) = 0$ (note that we cannot pick $x_0 = 0$ as the tangent there might be vertical). Thus, without loss of generality $f \geq 0$, $f(1) = 0$. This then also implies that f is monotonically increasing for $x \geq 1$, since with $1 < a < b$

$$f(b) \geq \frac{a-1}{b-1} f(b) = \frac{a-1}{b-1} f(b) + \left(1 - \frac{a-1}{b-1}\right) f(1) \tag{5.3.89}$$

$$\geq f\left(\frac{a-1}{b-1} b + 1 - \frac{a-1}{b-1}\right) = f\left(\frac{a-1}{b-1}(b-1) + 1\right) = f(a)\,. \tag{5.3.90}$$

Equivalently, f is monotonically decreasing for $x \leq 1$, as with $a < b < 1$ we have $1 - a > 1 - b$, and hence by the same argument

$$f(a) \geq \frac{1-b}{1-a} f(a) = \frac{1-b}{1-a} f(a) + \left(1 - \frac{1-b}{1-a}\right) f(1) \tag{5.3.91}$$

$$\geq f\left(\frac{1-b}{1-a} a + 1 - \frac{1-b}{1-a}\right) = f\left(\frac{1-b}{1-a}(a-1) + 1\right) = f(b)\,. \tag{5.3.92}$$

Hence, $f(x) \leq f(0)$ for all $x \in [0, 1]$, and thus, combining this with the monotonicity for $x \geq 1$, we get that for all $0 \leq a \leq b$

$$f(a) \leq f(b) + f(0)\,. \tag{5.3.93}$$

Additionally, we want to show that there exist constants $A, B \geq 0$ such that for every $\lambda \geq 1$ and $x \in [0, \infty)$,

$f(\lambda x) \leq \lambda^2(f(x) + Ax + B)$. For this we will use the operator convexity of f. It is well-known (see e.g. Hiai (2018) and references therein) that every operator convex function has the following integral representation on $x \in [0, \infty)$:

$$f(x) = a + b(x-1) + c(x-1)^2 + \int_{[0,\infty)} \frac{(x-1)^2}{x+s} d\xi(s)\,, \tag{5.3.94}$$

with $a, b \in \mathbb{R}$, $c \geq 0$, and ξ is a positive measure on $[0, \infty)$. Since we assumed $f(0) < \infty$ we know that

$$\int_{[0,\infty)} s^{-1} d\xi(s) =: D < \infty\,. \tag{5.3.95}$$

For all $\lambda \geq 1$, using $(\lambda x - 1)^2 = \lambda^2(x-1)^2 + 2\lambda(\lambda-1)x + (1-\lambda^2) \leq \lambda^2(x-1)^2 + 2\lambda^2 x$, and $(\lambda x + s)^{-1} \leq (x+s)^{-1}$ we then find:

$$f(\lambda x) \leq \lambda^2 f(x) + a(1-\lambda^2) + b(x(\lambda - \lambda^2) + (\lambda^2 - 1)) + c2\lambda^2 x + \int_{[0,\infty)} \frac{2\lambda^2 x}{\lambda x + s} d\xi(s) \tag{5.3.96}$$

$$\leq \lambda^2 f(x) + \lambda^2 |a| + \lambda^2 |b| (x+1) + c2\lambda^2 x + 2\lambda^2 D x \tag{5.3.97}$$

from which one can infer the existence of the desired constants A and B.

If we assume $n \geq 4$, then we have on $[0,1]$ that

$$\frac{g_n}{\tilde{h}_n} \leq \frac{2}{1 - \frac{1}{\sqrt{n}}} \frac{g}{h} \leq 4\frac{g}{h}\,, \tag{5.3.98}$$

and so, using the positivity of f as well as the two previous bounds, we see that

$$\left| \tilde{h}_n f\left(\frac{g_n}{\tilde{h}_n}\right) \right| = \tilde{h}_n f\left(\frac{g_n}{\tilde{h}_n}\right) \leq h\left[f\left(4\frac{g}{h}\right) + f(0)\right] \leq h4^2\left[f\left(\frac{g}{h}\right) + A\frac{g}{h} + B + f(0)\right]. \tag{5.3.99}$$

Integrating the right-hand side with respect to μ gives a finite value, since h and $gh^{-1}h = g$ are normalized (cancelling $h^{-1}h$ here is justified by (5.3.58) with $f(x) = x$), and $\widehat{S}_f(\rho\|\sigma) = S_f^\mu(g\|h) < \infty$. Note also that $\tilde{h}_n(t) > 0$ for all $t \in [0,1)$, and hence the continuity of f implies that for all $t \in [0,1)$

$$\tilde{h}_n(t) f\left(\frac{g_n(t)}{\tilde{h}_n(t)}\right) \xrightarrow{n\to\infty} h(t) f\left(\frac{g(t)}{h(t)}\right). \tag{5.3.100}$$

Additionally, at $t = 1$ we have

$$\tilde{h}_n(1) f\left(\frac{g_n(1)}{\tilde{h}_n(1)}\right) = g_n(1) f'(\infty) \xrightarrow{n\to\infty} g(1) f'(\infty) = h(1) f\left(\frac{g(1)}{h(1)}\right). \tag{5.3.101}$$

Hence, by dominated convergence we get

$$\int_0^1 \tilde{h}_n f\left(\frac{g_n}{\tilde{h}_n}\right) d\mu \to \int_0^1 h f\left(\frac{g}{h}\right) d\mu = S_f^\mu(g\|h)\,. \tag{5.3.102}$$

This completes the proof of (5.3.70). Finally, we still have to show that the formulation in terms of piecewise reverse tests is equivalent to the formulation in terms of discrete reverse tests. Let (Γ, p, q) be a discrete reverse test of (ρ, σ). Then, let μ be the Lebesgue measure on $[0,1]$, and $\{A_i\}_{i=1}^{\infty}$ be a countable collection of disjoint measurable subsets of $[0,1]$, such that $\mu(A_i) > 0$, for all i and $\sum_i \mu(A_i) = 1$. Define $g' \coloneqq \sum_i \frac{1_{A_i}}{\mu(A_i)} p_i$, $h' \coloneqq \sum_i \frac{1_{A_i}}{\mu(A_i)} q_i$ and $\Gamma'(f) \coloneqq \sum_i \Gamma_i \int_{A_i} f d\mu$ for any $f \in L^1$. Then, $\Gamma'(g') = \Gamma(p) = \rho$ and $\Gamma'(h') = \Gamma(q) = \sigma$, and so (Γ', g', h', μ) is a piecewise reverse test, and additionally also $S_f(p\|q) = S_f^{\mu}(g'\|h')$. Conversely, if (Γ', g', h', μ) is a piecewise reverse test as in (5.3.69), then define $p, q \in \ell_1$ by $p_i \coloneqq g'^{(i)}\mu(A_i)$, $q_i \coloneqq h'^{(i)}\mu(A_i)$, and $\Gamma_i \coloneqq \frac{1}{\mu(A_i)}\Gamma'(1_{A_i})$, which defines a discrete reverse test again achieving the same value in the optimization problem.

□

Lemma 5.3.8. *For any $\alpha \in (0,\infty)$, any states $\rho, \sigma \in \mathscr{P}(\mathcal{H}_A)$, and any two positive real numbers p, q the geometric trace function satisfies*

$$\widehat{S}_\alpha(p\rho\|q\sigma) = p^\alpha q^{1-\alpha}\widehat{S}_\alpha(\rho\|\sigma). \tag{5.3.103}$$

Proof. For any two functions g and h and any measure μ, it is easy to see that

$$\widehat{S}_\alpha^\mu(pg\|qh) = p^\alpha q^{1-\alpha}\widehat{S}_\alpha^\mu(g\|h)\,. \tag{5.3.104}$$

The statement then follows from the variational expression of $\widehat{S}_\alpha$ and the fact that if (Γ, g, h, μ) forms a reverse test of (ρ, σ), then (Γ, pg, ph, μ) form a reverse test of $(p\rho, q\sigma)$ and vice versa. □

Proof of Theorem 5.3.6. Let $\rho = \rho_A, \sigma = \sigma_A \in \mathcal{D}(\mathcal{H}_A)$ and fix $\varepsilon > 0$. Then, define $\tilde{\rho}_A = \rho_A + \varepsilon(\rho_A + \sigma_A)$, $\tilde{\sigma}_A = \sigma_A + \varepsilon(\rho_A + \sigma_A)$. Let (Γ, p, q) be a discrete reverse test of $(\tilde{\rho}_A, \tilde{\sigma}_A)$. We write $\Gamma_i = \Gamma(1_i) \in \mathcal{B}_1$, where $1_i \in \ell_1$ is the sequence that is one at index i and zero otherwise. We then have

$$\tilde{\rho}_A = \sum_i p_i \Gamma_i \qquad \tilde{\sigma}_A = \sum_i q_i \Gamma_i\,. \tag{5.3.105}$$

Hence, taking R an additional infinite-dimensional system with countable basis $\{|i\rangle\}_i$, we can define

$$\tilde{\rho}_{RA} = \sum_i p_i \left|i\right\rangle\!\left\langle i\right|_R \otimes \Gamma_i \qquad \tilde{\sigma}_{RA} = \sum_i q_i \left|i\right\rangle\!\left\langle i\right|_R \otimes \Gamma_i \tag{5.3.106}$$

such that $\tilde{\rho}_A = \mathrm{Tr}_R(\tilde{\rho}_{RA}), \tilde{\sigma}_A = \mathrm{Tr}_R(\tilde{\sigma}_{RA})$. We start with the case where $\alpha \in (1,2]$, for which it is

well-known that the function $f(x) = x^\alpha$ is operator convex. We have,

$$\begin{aligned}
\widehat{S}_\alpha(\mathcal{E}(\tilde{\rho}_A)\|\mathcal{F}(\tilde{\sigma}_A)) &\leq \widehat{S}_\alpha(\mathcal{E}(\tilde{\rho}_{RA})\|\mathcal{F}(\tilde{\sigma}_{RA})) && (5.3.107)\\
&= \widehat{S}_\alpha\left(\sum_i p_i \left|i\right\rangle\!\left\langle i\right| \otimes \mathcal{E}(\Gamma_i)\| \sum_i q_i \left|i\right\rangle\!\left\langle i\right| \otimes \mathcal{F}(\Gamma_i)\right) && (5.3.108)\\
&= \sum_i \widehat{S}_\alpha(p_i\mathcal{E}(\Gamma_i)\|q_i\mathcal{F}(\Gamma_i)) && (5.3.109)\\
&= \sum_i p_i^\alpha q_i^{1-\alpha} \widehat{S}_\alpha(\mathcal{E}(\Gamma_i)\|\mathcal{F}(\Gamma_i)) && (5.3.110)\\
&\leq \sum_i p_i^\alpha q_i^{1-\alpha} \sup_{\nu\in\mathcal{D}(A)} \widehat{S}_\alpha(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) && (5.3.111)\\
&= \widehat{S}_\alpha(p\|q) \sup_{\nu\in\mathcal{D}(A)} \widehat{S}_\alpha(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) && (5.3.112)
\end{aligned}$$

where the first inequality follows from the data-processing inequality for the geometric trace function (Hiai 2019, Theorem 2.9) (remember that the channels act only on system A), the third line follows from Lemma 5.3.13, and the fourth equality is Lemma 5.3.8.

We write $\widehat{S}_\alpha(\mathcal{E}\|\mathcal{F}) = \sup_{\nu\in\mathcal{D}(A)} \widehat{S}_\alpha(\mathcal{E}(\nu)\|\mathcal{F}(\nu))$. Taking the infimum over all discrete reverse tests, we find by Lemma 5.3.7 that

$$\widehat{S}_\alpha(\mathcal{E}(\tilde{\rho})\|\mathcal{F}(\tilde{\sigma})) \leq \widehat{S}_\alpha(\tilde{\rho}\|\tilde{\sigma})\widehat{S}_\alpha(\mathcal{E}\|\mathcal{F})\,. \tag{5.3.113}$$

It remains to show convergence in the limit $\varepsilon \to 0$. For the right-hand side, by Hiai (2019), Definition 2.8

$$\lim_{\varepsilon\to 0} \widehat{S}_\alpha(\tilde{\rho}\|\tilde{\sigma}) = \widehat{S}_\alpha(\rho\|\sigma)\,. \tag{5.3.114}$$

For the left-hand side, by Hiai (2019), Theorem 2.9, we have

$$\begin{aligned}
\widehat{S}_\alpha(\mathcal{E}(\tilde{\rho})\|\mathcal{F}(\tilde{\sigma})) &= \widehat{S}_\alpha(\mathcal{E}(\rho) + \varepsilon\mathcal{E}(\rho+\sigma))\|\mathcal{F}(\sigma) + \varepsilon\mathcal{F}(\rho+\sigma)) && (5.3.115)\\
&\leq \widehat{S}_\alpha(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) + \varepsilon\widehat{S}_\alpha(\mathcal{E}(\rho+\sigma)\|\mathcal{F}(\rho+\sigma)) && (5.3.116)\\
&\leq \widehat{S}_\alpha(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) + \varepsilon\widehat{S}_\alpha(\mathcal{E}\|\mathcal{F}). && (5.3.117)
\end{aligned}$$

Now, if $\widehat{S}_\alpha(\mathcal{E}\|\mathcal{F}) = \infty$ the statement of our theorem is empty, so we can assume it to be finite. In that case we get

$$\limsup_{\varepsilon\to 0} \widehat{S}_\alpha(\mathcal{E}(\tilde{\rho})\|\mathcal{F}(\tilde{\sigma})) \leq \widehat{S}_\alpha(\mathcal{E}(\rho)\|\mathcal{F}(\sigma))\,. \tag{5.3.118}$$

The opposite direction

$$\liminf_{\varepsilon\to 0} \widehat{S}_\alpha(\mathcal{E}(\tilde{\rho})\|\mathcal{F}(\tilde{\sigma})) \geq \widehat{S}_\alpha(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) \tag{5.3.119}$$

follows by the lower semi-continuity of $\widehat{S}_\alpha$ (Hiai 2019, Theorem 5.5). The statement then follows upon taking the logarithm and dividing by $\alpha - 1$ (which is positive).

For $\alpha \in [0, 1)$, $f(x) = -x^\alpha$ is well-known to be operator convex, and in that case we can apply all the above reasoning to $\widehat{S}_f = -\widehat{S}_\alpha$ (all the properties from Hiai (2019) we used apply to $\widehat{S}_f$ with an operator

convex function f). We then find that

$$\widehat{S}_f(\mathcal{E}(\rho_A)\|\mathcal{F}(\sigma_A)) \leq \widehat{S}_\alpha(\rho\|\sigma) \sup_{\nu\in\mathcal{D}(A)} \widehat{S}_f(\mathcal{E}(\nu)\|\mathcal{F}(\nu)) \tag{5.3.120}$$

which corresponds to

$$\widehat{S}_\alpha(\mathcal{E}(\rho_A)\|\mathcal{F}(\sigma_A)) \geq \widehat{S}_\alpha(\rho\|\sigma) \inf_{\nu\in\mathcal{D}(A)} \widehat{S}_\alpha(\mathcal{E}(\nu)\|\mathcal{F}(\nu))\,. \tag{5.3.121}$$

The desired statement follows again after taking logarithms and dividing by $\alpha - 1$ (which is now negative, so it turns around the inequality and changes the infimum into a supremum).

To see how the chain rule implies additivity, let $\mathcal{E}_1, \mathcal{F}_1 : \mathscr{P}(\mathcal{H}_{A_1}) \to \mathscr{P}(\mathcal{H}_{B_1})$, $\mathcal{E}_2, \mathcal{F}_2 : \mathscr{P}(\mathcal{H}_{A_2}) \to \mathscr{P}(\mathcal{H}_{B_2})$ be channels, and consider any joint input state $\nu = \nu_{RA_1A_2}$. Then, (supressing identities as before)

$$\begin{aligned}
\widehat{D}_\alpha\big((\mathcal{E}_1\otimes\mathcal{E}_2)(\nu)\|(\mathcal{F}_1\otimes\mathcal{F}_2)(\nu)\big) &= \widehat{D}_\alpha\big(\mathcal{E}_1(\mathcal{E}_2(\nu))\|\mathcal{F}_1(\mathcal{F}_2(\nu)\big) && (5.3.122)\\
&\leq \widehat{D}_\alpha(\mathcal{E}_1\|\mathcal{F}_1) + \widehat{D}_\alpha(\mathcal{E}_2(\nu)\|\mathcal{F}_2(\nu)) && (5.3.123)\\
&\leq \widehat{D}_\alpha(\mathcal{E}_1\|\mathcal{F}_1) + \widehat{D}_\alpha(\mathcal{E}_2\|\mathcal{F}_2) && (5.3.124)
\end{aligned}$$

where we used the chain rule in the first inequality. This implies

$$\widehat{D}_\alpha(\mathcal{E}_1\otimes\mathcal{E}_2\|\mathcal{F}_1\otimes\mathcal{F}_2) \leq \widehat{D}_\alpha(\mathcal{E}_1\|\mathcal{F}_1) + \widehat{D}_\alpha(\mathcal{E}_2\|\mathcal{F}_2), \tag{5.3.125}$$

and the other direction follows by just restricting the supremum over input states $\nu_{RA_1A_2}$ to product states. □

Remark 5.3.9 (Chain Rule for the Belavkin-Staszewski Relative Entropy)**.** By using a very similar argument as in Theorem 5.3.6, one can also show that the Belavkin-Staszewski relative entropy, which is defined as the maximal f-divergence for the operator convex function $f(x) = x\log x$, satisfies the chain rule. To see this, note that now instead of Lemma 5.3.8, we have $\sum_i \widehat{S}_f(p_i\rho_i\|q_i\sigma_i) = \widehat{S}_f(p\|q) + \sum_i p_i\widehat{S}_f(\rho_i\|\sigma_i)$, and the p_i are normalized. Using also an adaptation of Lemma 5.3.13, the remainder of the argument is then almost identical to Theorem 5.3.6.

Additional Properties of the Geometric Rényi Divergence

Below are some additional properties of the geometric Rényi divergence, some of which we have already used above, and some of which we will also make use of later on. We write $s(\rho)$ for the support of any operator ρ, and Π_ρ for the orthogonal projection onto the support of ρ.

Proposition 5.3.10 (Alternative definition of the geometric Rényi trace function, Hiai (2019))**.** *If $\rho, \sigma \in \mathscr{P}(\mathcal{H})$ are such that $\rho \leq c\sigma$ for some $0 < c < \infty$, then there exists a unique bounded operator A with support and image in $s(\sigma)$, such that $\rho^{1/2} = A\sigma^{1/2}$, and this operator also satisfies $A^\dagger A \leq c\Pi_\sigma$ and hence $\|A\|_\infty \leq \sqrt{c}$. If furthermore $b\sigma \leq \rho$ for some $0 < b < \infty$ we additionally have $b\Pi_\sigma \leq A^\dagger A$. We will subsequently also write $A = \rho^{1/2}\sigma^{-1/2}$ and $A^\dagger A = \sigma^{-1/2}\rho\sigma^{-1/2}$. We can then, for any $\alpha \in (0,\infty)$ define the geometric Rényi trace*

function between ρ and σ such that $b\sigma \leq \rho \leq c\sigma$ as

$$\widehat{S}_\alpha(\rho\|\sigma) = \mathrm{Tr}\Big(\sigma(A^\dagger A)^\alpha\Big) = \mathrm{Tr}\Big(\sigma(\sigma^{-1/2}\rho\sigma^{-1/2})^\alpha\Big)\,. \tag{5.3.126}$$

which can be extended to general ρ, σ, via

$$\widehat{S}_\alpha(\rho\|\sigma) = \lim_{\varepsilon\to 0} \widehat{S}_\alpha(\rho + \varepsilon(\rho+\sigma)\|\sigma + \varepsilon(\rho+\sigma))\,. \tag{5.3.127}$$

For $\alpha \in (0,1)\cup(1,2]$ this definition agrees with the one given above in terms of minimal reverse tests.

Lemma 5.3.11 (Martingale convergence, Hiai 2019, Theorem 5.6)**.** *Let $\{M_a\}_a$ be an increasing net of unital von Neumann subalgebras of $\mathcal{B}(\mathcal{H})$ such that $(\bigcup_a M_a)'' = \mathcal{B}(\mathcal{H})$. Then, for every $\rho, \sigma \in \mathscr{P}(\mathcal{H})$*

$$\lim_a \widehat{D}_\alpha(\rho|_{M_a}\|\sigma|_{M_a}) \nearrow \widehat{D}_\alpha(\rho\|\sigma)\,. \tag{5.3.128}$$

In particular, for every increasing sequence of projections $(e_n)_n$ in $\mathcal{B}(\mathcal{H})$ such that $e_n \xrightarrow{n\to\infty} \mathbb{1}$, one can construct such a sequence of subalgebras as $M_n = e_n\mathcal{B}(\mathcal{H})\,e_n \oplus \mathbb{C}(\mathbb{1}-e_n)$.

Lemma 5.3.12. *If $\sigma_1 \leq \sigma_2$, and $\alpha \in (0,1)\cup(1,2]$, then*

$$\widehat{D}_\alpha(\rho\|\sigma_1) \geq \widehat{D}_\alpha(\rho\|\sigma_2). \tag{5.3.129}$$

Proof. We thank Milán Mosonyi for pointing out that this can be seen easily from the above martingale convergence property, and so our original, more involved, proof is not necessary. It is well-known that that the desired statement holds in finite dimensions (see e.g. Matsumoto (2018)). Now, with some sequence of projections onto finite-dimensional subspaces $(e_n)_n$, and M_n defined through the construction in Lemma 5.3.11, we have for $i = 1, 2$: $\sigma_i|_{M_n} = e_n\sigma_i e_n \oplus \mathrm{Tr}(\sigma_i(\mathbb{1}-e_n))$. It is then immediate that $\sigma_1|_{M_n} \leq \sigma_2|_{M_n}$ for all n, and so the statement follows from the finite-dimensional result and Lemma 5.3.11. □

In Hiai (2019), Proposition 2.11, it was shown that if $(s(\rho_1)\cup s(\sigma_1)) \perp (s(\rho_2)\cup s(\sigma_2))$ then $\widehat{S}_\alpha(\rho_1+\rho_2\|\sigma_1+\sigma_2) = \widehat{S}_\alpha(\rho_1\|\sigma_1) + \widehat{S}_\alpha(\rho_2\|\sigma_2)$. We require a similar statement for an infinite orthogonal decomposition, which is not directly implied, so we give a proof here. The argument is still essentially equivalent to the one given for a finite orthogonal decomposition.

Lemma 5.3.13. *For $i \in \mathbb{N}$, let $\rho_i, \sigma_i \in \mathcal{B}_1$ be such that $(s(\rho_i)\cup s(\sigma_i)) \perp (s(\rho_j)\cup s(\sigma_j))$ for all i, j, and such that $\sum_i \rho_i$ and $\sum_i \sigma_i$ converge. Then, we have for $\alpha \in (0,1),(1,2]$:*

$$\widehat{S}_\alpha\Big(\sum_i \rho_i \Big\| \sum_i \sigma_i\Big) = \sum_i \widehat{S}_\alpha(\rho_i\|\sigma_i)\,. \tag{5.3.130}$$

Proof. In light of Proposition 5.3.10 we can show the statement in the case where there exists a $c < \infty$ such that $\rho_i \leq c\sigma_i$ for all i, and the full statement follows by taking limits. For each i pick A_i such that $\rho_i^{1/2} = A_i\sigma_i^{1/2}$. By Proposition 5.3.10 we can pick such an A_i with image and support contained in $s(\sigma_i)$

which are all orthogonal, and so we get the following three statements

$$\left(\sum_i \rho_i\right)^{1/2} = \left(\sum_i \rho_i^{1/2}\right) = \left(\sum_i A_i\right)\left(\sum_i \sigma_i^{1/2}\right) = \left(\sum_i A_i\right)\left(\sum_i \sigma_i\right)^{1/2}, \tag{5.3.131}$$

$$\left(\left(\sum_i A_i^\dagger\right)\left(\sum_j A_j\right)\right)^\alpha = \sum_i (A_i^\dagger A_i)^\alpha, \tag{5.3.132}$$

$$\sigma \sum_i (A_i^\dagger A_i)^\alpha = \sum_i \sigma_i (A_i^\dagger A_i)^\alpha, \tag{5.3.133}$$

which imply the desired equality. □

5.3.5 The (Smoothed) Max-Divergence

The max-divergence is defined analogously as in finite dimensions (Section 2.5)

$$D_{\max}(\rho\|\sigma) \coloneqq \log \inf\{\, \lambda \in \mathbb{R} \mid \rho \leq \lambda\sigma \,\}\,. \tag{5.3.134}$$

It also again satisfies the data-processing inequality (this follows immediately from the positivity of channels).

Similarly also, we again write the ε-ball in purified distance of normalized states around ρ as $B_\varepsilon^{P,\circ}(\rho) = B_\varepsilon^\circ(\rho) = B_\varepsilon(\rho)$, and then define the smoothed max-divergence between a normalized $\rho \in \mathcal{D}(\mathcal{H})$, and $\sigma \in \mathscr{P}(\mathcal{H})$ as:

$$D_{\max}^\varepsilon(\rho\|\sigma) \coloneqq \inf_{\tilde\rho \in B_\varepsilon(\rho)} D_{\max}(\rho\|\sigma)\,. \tag{5.3.135}$$

The data-processing inequality for the smoothed max-divergence also immediately follows from the data-processing inequality for the max-divergence, and for the fidelity.

5.3.6 The Hypothesis Testing Relative Entropy

The definition of a (finite) POVM in infinite dimensions is equivalent to the finite-dimensional definition, and so we can define again

$$D_H^\varepsilon(\rho\|\sigma) \coloneqq -\log\left[\min_{\substack{0\leq M\leq \mathbb{1}_{\mathcal{H}} \\ \mathrm{Tr}(M\rho)\geq 1-\varepsilon}} \mathrm{Tr}(M\sigma)\right] = \max_{\substack{0\leq M\leq \mathbb{1}_{\mathcal{H}} \\ \mathrm{Tr}(M\rho)\geq 1-\varepsilon}} -\log(\mathrm{Tr}(M\sigma))\,. \tag{5.3.136}$$

where the optimizations are still achieved since the functions $M \mapsto \mathrm{Tr}(M\rho)$, $M \mapsto \mathrm{Tr}(M\sigma)$ are weak*-continuous, and the set we optimize over is hence a weak*-closed subset of the unit-ball, which is weak*-compact. The hypothesis testing relative entropy again satisfies the data-processing inequality by the same exact operational argument as in finite dimensions (see e.g. L. Wang and Renner 2012). Also, the relation

$$D_H^\varepsilon(\rho\|\sigma) \leq \frac{1}{1-\varepsilon}(D(\rho\|\sigma) + h(\varepsilon)) \tag{5.3.137}$$

still holds in infinite dimensions, since it is only based on the data-processing inequality of the quantum relative entropy.

One can prove the quantum Stein's Lemma for state discrimination (i.e. the analogue of (2.4.14)) also in infinite dimensions, in fact even in general von Neumann algebras (Jakšić et al. 2012), with the strong-converse property (2.4.15) also holding under some additional finiteness assumptions (Mosonyi 2023). We will in fact not use these statements, and at least the weak-converse version can also be seen as a consequence of our results on quantum channel discrimination (Theorem 5.6.1).

5.4 Lemmas on Smooth and Rényi Entropies in Infinite Dimensions

One fairly recent approach to proving statements using smoothed max- and hypothesis testing relative entropies is to relate them to certain Rényi divergences (see e.g. Khatri and Wilde (2020) for an overview). In this section we prove that some of these relations also hold in separable Hilbert spaces.

This first lemma is a very famous lemma that states that if a measurement has one outcome with very high probability, then the post-measurement state is close to the original state (originally this was introduced as Tender measurement lemma in Winter 1999).

Lemma 5.4.1 (Gentle Measurement Lemma). *Let $\epsilon \in [0,1]$ and $\rho \in \mathcal{D}(\mathcal{H})$ be a quantum state. Let $F \in \mathcal{B}_{\mathrm{sa}}(\mathcal{H})$ be such that $0 \leq F \leq \mathbb{1}$ and $\mathrm{Tr}[F\rho] \geq 1-\epsilon$. Then the post-measurement state*

$$\rho' = \frac{\sqrt{F}\rho\sqrt{F}}{\mathrm{Tr}[F\rho]} \tag{5.4.1}$$

satisfies $P(\rho,\rho') \leq \sqrt{\epsilon}$.

Proof. The finite-dimensional proof of Khatri and Wilde (2020) is only based on the ideas of purification and the data-processing inequality for the fidelity, and hence also works in infinite dimensions: For a pure state $|\psi\rangle$ the squared fidelity between it and its post-measurement state is given by

$$\langle\psi| \frac{\sqrt{F}\,|\psi\rangle\langle\psi|\,\sqrt{F}}{\langle\psi|\,F\,|\psi\rangle}\,|\psi\rangle = \frac{\left|\,\langle\psi|\sqrt{F}|\psi\rangle\,\right|^2}{\langle\psi|F|\psi\rangle} \geq \langle\psi|F|\psi\rangle \geq 1-\varepsilon \tag{5.4.2}$$

where we used that $F \leq \mathbb{1}$ implies $F \leq \sqrt{F}$. Now, if $|\psi\rangle_{RA} \in \mathcal{H}\otimes\mathcal{H}$ is a purification of ρ, then

$$|\psi'\rangle_{RA} = \frac{\mathbb{1}_R \otimes \sqrt{F}\,|\psi\rangle_{RA}}{\sqrt{\langle\psi|\mathbb{1}_R\otimes F|\psi\rangle}} \tag{5.4.3}$$

is a purification of ρ', and hence

$$P(\rho,\rho') \leq P(\psi_{RA},\psi'_{RA}) \leq \sqrt{\varepsilon} \tag{5.4.4}$$

by the data-processing inequality for the purified-distance/fidelity. □

This second lemma forms the basis of many estimates involving the smoothed max divergence. It essentially

originally appeared in Datta (2009), Lemma 15, but with a different smoothing convention (smoothing over sub-normalized states close in trace-distance), hence we obtain a slightly different result.

Lemma 5.4.2. *For two states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ on a separable Hilbert space and for any $\lambda \in (-\infty, D_{\max}(\rho\|\sigma)]$, if $\epsilon = \operatorname{Tr}\left[(\rho - 2^{\lambda}\sigma)_+\right]$, then*

$$D_{\max}^{\sqrt{\epsilon}}(\rho\|\sigma) \leq \lambda - \log(1-\epsilon). \tag{5.4.5}$$

Additionally, for every $\varepsilon \in (1-\operatorname{Tr}(\rho\Pi_\sigma), 1)$ there exists a $\lambda \in (-\infty, D_{\max}(\rho\|\sigma))$ such that $\varepsilon = \operatorname{Tr}\left[(\rho - 2^{\lambda}\sigma)_+\right]$, where Π_σ is the projection onto the support of σ.

Proof. This proof was found together with Jan Kochanowski, following the lines of Tomamichel (2016), Lemma 6.21, and incorporating some ideas from Regula, Lami, and Datta (2025). Let $\Lambda := 2^{\lambda}\sigma$, $\Sigma := (\rho - 2^{\lambda}\sigma)_+$ and $B := \left((\Lambda+\Sigma)^{\frac{1}{2}}\Lambda(\Lambda+\Sigma)^{\frac{1}{2}}\right)^{\frac{1}{2}}$. Since obviously $\Lambda \leq \Lambda + \Sigma$ we also get $B \leq \Lambda + \Sigma$, and hence by the first part of Proposition 5.3.10 (that is essentially just Douglas' range inclusion theorem (Douglas 1966)) there exists an operator $A \in \mathcal{B}(\mathcal{H})$ with support and image contained in $\operatorname{supp}(\Lambda+\Sigma)$ such that $\|A\| \leq 1$ and $A(\Lambda+\Sigma)^{\frac{1}{2}} = B^{\frac{1}{2}}$. We define $G := A^\dagger A$, and find that $0 \leq G \leq \mathbb{1}$. Now, observe that

$$\Lambda + \Sigma - \rho = (2^{\lambda}\sigma - \rho) + (\rho - 2^{\lambda}\sigma)_+ = (2^{\lambda}\sigma - \rho)_+ - (2^{\lambda}\sigma - \rho)_- + (2^{\lambda}\sigma - \rho)_- = (2^{\lambda}\sigma - \rho)_+ \geq 0, \tag{5.4.6}$$

and hence $\rho \leq \Lambda + \Sigma$. Thus, it follows that

$$G\rho G^\dagger \leq G(\Lambda+\Sigma)G^\dagger = A^\dagger B A\,. \tag{5.4.7}$$

Now, let $|\psi_1\rangle, |\psi_2\rangle \in \operatorname{im}\left((\Lambda+\Sigma)^{\frac{1}{2}}\right)$ be two vectors in the image of $(\Lambda+\Sigma)^{\frac{1}{2}}$, so there exist $|\phi_1\rangle, |\phi_2\rangle \in \mathcal{H}$ such that $|\psi_1\rangle = (\Lambda+\Sigma)^{\frac{1}{2}}|\phi_1\rangle$, $|\psi_2\rangle = (\Lambda+\Sigma)^{\frac{1}{2}}|\phi_2\rangle$. We then find

$$\langle\psi_1| A^\dagger B A |\psi_2\rangle = \langle\phi_1|(\Lambda+\Sigma)^{\frac{1}{2}} A^\dagger B A (\Lambda+\Sigma)^{\frac{1}{2}}|\phi_2\rangle = \langle\phi_1| B^2 |\phi_2\rangle \tag{5.4.8}$$

$$= \langle\phi_1|(\Lambda+\Sigma)^{\frac{1}{2}}\Lambda(\Lambda+\Sigma)^{\frac{1}{2}}|\phi_2\rangle \tag{5.4.9}$$

$$= \langle\psi_1|\Lambda|\psi_2\rangle\,. \tag{5.4.10}$$

By continuity this extends to the closure of the image, i.e. to all all $|\psi_1\rangle, |\psi_2\rangle \in \overline{\operatorname{im}(\Lambda+\Sigma)}$. Since

$$\overline{\operatorname{im}(\Lambda+\Sigma)} = [\ker(\Lambda^\dagger + \Sigma^\dagger)]^\perp = [\ker(\Lambda+\Sigma)]^\perp = \operatorname{supp}(\Lambda+\Sigma) \tag{5.4.11}$$

and since both $A^\dagger B A$ and Λ are self-adjoint and supported only on $\operatorname{supp}(\Lambda+\Sigma)$, this implies that $A^\dagger B A = \Lambda$ everywhere, and hence $G(\Lambda+\Sigma)G^\dagger = A^\dagger B A = \Lambda$ as well as

$$G\rho G^\dagger \leq A^\dagger B A = \Lambda = 2^{\lambda}\sigma. \tag{5.4.12}$$

Furthermore,

$$1 - \operatorname{Tr}\left[G^\dagger G\rho\right] = \operatorname{Tr}\left[(\mathbb{1} - G^\dagger G)\rho\right] \overset{\rho\leq\Lambda+\Sigma}{\leq} \operatorname{Tr}\left[(\mathbb{1} - G^\dagger G)(\Lambda + \Sigma)\right] \tag{5.4.13}$$

$$= \operatorname{Tr}[\Lambda + \Sigma] - \operatorname{Tr}\left[G^\dagger G(\Lambda + \Sigma)\right] = \operatorname{Tr}[\Lambda + \Sigma] - \operatorname{Tr}\left[G(\Lambda + \Sigma)G^\dagger\right] \tag{5.4.14}$$

$$= \operatorname{Tr}[\Lambda + \Sigma] - \operatorname{Tr}(\Lambda) = \operatorname{Tr}(\Sigma) \tag{5.4.15}$$

where we used the cyclicity of the trace in the third-last step, which is allowed since $G^\dagger$ is bounded and $G(\Lambda + \Sigma)$ is trace-class. Thus, $\operatorname{Tr}\left[G^\dagger G\rho\right] \geq 1 - \operatorname{Tr}[\Sigma] = 1 - \epsilon$. Defining the state $\tilde{\rho} := \frac{G\rho G^\dagger}{\operatorname{Tr}\left[G^\dagger G\rho\right]} \in \mathcal{D}(\mathcal{H})$, it then holds by the Gentle Measurement Lemma (Lemma 5.4.1) that $P(\rho, \tilde{\rho}) \leq \sqrt{\epsilon}$. We also have that

$$\tilde{\rho} \leq \frac{2^\lambda \sigma}{\operatorname{Tr}[G^\dagger G\rho]} \leq \frac{2^\lambda}{1-\epsilon}\sigma = 2^{\lambda - \log(1-\epsilon)}\sigma \tag{5.4.16}$$

and thus the desired result follows immediately:

$$D_{\max}^{\sqrt{\epsilon}}(\rho\|\sigma) \leq D_{\max}(\tilde{\rho}\|\sigma) \leq \lambda - \log(1-\epsilon)\,. \tag{5.4.17}$$

For the second part of the statement, note that $\lambda \mapsto \operatorname{Tr}\left[(\rho - 2^\lambda\sigma)_+\right]$ is continuous (this follows from the continuity of $x \mapsto (x)_+ = \max\{0, x\}$ on $\mathbb{R}$), and additionally

$$\lim_{\lambda\to-\infty} \operatorname{Tr}\left[(\rho - 2^\lambda\sigma)_+\right] = 1 \tag{5.4.18}$$

$$\lim_{\lambda\to D_{\max}(\rho\|\sigma)} \operatorname{Tr}\left[(\rho - 2^\lambda\sigma)_+\right] = 1 - \operatorname{Tr}(\rho\Pi_\sigma)\,. \tag{5.4.19}$$

□

We can use this result to derive upper bounds on the smoothed max divergence in terms of the Petz-Rényi and the hypothesis testing divergences. The following lemma is the infinite-dimensional analogue of the upper bound in Lemma 2.5.1, see also O. Fawzi, Gao, and Rahaman (2023) for an alternative infinite-dimensional proof of such a statement.

Lemma 5.4.3. *For any $\epsilon \in (0, 1)$ and two states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ on a separable Hilbert space $\mathcal{H}$ the following inequality holds:*

$$D_{\max}^{\epsilon}(\rho\|\sigma) \leq D_H^{1-\epsilon^2}(\rho\|\sigma) - \log\left(1 - \epsilon^2\right) \tag{5.4.20}$$

Proof. This proof was found by Jan Kochanowski. First of all, if $\operatorname{Tr}(\rho\Pi_\sigma) = 1 - \delta < 1$, then for $\varepsilon \leq \delta$ we have

$$D_H^{1-\varepsilon^2}(\rho\|\sigma) \geq D_H^{1-\varepsilon}(\rho\|\sigma) \geq D_H^{1-\delta}(\rho\|\sigma) = \infty \tag{5.4.21}$$

since $\{\Pi_\sigma, \mathbb{1} - \Pi_\sigma\}$ is a suitable measurement that achieves zero type-II error. Hence, in this case the statement holds, and we can in the following restrict to $\varepsilon > 1 - \operatorname{Tr}(\rho\Pi_\sigma)$. Then, by Lemma 5.4.2 there exists some λ such that $\epsilon = \operatorname{Tr}\left[(\rho - 2^\lambda\sigma)_+\right]$. Let $P_+ = \{\rho - 2^\lambda\sigma\}_+$ be the projector onto the support of $(\rho - 2^\lambda\sigma)_+$. Now consider the POVM $\{P_+, P_-\}$ as the decision rule for a hypothesis test between ρ and σ.

The associated type I and II errors are, respectively,

$$\alpha = \operatorname{Tr}[P_-\rho] = 1 - \operatorname{Tr}[P_+\rho] \overset{2^\lambda\sigma\geq 0}{\leq} 1 - \operatorname{Tr}\left[P_+(\rho - 2^\lambda\sigma)\right] = 1 - \operatorname{Tr}[\Sigma] = 1 - \epsilon \tag{5.4.22}$$

$$\beta = \operatorname{Tr}[P_+\sigma] \underset{P_+(\rho-2^\lambda\sigma)\geq 0}{\leq} 2^{-\lambda} \operatorname{Tr}[P_+\rho] \leq 2^{-\lambda}\,. \tag{5.4.23}$$

Therefore, it follows that

$$D_H^{1-\epsilon}(\rho\|\sigma) = -\log \inf_{0\leq F\leq \mathbb{1}} \{\operatorname{Tr}[F\sigma] | \operatorname{Tr}[(\mathbb{1}-F)\rho] \leq 1-\epsilon\} \geq -\log 2^{-\lambda} = \lambda. \tag{5.4.24}$$

Thus the claim follows immediately from Lemma 5.4.2 when substituting ϵ by ϵ^2:

$$D_{\max}^{\sqrt{\epsilon}}(\rho\|\sigma) \leq \lambda - \log(1-\epsilon) \leq D_H^{1-\epsilon}(\rho\|\sigma) + \log\left(\frac{1}{1-\epsilon}\right). \tag{5.4.25}$$

□

Lemma 5.4.4. *For any $\epsilon \in (0,1)$ and $\alpha \in (1,\infty)$ and two states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ on a separable Hilbert space $\mathcal{H}$ the following inequality holds:*

$$D_{\max}^{\epsilon}(\rho\|\sigma) \leq \mathbb{D}_\alpha(\rho\|\sigma) + \frac{2}{\alpha-1}\log\left(\frac{1}{\epsilon}\right) + \log\left(\frac{1}{1-\epsilon^2}\right). \tag{5.4.26}$$

where $\mathbb{D}_\alpha$ can be any quantum α-Rényi divergence, i.e. any function on quantum states that satisfies the data-processing inequality, and reduces to the classical α-Rényi divergence when evaluated on commuting states. Specifically, the result will hold for the Petz-Rényi divergence.

Proof. The proof follows the lines of Tomamichel (2016), Proposition 6.5, although the core of the argument can also already be found much earlier, as e.g. in Ogawa and Nagaoka (2000), Section II. Let $\mathbb{D}_\alpha$ be a quantum Rényi divergence, and let us assume that $\mathbb{D}_\alpha(\rho\|\sigma) < \infty$, as otherwise there is nothing to prove. This then implies that $\operatorname{Tr}(\rho\Pi_\sigma) = 1$, and hence by Lemma 5.4.2 there exists some λ such that $\epsilon = \operatorname{Tr}\left[(\rho - 2^\lambda\sigma)_+\right]$. We will start by constructing classical probability distributions from ρ and σ, by pinching them in the eigenbasis of $(\rho - 2^\lambda\sigma)$. Since $\rho - 2^\lambda\sigma$ is a trace-class (and hence compact) operator, it has a countable eigenbasis, which we shall denote by $\{|\nu_i\rangle\}_{i\in S}$ and write $S_+ := \{i \in S | \langle\nu_i|(\rho - 2^\lambda\sigma)|\nu_i\rangle \geq 0\}$. Set $p_i := \langle\nu_i|\rho|\nu_i\rangle$ and $q_i := \langle\nu_i|\sigma|\nu_i\rangle$, then $P := \{p_i\}_{i\in S}$ and $Q := \{q_i\}_{i\in S}$ are probability distributions on S. Additionally, we have with the channel that implements the measurement in the basis $\{|\nu_i\rangle\}$

$$\mathcal{M}(\omega) := \sum_{i\in S} |\nu_i\rangle\langle\nu_i|\,\omega\,|\nu_i\rangle\langle\nu_i|\,, \tag{5.4.27}$$

that $\mathcal{M}(\rho) = \sum_i p_i\,|\nu_i\rangle\langle\nu_i|$ and $\mathcal{M}(\sigma) = \sum_i q_i\,|\nu_i\rangle\langle\nu_i|$, and hence

$$D_\alpha(P\|Q) = \frac{1}{\alpha-1}\log\left(\sum_{i\in S} p_i^\alpha q_i^{1-\alpha}\right) = \mathbb{D}_\alpha(\mathcal{M}(\rho)\|\mathcal{M}(\sigma)) \leq \mathbb{D}_\alpha(\rho\|\sigma) < \infty\,, \tag{5.4.28}$$

where D_α is the classical α-Rényi divergence. Now, if $i \in S_+$, then $p_i - 2^\lambda q_i \geq 0$ and so $\frac{p_i}{q_i} 2^{-\lambda} \geq 1$. For $\alpha \in (1, \infty)$ we have that

$$\epsilon = \operatorname{Tr}\left[(\rho - 2^\lambda \sigma)_+\right] = \sum_{i \in S_+} p_i - 2^\lambda q_i \leq \sum_{i \in S_+} p_i \leq \sum_{i \in S_+} p_i \left(\frac{p_i}{q_i} 2^{-\lambda}\right)^{\alpha - 1} \tag{5.4.29}$$

$$= 2^{\lambda(1-\alpha)} \sum_{i \in S_+} p_i^\alpha q_i^{1-\alpha} \leq 2^{\lambda(1-\alpha)} \sum_{i \in S} p_i^\alpha q_i^{1-\alpha}\,, \tag{5.4.30}$$

and this is finite since $D_\alpha(P\|Q)$ is finite. Thus,

$$\log \epsilon \leq -\lambda(\alpha - 1) + \log\left(\sum_{i \in S} p_i^\alpha q_i^{1-\alpha}\right) \tag{5.4.31}$$

and hence

$$\lambda \leq \frac{1}{\alpha - 1} \log\left(\sum_{i \in S} p_i^\alpha q_i^{1-\alpha}\right) + \frac{1}{\alpha - 1} \log\left(\frac{1}{\epsilon}\right) \leq D_\alpha(P\|Q) + \frac{1}{\alpha - 1} \log\left(\frac{1}{\epsilon}\right). \tag{5.4.32}$$

With the previous Lemma 5.4.2 we thus have

$$D_{\max}^{\sqrt{\epsilon}}(\rho\|\sigma) \leq \lambda + \log\left(\frac{1}{1-\epsilon}\right) \leq \mathbb{D}_\alpha(\rho\|\sigma) + \log\left(\frac{1}{1-\epsilon}\right) + \frac{1}{\alpha - 1} \log\left(\frac{1}{\epsilon}\right), \tag{5.4.33}$$

and the result then follows after replacing ϵ with ϵ^2. □

We can also construct the following lower bound of the smoothed max divergence by a Rényi divergence, which is an infinite-dimensional variant of *Proposition 4* from X. Wang and Wilde (2019a).

Lemma 5.4.5. *For any two states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ on a separable Hilbert space $\mathcal{H}$ and any $\epsilon \in (0,1), \alpha \in [0,1)$ it holds that*

$$D_{\max}^\epsilon(\rho\|\sigma) \geq D_\alpha(\rho\|\sigma) + \frac{2}{\alpha - 1} \log\left(\frac{1}{1-\epsilon}\right). \tag{5.4.34}$$

Proof. This is a very slight adaptation of the proof in X. Wang and Wilde (2019a). Let us start by showing that if $\rho_0, \rho_1, \sigma \in \mathcal{D}(\mathcal{H})$ are s.t. $\rho \ll \sigma$, $\alpha \in (0,1)$, and $\beta := 2 - \alpha \in (1,2)$, then it holds that

$$D_\beta(\rho_0\|\sigma) - D_\alpha(\rho_1\|\sigma) \geq \frac{2}{1-\alpha} \log\left(1 - \frac{1}{2}\|\rho_0 - \rho_1\|_1\right) \geq \frac{2}{1-\alpha} \log(1 - P(\rho_0, \rho_1)). \tag{5.4.35}$$

For this we can assume that $D_\beta(\rho_0\|\sigma)$ is finite, as otherwise there is nothing to prove. Hence then, by Lemma 5.3.3 there exists an $\eta \in \mathcal{B}_2(\mathcal{H})$ such that $\rho_0^{\frac{\beta}{2}} = \eta \sigma^{\frac{\beta-1}{2}}$. Moreover, since $\frac{\beta}{2} \in [0,1]$, we can use the famous result from Audenaert et al. (2007)

$$\log[1 - \frac{1}{2}\|\rho_0 - \rho_1\|_1] \leq \log \operatorname{Tr}\left[\rho_0^{\frac{\beta}{2}} \rho_1^{1-\frac{\beta}{2}}\right], \tag{5.4.36}$$

which (as stated in Audenaert et al. (2007), see also Audenaert et al. (2008) and Jakšić et al. (2012)) holds also in infinite dimensions. Also, using the Cauchy-Schwarz inequality we get

$$\log \operatorname{Tr}\left[\rho_0^{\frac{\beta}{2}} \rho_1^{1-\frac{\beta}{2}}\right] = \frac{1}{2}\log\left(\operatorname{Tr}\left[\rho_0^{\frac{\beta}{2}} \rho_1^{1-\frac{\beta}{2}}\right]\right)^2 = \frac{1}{2}\log\left(\operatorname{Tr}\left[\eta\sigma^{\frac{\beta-1}{2}} \rho_1^{1-\frac{\beta}{2}}\right]\right)^2 \tag{5.4.37}$$

$$\leq \frac{1}{2}\log\left(\|\eta\|_2^2 \|\sigma^{\frac{\beta-1}{2}} \rho_1^{1-\frac{\beta}{2}}\|_2^2\right) = \frac{1}{2}\log\left(\|\eta\|_2^2 \|\sigma^{\frac{1-\alpha}{2}} \rho_1^{\frac{\alpha}{2}}\|_2^2\right) \tag{5.4.38}$$

Hence, using also the Fuchs-van der Graaf inequality (5.2.13) we get

$$\frac{2}{\beta-1}\log[1-P(\rho_0,\rho_1)] \leq \frac{2}{\beta-1}\log[1-\frac{1}{2}\|\rho_0-\rho_1\|_1] \leq \frac{1}{\beta-1}\log\left(\|\eta\|_2^2 \|\sigma^{\frac{1-\alpha}{2}} \rho_1^{\frac{\alpha}{2}}\|_2^2\right) \tag{5.4.39}$$

$$= \frac{1}{\beta-1}\log\|\eta\|_2^2 - \frac{1}{\alpha-1}\log\operatorname{Tr}\left[\rho_1^\alpha \sigma^{1-\alpha}\right] = D_\beta(\rho_0\|\sigma) - D_\alpha(\rho_1\|\sigma)\,. \tag{5.4.40}$$

Secondly, let us show that for any $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, $D_{\max}(\rho\|\sigma) \geq D_2(\rho\|\sigma)$. For this assume that $D_{\max}(\rho\|\sigma) = \log(c) < \infty$, i.e. there exists a c such that $\rho \leq c\sigma$. Then, by the first part of Proposition 5.3.10, there exists an $A \in \mathcal{B}(\mathcal{H})$ such that $\sqrt{\rho} = A\sqrt{\sigma}$ and $\|A\| \leq c$. Now, $\eta = \sqrt{\rho}A$ satisfies $\eta\sqrt{\sigma} = \rho$, and also $\operatorname{Tr}(\eta^\dagger\eta) = \operatorname{Tr}(\sqrt{\rho}A\sqrt{\rho}) \leq c$, so $\eta \in \mathcal{B}_2(\mathcal{H})$. Hence, η is the unique η of Lemma 5.3.3 for $\alpha = 2$, and thus we have

$$D_2(\rho\|\sigma) \leq \log(c) = D_{\max}(\rho\|\sigma)\,. \tag{5.4.41}$$

Finally, for the actual statement of the Lemma, fix $\alpha \in (0,1)$ and $\tilde{\rho} \in B^\epsilon(\rho)$, then for $\beta = 2-\alpha \in (1,2)$ we have, by the monotonicity of $\alpha \mapsto D_\alpha(\rho\|\sigma)$ on $(1,2]$, that

$$D_{\max}(\tilde{\rho}\|\sigma) \geq D_2(\tilde{\rho}\|\sigma) \geq D_\beta(\tilde{\rho}\|\sigma) \geq D_\alpha(\rho\|\sigma) + \frac{2}{1-\alpha}\log(1-P(\rho,\tilde{\rho})) \tag{5.4.42}$$

$$\geq D_\alpha(\rho\|\sigma) + \frac{2}{\alpha-1}\log\left(\frac{1}{1-\epsilon}\right). \tag{5.4.43}$$

Optimizing over all $\tilde{\rho} \in B^\epsilon(\rho)$ yields the desired statement for $\alpha \in (0,1)$. Taking the limit $\alpha \to 0$ gives the statement for $\alpha = 0$. □

5.4.1 Asymptotic Equipartition Property

Using all of the above, we can conclude that our finite-n AEP error bound (Lemma 2.5.2) also holds in infinite dimensions:

Lemma 5.4.6 (Finite n AEP). *Lemma 2.5.2 also holds in infinite dimensions. In particular, if $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ are states on a separable Hilbert space $\mathcal{H}$, and for $\gamma \in (0,1]$ let $c_\gamma(\rho\|\sigma)$ be defined as in (5.3.29), then, if $c_\gamma(\rho\|\sigma) < \infty$, it holds that for all $\epsilon \in (0,1]$ and $n \in \mathbb{N}$*

$$\frac{1}{n}D^\epsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D(\rho\|\sigma) - \frac{4c_\gamma(\rho\|\sigma)}{\sqrt{n}}\log\left(\frac{2}{1-\epsilon}\right), \tag{5.4.44}$$

$$\frac{1}{n}D^\epsilon_{\max}(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq D(\rho\|\sigma) + \frac{4c_\gamma(\rho\|\sigma)}{\sqrt{n}}\log\left(\frac{2}{\epsilon}\right) + \frac{1}{n}\log\left(\frac{1}{1-\epsilon^2}\right). \tag{5.4.45}$$

Proof. The proof is identical to the proof of Lemma 2.5.2, since, as established in this section, all the statements used in that proof also hold in infinite dimensions. □

Remark 5.4.7. The authors of O. Fawzi, Gao, and Rahaman (2023) also recently proved a different version of an AEP for the smoothed max divergence in infinite dimensions. However, their version is not suitable for our applications. Specifically, they show what is known as second-order asymptotics, where the second order term is controlled by the relative entropy variance $V(\rho\|\sigma)$. For our result we subsequently need to upper bound the error terms in the AEP by something that satisfies a chain rule (we will use the geometric Rényi divergence for this purpose), and we are not aware of any way to do this for the relative entropy variance, which is why we cannot make use of the results from O. Fawzi, Gao, and Rahaman (2023).

5.5 One-Shot Error Exponents of Adaptive and Parallel Strategies

Throughout this section, we will assume that $\mathcal{H}_A, \mathcal{H}_B$ are two separable Hilbert spaces and $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}_A) \to \mathscr{P}(\mathcal{H}_B)$ are two CPTP maps such that $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$, where we will discuss this last condition in detail towards the end of this section. We then get the following result regarding the discrimination of these two quantum channels in infinite dimensions, which is the infinite-dimensional analogue of Corollary 3.3.1:

Corollary 5.5.1 (Infinite-dimensional version of Corollary 3.3.1)**.** *Let $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}_A) \to \mathscr{P}(\mathcal{H}_B)$ be two quantum channels between separable Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, s.t. $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$. Let there be an adaptive discrimination protocol with n channel uses, that – for an arbitrary type I error $\alpha_a \in [0,1]$ – achieves type II error $\beta_a(\alpha_a)$, and thus type-II error rate $\frac{-1}{n}\log(\beta_a(\alpha_a))$. Then, for all $\alpha_p \in (0,1]$ there exists a parallel protocol with m channel uses and type II error $\beta_p(\alpha_p)$ such that for all $\alpha_a \in [0,1]$ the type II error rates per channel use obey the following relation:*

$$-\frac{1}{m}\log(\beta_p(\alpha_p)) \geq -\frac{1-\alpha_a}{n}\log(\beta_a(\alpha_a)) - \frac{Cn}{\sqrt{m}}\log\left(\frac{8}{\alpha_p}\right) - \frac{1}{n}. \tag{5.5.1}$$

That is, the type II error rate of the parallel protocol is essentially at least as good as the adaptive one modulo an additional error term, which decays as $m \to \infty$. The constant C is given by

$$C := 7 \inf_{\gamma\in(0,1]} \frac{1}{\gamma}\log\left(2^{\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\right) \tag{5.5.2}$$

Similarly to our finite-dimensional result, this corollary comes with an associated theorem that states the tight bound we are able to obtain. This theorem is the infinite-dimensional version of Theorem 3.3.4, and essentially identical, with the exception of one additional term ($\frac{1}{m}\mu$) coming from the fact that optimizations over input states are in general not achieved in infinite dimensions (see also Lemma 5.5.3 below).

Theorem 5.5.2 (Infinite-dimensional version of Theorem 3.3.4)**.** *Let $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}_A) \to \mathscr{P}(\mathcal{H}_B)$ be two quantum channels between separable Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, s.t. $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$. Given an arbitrary adaptive protocol with n channel uses, we write $\rho_i, \sigma_i \in \mathcal{D}(R_a \otimes A)$, $i \in \{1, \ldots, n\}$ for the states*

that are input into the channel during the adaptive protocol (ρ_i if the channel is $\mathcal{E}$, and σ_i if the channel is $\mathcal{F}$; see Section 3.1 above for a more detailed explanation of this notation). We define $\ell \in \{1, \dots, n\}$ as the step in the protocol where the distinguishability increases the most, i.e.,

$$\ell \coloneqq \underset{k\in\{1,\dots,n\}}{\arg\max} \Big[D(\mathcal{E}(\rho_k)\|\mathcal{F}(\sigma_k)) - D(\rho_k\|\sigma_k) \Big] . \tag{5.5.3}$$

Then, for all $\alpha_p \in (0,1]$, $m \in \mathbb{N}$ and $\mu > 0$, there exists a separable Hilbert space $\mathcal{H}_R$ and a state $\nu \in \mathcal{D}\left(R \otimes A^{\otimes m}\right)$ such that for all $\alpha_a \in [0,1]$:

$$\begin{aligned} \frac{1}{m} D_H^{\alpha_p}\left(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\right) \geq \frac{1-\alpha_a}{n} D_H^{\alpha_a}\left(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)\right) - \frac{c'_\ell}{\sqrt{m}}\left[\log\left(\frac{4}{\alpha_p}\right) + K\right] \\ - \frac{1}{m}\left[\log\left(\frac{1}{\alpha_p}\right) - \log\left(1-\frac{\alpha_p}{4}\right) + \mu\right] - \frac{h(\alpha_a)}{n} , \end{aligned} \tag{5.5.4}$$

where

$$K \coloneqq \frac{\ln(2)\log^2(3)}{8} \cosh\left(\frac{\log(3)}{2}\right) \leq 0.29 \tag{5.5.5}$$

and c'_ℓ depends on the pair of channels $\mathcal{E}$, $\mathcal{F}$ and can be bounded as follows:

$$\begin{aligned} c'_\ell &\coloneqq \frac{4}{\log(3)} \inf_{\gamma_1,\gamma_2\in(0,1]} \left[c_{\gamma_1}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) + c_{\gamma_2}(\rho_\ell\|\sigma_\ell)\right] && (5.5.6)\\ &\leq \frac{8\ell}{\log(3)} \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) && (5.5.7)\\ &\leq \frac{8n}{\log(3)} \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) , && (5.5.8) \end{aligned}$$

where

$$c_\gamma(\rho\|\sigma) \coloneqq \frac{1}{\gamma}\log\left(2^{\gamma D_{1+\gamma}(\rho\|\sigma)} + 2^{-\gamma D_{1-\gamma}(\rho\|\sigma)} + 1\right), \tag{5.5.9}$$

$$\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \coloneqq \frac{1}{\gamma}\log\left(2^{\gamma \widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\right) . \tag{5.5.10}$$

Moreover, if

$$m \geq \log\left(\frac{4}{\alpha_p}\right)\left(\frac{4}{\log(3)\sqrt{2\ln(2)}}\right)^2 , \tag{5.5.11}$$

then we have the following tighter bound

$$\begin{aligned} \frac{1}{m} D_H^{\alpha_p}\left(\mathcal{E}^{\otimes m}(\nu)\|\mathcal{F}^{\otimes m}(\nu)\right) \geq \frac{1-\alpha_a}{n} D_H^{\alpha_a}\left(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)\right) \\ - \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{4}{\alpha_p}\right)} - \frac{1}{m}\left[\log\left(\frac{1}{\alpha_p}\right) - \log\left(1-\frac{\alpha_p}{4}\right) + \mu\right] - \frac{h(\alpha_a)}{n} , \end{aligned} \tag{5.5.12}$$

where c_ℓ is defined in a similar way as c'_ℓ but with different numerical constants K_1 and K_2:

$$K_1 \coloneqq 2\sqrt{2\ln(2)\cosh\left(\frac{\log(3)}{2}\right)} \leq 2.72, \tag{5.5.13}$$

$$K_2 \coloneqq 2\sqrt{2\ln(2)} \leq 2.36, \tag{5.5.14}$$

$$c_\ell \coloneqq \inf_{\gamma_1,\gamma_2\in(0,1]} [K_1 c_{\gamma_1}(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) + K_2 c_{\gamma_2}(\rho_\ell\|\sigma_\ell)] \tag{5.5.15}$$

$$\leq \ell(K_1+K_2) \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \tag{5.5.16}$$

$$\leq n(K_1+K_2) \inf_{\gamma\in(0,1]} \widehat{c}_\gamma(\mathcal{E}\|\mathcal{F})\,. \tag{5.5.17}$$

The parallel input state ν can be chosen as either:

- *A μ-close optimizer in the smoothing of the max-divergence $D^\varepsilon_{\max}(\rho_\ell^{\otimes m}\|\sigma_\ell^{\otimes m})$, where $\varepsilon = \frac{1}{2}\left(1-\sqrt{1-\alpha_p}\right)$, i.e.,*

$$D_{\max}(\nu\|\sigma_\ell^{\otimes m}) = D^\varepsilon_{\max}(\rho_\ell^{\otimes m}\|\sigma_\ell^{\otimes m}) + \mu. \tag{5.5.18}$$

  In this case the reference system is equal to m tensor copies of the reference system of the original adaptive strategy. We refer to this input state as $\tilde{\nu}$.

- *Any purification of the A^m-marginal $\tilde{\nu}_{A^m} = \mathrm{Tr}_{R_a^m}(\tilde{\nu}_{R_a^m A^m})$, in which case the reference system is the purifying system.*

We will use Theorem 5.5.2 to prove the asymptotic equivalence of adaptive and parallel strategies in Theorem 5.6.1, which means that unlike the original argument in finite dimensions (Fang et al. 2020, X. Wang and Wilde 2019b), we do not go the way of establishing an amortized expression for the adaptive exponent and then applying a chain rule. However, we still consider the chain rule for the quantum relative entropy to be of independent interest, and will establish it separately in Theorem 5.6.5.

Proof of Theorem 5.5.2

Our way of proving our one-shot result proceeds by a similar strategy as the proof of the finite-dimensional result Theorem 3.3.4. In the previous sections we have established most of the necessary lemmas in the infinite-dimensional setting, the last one that is remaining is an infinite-dimensional version of the one-shot chain rule:

Lemma 5.5.3 (An extended infinite-dimensional version of Lemma 3.3.7, see also Fang et al. (2020), Prop. 3.2)**.** *Let $\mathcal{E},\mathcal{F}:\mathscr{P}(\mathcal{H})\to\mathscr{P}(\mathcal{K})$ be two quantum channels, where $\mathcal{H},\mathcal{K}$ are separable Hilbert spaces and $\rho,\sigma\in\mathcal{D}(\mathcal{H})$ be some states. Then for any $\epsilon,\epsilon',\mu>0$ and any $m\in\mathbb{N}$, there exists a state $\nu\equiv\nu(m,\mu,\rho,\sigma)\in B_\epsilon(\rho)$ s.t.*

$$D^{\epsilon+\epsilon'}_{\max}\left((\mathcal{E}(\rho))^{\otimes m}\|(\mathcal{F}(\sigma))^{\otimes m}\right) \leq D^{\epsilon'}_{\max}\left((\mathcal{E}(\nu))^{\otimes m}\|(\mathcal{F}(\nu))^{\otimes m}\right) + mD^{\epsilon/m}_{\max}(\rho\|\sigma) + \mu\,. \tag{5.5.19}$$

Additionally it holds that:

$$D_{\max}^{\epsilon+\epsilon'}\big((\mathcal{E}(\rho))^{\otimes m}\|\mathcal{F}(\sigma))^{\otimes m}\big) \leq \sup_{\nu\in\mathcal{D}(\mathcal{H})} D_{\max}^{\epsilon'}\big((\mathcal{E}(\nu))^{\otimes m}\|(\mathcal{F}(\nu))^{\otimes m}\big) + mD_{\max}^{\epsilon/m}(\rho\|\sigma)\,. \tag{5.5.20}$$

Proof. The proof is essentially the same as in Lemma 3.3.7, just that we cannot assume that the infimum in the smoothed max-relative entropy is achieved, as the underlying Hilbert space being infinite-dimensional means that $B_\epsilon(\rho)$ is not a (norm-)compact set.

Assume the requirements on $\mathcal{E}, \mathcal{F}, \rho, \sigma, \epsilon, \epsilon', \mu$ as stated in the Lemma. Pick $\nu \in B^{\epsilon/m}(\rho)$ s.t.

$$D_{\max}^{\epsilon/m}(\rho\|\sigma) \geq D_{\max}(\nu\|\sigma) - \frac{\mu}{2m}\,, \tag{5.5.21}$$

i.e. close to 'the optimal smoothing state'. Note that this implies

$$\nu^{\otimes m} \leq 2^{mD_{\max}^{\epsilon/m}(\rho\|\sigma)+\frac{\mu}{2}}\sigma^{\otimes m}\,. \tag{5.5.22}$$

Similarly pick $\tau \in B_{\epsilon'}(\mathcal{E}(\nu)^{\otimes m})$ s.t.

$$D_{\max}^{\epsilon'}(\mathcal{E}(\nu)^{\otimes m}\|\mathcal{F}(\nu)^{\otimes m}) \geq D_{\max}(\tau\|\mathcal{F}(\nu)^{\otimes m}) - \frac{\mu}{2}\,. \tag{5.5.23}$$

Now,

$$\tau \leq 2^{D_{\max}(\tau\|\mathcal{F}(\nu))^{\otimes m}}\mathcal{F}(\nu)^{\otimes m} \leq 2^{D_{\max}^{\epsilon'}(\mathcal{E}(\nu)^{\otimes m}\|\mathcal{F}(\nu)^{\otimes m})+\frac{\mu}{2}}2^{mD_{\max}^{\epsilon/m}(\rho\|\sigma)+\frac{\mu}{2}}\mathcal{F}(\sigma)^{\otimes m}\,, \tag{5.5.24}$$

where we used that $\omega \geq \tilde{\omega} \implies \mathcal{E}(\omega) \geq \mathcal{E}(\tilde{\omega})$ for any two states $\omega, \tilde{\omega}$. Now, by the triangle inequality and DPI for the purified distance we have

$$P(\tau, \mathcal{E}(\rho)^{\otimes m}) \leq P(\tau, \mathcal{E}(\nu)^{\otimes m}) + P(\mathcal{E}(\nu)^{\otimes m}, \mathcal{E}(\rho)^{\otimes m}) \leq \epsilon' + P(\nu^{\otimes m}, \rho^{\otimes m})\,. \tag{5.5.25}$$

Now, we have (this in fact works for any distance that satisfies the triangle inequality and monotonicity under CPTP maps)

$$P(\nu^{\otimes m}, \rho^{\otimes m}) \leq P(\nu^{\otimes m}, \nu\otimes\rho^{\otimes m-1}) + P(\nu\otimes\rho^{\otimes m-1}, \rho^{\otimes m}) \leq P(\nu^{\otimes m-1}, \rho^{\otimes m-1}) + P(\nu,\rho) \tag{5.5.26}$$

$$\leq \ldots \leq mP(\nu,\rho) = \varepsilon \tag{5.5.27}$$

where we used that $P(\rho\otimes\nu, \nu\otimes\nu) \leq P(\rho,\nu)$, since the states on the left-hand side can be generated from the states on the right-hand side by the application of a channel that adds a system in the state ν. Note that by using a different argument involving directly the definition of the purified distance, this can actually be tightened to $\sqrt{m}P(\nu,\rho)$, which would then also improve our result. However, since the only application of Lemma 5.5.3 where we don't set $m = 1$ is Theorem 5.6.5, where this tightening is of no consequence (and only leads to slightly more writing effort), we keep the statement as it is.

The second inequality of the statement then also immediately follows by taking the supremum over ν on

the right-hand side and then the limit $\mu \to 0$. □

With this we can then complete the proof of Theorem 5.5.2:

Proof of Theorem 5.5.2. As the proof is almost identical to the proof of Theorem 3.3.4 we will not repeat everything but just highlight the one occasion where things slightly differ. First note that the assumed finiteness condition implies that $c_\ell < \infty$. We can follow the proof of Theorem 3.3.4 completely analogously (using the corresponding infinite-dimensional technical lemmas established previously in this chapter) until (3.3.42), after which we use Lemma 5.5.3 instead of Lemma 3.3.7 to obtain a slight variation of (3.3.43), namely that for every $\mu > 0$ we obtain a $\tilde{\nu} \in B_\varepsilon(\rho_\ell^{\otimes m})$ such that:

$$D(\mathcal{E}(\rho_\ell)\|\mathcal{F}(\sigma_\ell)) - D(\rho_\ell\|\sigma_\ell) \leq \frac{1}{m} D_{\max}^{1-2\varepsilon}\left(\mathcal{E}^{\otimes m}(\tilde{\nu})\|\mathcal{F}^{\otimes m}(\tilde{\nu})\right) + \frac{1}{m}\mu + \frac{c_\ell}{\sqrt{m}}\sqrt{\log\left(\frac{1}{\varepsilon}\right)} + \frac{1}{m}\log\left(\frac{1}{1-\varepsilon^2}\right). \tag{5.5.28}$$

The subsequent discussion regarding purifications of the A^m-marginal of $\tilde{\nu}$ still applies also to any purification of this infinite-dimensional $\tilde{\nu}$. The rest of the proof the proceeds identically, and the additional term $\frac{1}{m}\mu$ stays around until the end. □

Proof of Corollary 5.5.1. Note that in the statement of Theorem 5.5.2, $\mu > 0$ can be chosen at will and can depend on any of the other parameters. Hence, for Corollary 5.5.1 it can be absorbed into any estimate that is a strict inequality. Looking at the proof of Corollary 3.3.1 (the corresponding finite-dimensional statement), this is the case for example for $K \leq 0.29 < 1$. Corollary 5.5.1 then follows from Theorem 5.5.2 by the same estimates as Corollary 3.3.1. □

5.6 Asymptotic Error Exponents

In this section we use the finite-n results and general techniques established in the previous sections to prove some asymptotic results for quantum channel discrimination, as these have also not previously been established in the infinite-dimensional setting.

5.6.1 Quantum Stein's Lemma for Channels in Infinite Dimensions

Our main asymptotic result is the following characterization of the Stein exponent for asymmetric channel discrimination, which establishes the role of the regularized channel divergence as the asymptotic asymmetric error exponent of asymmetric channel discrimination also in infinite dimensions.

Theorem 5.6.1 (Stein's lemma for infinite-dimensional channels)**.** *Let $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}_A) \to \mathscr{P}(\mathcal{H}_B)$ be two quantum channels between separable Hilbert spaces $\mathcal{H}_A, \mathcal{H}_B$. Then, the asymmetric asymptotic error exponent of discriminating these two channels with a parallel strategy is given by*

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty}\frac{1}{n}D_H^\varepsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) = D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.1}$$

Moreover, if $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$, then this is also equal to the asymptotic error exponent using adaptive strategies and hence adaptive strategies offer no asymptotic advantage.

To prove this, we will first establish that this rate is achievable and optimal for parallel strategies, and then use the one-shot conversion theorem Theorem 5.5.2 to establish the same also for adaptive strategies (assuming the finiteness condition).

The achievability part for parallel strategies goes via the path of constructing a blocked strategy, and then exploiting the emerging i.i.d. structure using the additivity of Rényi divergences:

Proposition 5.6.2. *Let $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{K})$ be two quantum channels, where $\mathcal{H}, \mathcal{K}$ are arbitrary separable Hilbert spaces. Then the asymptotic error exponent with parallel strategies satisfies:*

$$\lim_{\epsilon\to 0} \liminf_{n\to\infty} \frac{1}{n} D_H^\epsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) \geq D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.2}$$

Proof. Combining Lemma 5.4.3 and Lemma 5.4.5 we get for any two states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, and any $\alpha \in [0, 1)$,

$$D_H^\epsilon(\rho\|\sigma) \geq D_\alpha(\rho\|\sigma) - \log\left(\frac{1}{\epsilon}\right) + \frac{2}{\alpha - 1}\log\left(\frac{1}{1-\sqrt{1-\epsilon}}\right) \tag{5.6.3}$$

for any $\alpha \in (0, 1)$ and for any states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$. Since the α-Petz-Rényi divergence is additive it follows that

$$\frac{1}{n} D_H^\epsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D_\alpha(\rho\|\sigma) + \frac{1}{n}\left(\frac{2}{\alpha - 1}\log\left(\frac{1}{1-\sqrt{1-\epsilon}}\right) - \log\left(\frac{1}{\epsilon}\right)\right). \tag{5.6.4}$$

The trick is now to apply this to a blocked input state, i.e. an entangled state on m systems that is repeated n times. Since all separable Hilbert spaces are isomorphic, we can restrict ourselves to a single reference system $\mathcal{H}_R$ in this derivation. We then get:

$$\begin{aligned}
\frac{1}{nm} D_H^\epsilon(\mathcal{E}^{\otimes nm}\|\mathcal{F}^{\otimes nm}) &= \frac{1}{nm} \sup_{\nu_{RA^{nm}}} D_H^\epsilon(\mathcal{E}^{\otimes nm}(\nu_{RA^{nm}})\|\mathcal{F}^{\otimes nm}(\nu_{RA^{nm}})) && (5.6.5)\\
&\geq \frac{1}{nm} \sup_{\nu_{RA^m}} D_H^\epsilon\Big(\big(\mathcal{E}^{\otimes m}(\nu_{RA^m})\big)^{\otimes n} \,\Big\|\, \big(\mathcal{F}^{\otimes m}(\nu_{RA^m})\big)^{\otimes n}\Big) && (5.6.6)\\
&\geq \frac{1}{m} D_\alpha(\mathcal{E}^{\otimes m}\|\mathcal{F}^{\otimes m}) + \frac{1}{nm}\left(\frac{2}{\alpha - 1}\log\left(\frac{1}{1-\sqrt{1-\epsilon}}\right) - \log\left(\frac{1}{\epsilon}\right)\right). && (5.6.7)
\end{aligned}$$

Taking the $\liminf_{nm\to\infty}$ on both sides yields

$$\liminf_{nm\to\infty} \frac{1}{nm} D_H^\epsilon(\mathcal{E}^{\otimes nm}\|\mathcal{F}^{\otimes nm}) \geq D_\alpha^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) \tag{5.6.8}$$

for any $\alpha \in [0, 1)$. Now, as already mentioned, also in infinite dimensions, the Petz–Rényi divergence is increasing in α, and converges to the quantum relative entropy as $\alpha \uparrow 1$, i.e.

$$\lim_{\alpha\uparrow 1} D_\alpha(\rho\|\sigma) = \sup_{\alpha\in(0,1)} D_\alpha(\rho\|\sigma) = D(\rho\|\sigma) \tag{5.6.9}$$

and hence it follows that

$$\lim_{\alpha\uparrow 1} D_\alpha^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) = \sup_{\alpha\in(0,1)} D_\alpha^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) = \sup_{\alpha\in(0,1)} \sup_{m\in\mathbb{N}} \sup_{\nu_{RA^m}} \frac{1}{m} D_\alpha(\mathcal{E}^{\otimes m}(\nu_{RA^m})\|\mathcal{F}^{\otimes m}(\nu_{RA^m})) \tag{5.6.10}$$

$$= \sup_{m\in\mathbb{N}} \sup_{\nu_{RA^m}} \sup_{\alpha\in(0,1)} \frac{1}{m} D_\alpha(\mathcal{E}^{\otimes m}(\nu_{RA^m})\|\mathcal{F}^{\otimes m}(\nu_{RA^m})) = \sup_{m\in\mathbb{N}} \frac{1}{m} D(\mathcal{E}^{\otimes m}\|\mathcal{F}^{\otimes m}) \tag{5.6.11}$$

$$= D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.12}$$

Thus, we have the desired result

$$\liminf_{nm\to\infty} \frac{1}{nm} D_H^\epsilon(\mathcal{E}^{\otimes nm}\|\mathcal{F}^{\otimes nm}) \geq \sup_{\alpha\in(0,1)} D_\alpha^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) = D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.13}$$

□

The converse bound is an immediate application of (5.3.137):

Lemma 5.6.3. *Let $\mathcal{E},\mathcal{F} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{K})$ be two quantum channels, where $\mathcal{H},\mathcal{K}$ are arbitrary separable Hilbert spaces. Then*

$$\lim_{\epsilon\to 0} \limsup_{n\to\infty} \frac{1}{n} D_H^\epsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) \leq D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.14}$$

Proof. For any input state ν_{RA^n}, (5.3.137) gives

$$\frac{1}{n} D_H^\varepsilon(\mathcal{E}^{\otimes n}(\nu_{RA^n})\|\mathcal{F}^{\otimes n}(\nu_{RA^n})) \leq \frac{1}{n}\frac{1}{1-\varepsilon}\left[D(\mathcal{E}^{\otimes n}(\nu_{RA^n})\|\mathcal{F}^{\otimes n}(\nu_{RA^n})) + h(\varepsilon)\right] \tag{5.6.15}$$

and hence

$$\frac{1}{n} D_H^\varepsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) \leq \frac{1}{n}\frac{1}{1-\varepsilon}\left[D(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) + h(\varepsilon)\right] \tag{5.6.16}$$

from which the statement follows by taking $\lim_{\varepsilon\to 0} \limsup_{n\to\infty}$ in that order. □

These two statements together prove the parallel quantum channel Stein's lemma. The statement of the adaptive quantum Stein's lemma for channels is that the same error exponent is also optimal for adaptive discrimination strategies. In the finite-dimensional setting it was proved by combining a statement that the asymptotic optimal adaptive scaling is given by the amortized relative channel divergence (X. Wang and Wilde 2019b), instead of the regularized one, and a chain rule for the quantum relative entropy (Fang et al. 2020) which implies that the amortized and regularized quantum channel divergences are indeed equivalent. Here we use our one-shot result Theorem 5.5.2 to show this equivalence directly and without going through the amortized channel divergence. Note that as mentioned and discussed in Section 5.6.3, in infinite dimensions, we are not able to show this equivalence for all channels, but only under the condition that the geometric Rényi divergence between the channels is finite for some $\alpha > 1$.

Proposition 5.6.4 (Quantum Stein's Lemma for adaptive strategies in infinite dimensions)**.** *Let $\mathcal{E},\mathcal{F} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{K})$ be two quantum channels (CPTP maps) where $\mathcal{H},\mathcal{K}$ are arbitrary separable Hilbert spaces. If*

there exists $\alpha > 1$, *s.t.* $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$, *then*

$$\lim_{\epsilon\to 0}\lim_{n\to\infty}\sup_{\rho_1=\sigma_1;\{\Lambda_i\}_i}\frac{1}{n}D_H^\epsilon(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) = D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.17}$$

Proof. Since adaptive strategies are more general than parallel ones, it follows that

$$\sup_{\rho_1=\sigma_1;\{\Lambda_i\}_i}\frac{1}{n}D_H^\epsilon(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) \geq \sup_{\rho}\frac{1}{n}D_H^\epsilon(\mathcal{E}^{\otimes n}(\rho)\|\mathcal{F}^{\otimes n}(\rho)) = \frac{1}{n}D_H^\epsilon(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}). \tag{5.6.18}$$

which implies

$$D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) \leq \lim_{\epsilon\to 0}\liminf_{n\to\infty}\sup_{\rho_1=\sigma_1;\{\Lambda_i\}_i}\frac{1}{n}D_H^\epsilon(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n))\,. \tag{5.6.19}$$

If $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$, then Theorem 5.5.2 holds. In Theorem 5.5.2, taking the supremum over parallel input states ν_m, and then the limits $m \to \infty$, $n \to \infty$, $\alpha_a \to 0$, $\alpha_p \to 0$ in this order, we find that:

$$\lim_{\alpha_a\to 0}\limsup_{n\to\infty}\frac{1}{n}D_H^{\alpha_a}(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) \leq \lim_{\alpha_p\to 0}\limsup_{m\to\infty}\sup_{\nu_m}D_H^{\alpha_p}(\mathcal{E}^{\otimes m}(\nu_m)\|\mathcal{F}^{\otimes m}(\nu_m)) = D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) \tag{5.6.20}$$

which allows us to conclude that the limit exists without requiring $\liminf$ or $\limsup$, and:

$$D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}) \leq \lim_{\epsilon\to 0}\lim_{n\to\infty}\sup_{\rho_1=\sigma_1;\{\Lambda_i\}_i}\frac{1}{n}D_H^\epsilon(\mathcal{E}(\rho_n)\|\mathcal{F}(\sigma_n)) \leq D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.21}$$

□

These two propositions together prove Theorem 5.6.1.

5.6.2 Chain Rule for the Quantum Relative Entropy

While we did not require it to prove the asymptotic equivalence of adaptive and parallel strategies, we still believe the chain rule for the quantum relative entropy to be of independent interest. Also in infinite dimensions, it can be be proved quite straightforwardly combining Lemma 5.4.6 and Lemma 5.5.3:

Theorem 5.6.5. *Let* $\mathcal{H}, \mathcal{K}$ *be separable Hilbert spaces, and* $\mathcal{E}, \mathcal{F} : \mathscr{P}(\mathcal{H}) \to \mathscr{P}(\mathcal{K})$ *be two quantum channels, such that* $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ *for some* $\alpha > 1$. *Then, for any two states* $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ *it holds that:*

$$D(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) \leq D(\rho\|\sigma) + D^{\mathrm{reg}}(\mathcal{E}\|\mathcal{F}). \tag{5.6.22}$$

Proof. Applying the second statement in Lemma 5.5.3 with $\rho \leftarrow \rho^{\otimes n}$, $\sigma \leftarrow \sigma^{\otimes n}$, $\mathcal{E} \leftarrow \mathcal{E}^{\otimes n}$ and $\mathcal{F} \leftarrow \mathcal{F}^{\otimes n}$ we get

$$D_{\max}^{\epsilon+\epsilon'}\big((\mathcal{E}(\rho))^{\otimes nm}\|(\mathcal{F}(\sigma))^{\otimes nm}\big) \leq \sup_{\nu\in\mathcal{D}(\mathcal{H}^{\otimes n})}D_{\max}^{\epsilon'}\big((\mathcal{E}^{\otimes n}(\nu))^{\otimes m}\|(\mathcal{F}^{\otimes n}(\nu))^{\otimes m}\big) + mD_{\max}^{\epsilon/m}(\rho^{\otimes n}\|\sigma^{\otimes n})\,. \tag{5.6.23}$$

Now, dividing my nm and using Lemma 5.4.6 we get

$$\begin{aligned} &D(\mathcal{E}(\rho)\|\mathcal{F}(\sigma)) - \mathcal{O}\left(\frac{1}{\sqrt{nm}}\right) \\ &\quad\leq \frac{1}{n}\sup_{\nu\in\mathcal{D}(\mathcal{H}^{\otimes n})}\left(D(\mathcal{E}^{\otimes n}(\nu)\|\mathcal{F}^{\otimes n}(\nu)) + \frac{4c_\gamma(\mathcal{E}^{\otimes n}(\nu)\|\mathcal{F}^{\otimes n}(\nu))}{\sqrt{m}}\log\left(\frac{2}{\epsilon}\right) + \mathcal{O}\left(\frac{1}{nm}\right)\right) \\ &\qquad\qquad + D(\rho\|\sigma) + \mathcal{O}\left(\log\left(\frac{2m}{\epsilon}\right)\frac{1}{\sqrt{n}}\right) + \mathcal{O}\left(\frac{1}{n}\right) \end{aligned} \tag{5.6.24}$$

Similarly to Theorem 3.3.4, we now have that

$$c_\gamma(\mathcal{E}^{\otimes n}(\nu)\|\mathcal{F}^{\otimes n}(\nu)) \leq \widehat{c}_\gamma(\mathcal{E}^{\otimes n}\|\mathcal{F}^{\otimes n}) = n\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) \leq n\widehat{c}(\mathcal{E}\|\mathcal{F}) \tag{5.6.25}$$

where

$$\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F}) := \frac{1}{\gamma}\log\left(2^{\gamma\widehat{D}_{1+\gamma}(\mathcal{E}\|\mathcal{F})} + 2\right) \qquad \widehat{c}(\mathcal{E}\|\mathcal{F}) := \inf_{\gamma\in(0,1]}\widehat{c}_\gamma(\mathcal{E}\|\mathcal{F})\,. \tag{5.6.26}$$

and the equality in the first line follows from the chain rule for the geometric Rényi divergence (Theorem 5.3.6), which implies additivity of the channel divergence. Now, by assumption, $\widehat{c}(\mathcal{E}\|\mathcal{F})$ is finite. Hence, we can take $m = n^3$ and then in the limit $n \to \infty$ all error terms will disappear, so that we get the desired expression. □

Remark 5.6.6. In the updated version of O. Fawzi, Gao, and Rahaman (2023) the authors prove this chain rule under the stronger condition $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$, and hence Theorem 5.6.5 can be seen as an improvement of their result. Please see our following remarks on why this less restrictive finiteness condition can be relevant.

5.6.3 Finiteness Condition

As mentioned previously, we are able to show the equivalence of adaptive and parallel discrimination strategies in infinite dimensions only under the condition that $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$, whereas in finite dimensions no such condition was needed. The reason for this is that in finite dimensions, most relative entropies are infinite at the same time, and only if a support-condition is violated. Specifically,

$$D(\rho\|\sigma) = \infty \Leftrightarrow \mathbb{D}_\alpha(\rho\|\sigma) = \infty\ \forall\alpha > 1 \Leftrightarrow D_{\max}(\rho\|\sigma) = \infty \Leftrightarrow \rho \not\ll \sigma \tag{5.6.27}$$

where $\mathbb{D}_\alpha$ is any quantum α-Rényi divergence (i.e. any function that satisfies the data-processing inequality (DPI) and reduces to the classical α-Rényi divergence on commuting states). Hence, if our finiteness condition $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ is violated in finite dimensions, also $D(\mathcal{E}\|\mathcal{F}) = \infty$, which then implies that a parallel strategy (even a product strategy that just repeats the same input state) can achieve an asymmetric error exponent of infinity, and hence there cannot be any asymptotic adaptive advantage.

In infinite dimensions, however, while the divergences in (5.6.27) will also all be infinite if the support condition is violated, this is not the only possibility, and additionally infinity will occur if "one of the two states decays faster than the other". This appears already classically. As an example, consider the following

three (unnormalized) classical probability distributions on $\mathbb{N}$: $p = \{p_n\}, q = \{q_n\}$ and $r = \{p_n\}$, with

$$p_n = \frac{2^{-n}}{n^2} \qquad q_n = \frac{p_n}{n} \qquad r_n = p_n 2^{-2^n} \tag{5.6.28}$$

We could easily make these three probability distributions normalized by redefining them with a normalization factor, however this does not change the finiteness of the divergences, so we omit the normalization factors to keep the following calculations simpler. We have $p_n, q_n, r_n > 0$ for all n, and hence they satisfy the support condition. On the other hand, it is easy to see that $p_n/q_n = n$ is unbounded, and hence $D_{\max}(p\|q) = \infty$, whereas

$$D_\alpha(p\|q) = \frac{1}{\alpha - 1} \log \sum_n q_n \left(\frac{p_n}{q_n}\right)^\alpha = \frac{1}{\alpha - 1} \log \sum_n 2^{-n} n^{\alpha - 3} \tag{5.6.29}$$

is finite for all $\alpha > 1$. Similarly,

$$D_\alpha(p\|r) = \frac{1}{\alpha - 1} \log \sum_n 2^{(\alpha-1)2^n - n} n^{-2} \tag{5.6.30}$$

is clearly infinite for all $\alpha > 1$, whereas

$$D(p\|r) = \sum_n p_n \log \left(\frac{p_n}{r_n}\right) = \sum_n \frac{1}{n^2}(1 - 2^{-n}) \tag{5.6.31}$$

is finite. Roughly speaking, in infinite dimensions the finiteness of these divergences is governed by how the tails of the distributions p and q are related to each other. Finiteness of $D_{\max}$ requires q_n to be at most a constant factor smaller than p_n. Assuming, that p_n is such that also the sequence $p'_n = p_n n^\beta$ is summable for some $\beta > 0$, then $D_\alpha(p\|q)$ will be finite for some $\alpha > 1$ if q_n is at most a polynomial in n smaller than p_n (i.e. $p_n/q_n = \mathcal{O}(n^{\beta/\alpha})$), and $D(p\|q)$ will be finite if q_n is at most exponentially smaller than p_n, i.e. $p_n/q_n = \mathcal{O}(2^{n^\beta})$.

Different conditions for quantum channel discrimination

Given that in infinite dimensions finiteness conditions involving different relative entropies are no longer equivalent, we would like to demonstrate, in this section and the next, that there are examples of quantum channels (which are interesting in the context of channel discrimination) that satisfy some but not all of the finiteness conditions. Specifically, given two quantum channels $\mathcal{E}$ and $\mathcal{F}$ we will be looking at the the following three finiteness conditions

(a) $D_{\max}(\mathcal{E}\|\mathcal{F}) < \infty$

(b) $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$

(c) $D^{\text{reg}}(\mathcal{E}\|\mathcal{F}) < \infty$

In the context of comparing adaptive and parallel strategies, the third condition would be the desired one, as it amounts to imposing no condition at all. This follows from the fact that $D^{\text{reg}} = \infty$ implies that both

adaptive and parallel strategies are able to achieve a rate of infinity, and so one then obtains asymptotic equivalence of adaptive and parallel strategies for all channels.

By choosing replacer channels (i.e. channels that output a specific state independent of the input) outputting classical states corresponding to the probability distributions p, q and r defined above, we see that there exist classical channels for which these three finiteness conditions are not equivalent. However, these channels are not particularly interesting from the perspective of channel discrimination and for studying the relation between adaptive and parallel strategies. This is because for replacer channels, the input state and hence also the chosen strategy is irrelevant already on the one-shot level. Additionally, for classical channels also on continuous systems (which form the classical analogue of an infite dimensional quantum systems) it is known that adaptivity does not give any asymptotic advantage, without requiring any finiteness conditions (Hayashi 2009).

Hence, in the next subsection, we give examples of fully quantum channels that are not replacer channels for which the three finiteness conditions are not equivalent. This illustrates first, that there are interesting channels satisfying only condition (b) but not (a), and hence being able to establish the equivalence of adaptive and parallel strategies under condition (b) is significant. Additionally, we find channels for which only condition (c) holds and for which we are currently unable to show the equivalence.

We conjecture that, in fact, no condition is necessary also in infinite dimensions, and the asymptotic equivalence of adaptive and parallel strategies does hold for all channels also in infinite dimensions. However, we currently do not have the necessary tools to prove this conjecture. We would also like to highlight that we believe the condition (a) to be very restrictive in practice. Specifically, with this condition the data-processing inequality implies that in *all bases* the diagonal elements of $\mathcal{E}(\nu)$ and $\mathcal{F}(\nu)$ have to decay while differing at most by a constant factor, and this too for all input states ν. For channels which do not just output states on some finite-dimensional subspace we expect this to be the case only for very specific examples.

Fully quantum examples

For any state $\tau \in \mathcal{D}(\mathcal{H}_A)$ and any $\lambda \in [0,1]$, define the generalized depolarizing channel $\Lambda_\tau^\lambda : \mathscr{P}(\mathcal{H}_A) \to \mathscr{P}(\mathcal{H}_A)$ as

$$\Lambda_\tau^\lambda(\omega) \coloneqq (1-\lambda)\omega + \lambda \operatorname{Tr}(\omega)\tau\,. \tag{5.6.32}$$

Lemma 5.6.7. *For any two states $\rho, \sigma \in \mathcal{D}(\mathcal{H}_A)$, any $\lambda \in [0,1]$ and any quantum divergence $\mathbf{D}$ (i.e. any function of two positive trace-class operators that satisfies the data-processing inequality) and any $n \in \mathbb{N}$, it holds that*

$$\mathbf{D}((\Lambda_\rho^\lambda)^{\otimes n} \| (\Lambda_\sigma^\lambda)^{\otimes n}) \leq \mathbf{D}(\rho^{\otimes n} \| \sigma^{\otimes n}) \tag{5.6.33}$$

where the channel divergence on the left-hand side is defined as in (5.3.1).

Proof. For the sake of illustration, consider first the case $n = 1$. For any normalized density matrix $\omega = \omega_{RA}$

we then have

$$\begin{aligned}\mathbf{D}(\rho\|\sigma) &= \mathbf{D}(\omega_R\otimes\rho\|\omega_R\otimes\sigma) && (5.6.34)\\ &\geq \mathbf{D}(\Lambda^{1-\lambda}_{\omega_{RA}}(\omega_R\otimes\rho)\|\Lambda^{1-\lambda}_{\omega_{RA}}(\omega_R\otimes\sigma)) && (5.6.35)\\ &= \mathbf{D}(\lambda(\omega_R\otimes\rho)+(1-\lambda)\omega_{RA}\|\lambda(\omega_R\otimes\sigma)+(1-\lambda)\omega_{RA}) && (5.6.36)\\ &= \mathbf{D}((\mathrm{id}_R\otimes\Lambda^{\lambda}_{\rho})(\omega_{RA})\|(\mathrm{id}_R\otimes\Lambda^{\lambda}_{\sigma})(\omega_{RA})), && (5.6.37)\end{aligned}$$

where the first line also follows from the data-processing inequality, by using channels that add or trace out the additional state ω_R. As this holds for all states ω_{RA}, it then also holds for the supremum over all states.

For $n>1$, the argument is essentially the same. For a subset $S\subset\{1,2,...,n\}$ and a state $\tau_{A^n}=\tau_{A_1\ldots A_n}$ we write τ_{A_S} for $\mathrm{Tr}_{A_{S^c}}(\tau_{A^n})$, where we trace out all the systems A_i whose index is not in S, and S^c is the complement of S. Given any state $\omega_{R^nA^n}$, we can then define a channel $\mathcal{M}:A^n\to R^nA^n$ via

$$\mathcal{M}(\tau_{A^n}) \coloneqq \sum_{S\subset\{1,\ldots,n\}} \lambda^{|S|}(1-\lambda)^{|S^c|}\tau_{A_S}\otimes\omega_{R_S}\otimes\omega_{R_{S^c}A_{S^c}}\,, \tag{5.6.38}$$

with the idea being that

$$\begin{aligned}(\mathrm{id}_R\otimes\Lambda^{\lambda}_{\rho})^{\otimes n}(\omega_{R^nA^n}) &= \mathcal{M}(\rho^{\otimes n}) && (5.6.39)\\ (\mathrm{id}_R\otimes\Lambda^{\lambda}_{\sigma})^{\otimes n}(\omega_{R^nA^n}) &= \mathcal{M}(\sigma^{\otimes n}) && (5.6.40)\end{aligned}$$

and hence the claim follows again from an application of the data-processing inequality:

$$\mathbf{D}(\rho^{\otimes n}\|\sigma^{\otimes n})\geq\mathbf{D}(\mathcal{M}(\rho^{\otimes n})\|\mathcal{M}(\sigma^{\otimes n}))=\mathbf{D}((\mathrm{id}_R\otimes\Lambda^{\lambda}_{\rho})^{\otimes n}(\omega_{R^nA^n})\|(\mathrm{id}_R\otimes\Lambda^{\lambda}_{\sigma})^{\otimes n}(\omega_{R^nA^n}))\,. \tag{5.6.41}$$

□

Note specifically that for any divergence $\mathbf{D}$ that is additive on states (e.g. $\mathbf{D}=D$), this implies

$$\mathbf{D}^{\mathrm{reg}}(\Lambda^{\lambda}_{\rho}\|\Lambda^{\lambda}_{\sigma})\leq\mathbf{D}(\rho\|\sigma)\,. \tag{5.6.42}$$

To continue constructing our examples, given any orthonormal basis $\{|a_i\rangle\}_{i=0}^{\infty}$ of a Hilbert space, we construct the basis $\{|b_i\rangle\}_{i=0}^{\infty}$ as

$$|b_i\rangle \coloneqq \begin{cases}\frac{1}{\sqrt{2}}(|a_i\rangle+|a_{i+1}\rangle) & i\text{ even}\\ \frac{1}{\sqrt{2}}(|a_{i-1}\rangle-|a_i\rangle) & i\text{ odd}\,.\end{cases} \tag{5.6.43}$$

For any classical probability distribution $p\in\ell_1$ we then define the states

$$\rho_p\coloneqq\sum_i p_i\,|a_i\rangle\langle a_i| \qquad \sigma_p\coloneqq\sum_i p_i\,|b_i\rangle\langle b_i|\,. \tag{5.6.44}$$

Lemma 5.6.8. *Let* $\mathbf{D}$ *be one of* $D, \widehat{D}_\alpha$ *(with* $\alpha \in (1,2]$*) or* $D_{\max}$*, and* $p, q \in \ell_1$ *be normalized probability distributions. Then, for all* $\lambda \in (0,1]$

$$\mathbf{D}(\rho_p \| \sigma_q) = \infty \Rightarrow \mathbf{D}(\Lambda^\lambda_{\rho_p} \| \Lambda^\lambda_{\sigma_q}) = \infty\,. \tag{5.6.45}$$

Proof. Throughout this proof we will write $\rho \coloneqq \rho_p$ and $\sigma \coloneqq \sigma_q$ for simplicity. Let us start with just choosing the reference system of the input state as trivial and picking the input state $\omega = \sigma$. Then

$$\mathbf{D}(\Lambda^\lambda_\rho(\omega) \| \Lambda^\lambda_\sigma(\omega)) = \mathbf{D}((1-\lambda)\sigma + \lambda\rho \| \sigma) \tag{5.6.46}$$

If $\mathbf{D} = D_{\max}$ we can use the monotonicity of $D_{\max}$ in the first variable to conclude that

$$\mathbf{D}((1-\lambda)\sigma + \lambda\rho \| \sigma) \geq \mathbf{D}(\lambda\rho \| \sigma) = \infty\,. \tag{5.6.47}$$

If $\mathbf{D} = D$, we can use the almost-concavity of D in the first argument (we provide a proof in Lemma 5.3.2) to find

$$\mathbf{D}((1-\lambda)\sigma + \lambda\rho \| \sigma) \geq (1-\lambda)\mathbf{D}(\sigma\|\sigma) + \lambda \mathbf{D}(\rho\|\sigma) - h(\lambda) = \infty \tag{5.6.48}$$

where h is the binary entropy. Note that so far we have not used any of the structure of the states ρ and σ. For $\mathbf{D} = \widehat{D}_\alpha$ we could use a similar monotonicity argument for $\alpha \in (0,1]$, but for $\alpha > 1$, the function $t \mapsto t^\alpha$ is unfortunately not operator monotone, so we need a different argument, which is where the assumptions on the states come in. Let j be such that $q_j > 0$ and then pick the state $\omega = |b_j\rangle\langle b_j|$. Let also k_1, k_2 be the odd and even index of the block associated to j, i.e. $k_1 = 2\left\lfloor \frac{j}{2} \right\rfloor$, $k_2 = 2\left\lfloor \frac{j}{2} \right\rfloor + 1$. Consider the subspace of our Hilbert space spanned by $\{|a_i\rangle\}_{i=k_1,k_2}$, and its orthogonal complement (this is the same as the span of $\{|b_i\rangle\}_{i=k_1,k_2}$ and its orthogonal complement), and let $\mathcal{M}$ be the POVM measurement channel between these two subspaces, i.e. $\mathcal{M}(\nu) = \sum_{i=1,2} \Pi_i \nu \Pi_i$, where the Π_1 projects onto the subspace, and Π_2 projects onto its orthogonal complement. Then,

$$\begin{aligned}
\mathbf{D}(\Lambda^\lambda_\rho(\omega) \| \Lambda^\lambda_\sigma(\omega)) &\geq \mathbf{D}(\mathcal{M} \circ \Lambda^\lambda_\rho(\omega) \| \mathcal{M} \circ \Lambda^\lambda_\sigma(\omega)) && (5.6.49)\\
&\geq \mathbf{D}(\Pi_2 \Lambda^\lambda_\rho(\omega) \Pi_2 \| \Pi_2 \Lambda^\lambda_\sigma(\omega) \Pi_2) && (5.6.50)\\
&= \mathbf{D}(\lambda \Pi_2 \rho \Pi_2 \| \lambda \Pi_2 \sigma \Pi_2), && (5.6.51)
\end{aligned}$$

where we used the data-processing inequality, Lemma 5.3.13, and the fact that ω lies in the kernel of Π_2. Since $\mathbf{D}(\rho\|\sigma)$ is infinite and ρ and σ are block-diagonal in the decomposition, either the term in (5.6.51) has to be infinite (which proves our statement), or $\mathbf{D}(\Pi_1\rho\Pi_1\|\Pi_1\sigma\Pi_1)$ is infinite, which (since the subspace is finite-dimensional) implies that $\Pi_1\rho\Pi_1 \not\ll \Pi_1\sigma\Pi_1$, which implies $\rho \not\ll \sigma$. Since $q_j > 0$, we have $\omega \ll \sigma$, and hence this also implies $\Lambda^\lambda_\rho(\omega) \not\ll \Lambda^\lambda_\sigma(\omega)$ and so also the channel divergence has to be infinite. □

Lemma 5.6.9. *Let* $\mathbf{D}$ *be one of* $D, \widehat{D}_\alpha$ *(with* $\alpha \in (1,2]$*) or* $D_{\max}$*. Let* $p, q \in \ell_1$ *be positive, and define the following variants of q, which take the minimum/maximum over a block of two indices:*

$$q_i^\uparrow \coloneqq \max\{q_{2\lfloor \frac{i}{2} \rfloor}, q_{2\lfloor \frac{i}{2} \rfloor + 1}\} \qquad q_i^\downarrow \coloneqq \min\{q_{2\lfloor \frac{i}{2} \rfloor}, q_{2\lfloor \frac{i}{2} \rfloor + 1}\}\,. \tag{5.6.52}$$

Then,

$$\mathbf{D}(p\|q^{\uparrow}) \leq \mathbf{D}(\rho_p\|\sigma_q) \leq \mathbf{D}(p\|q^{\downarrow})\,. \tag{5.6.53}$$

Proof. All the three divergences we consider are anti-monotonous in the second variable. For $D_{\max}$ this is obvious from the definition, for D this is shown in Hiai (2019), Theorem 4.1, and for $\widehat{D}_\alpha$ we show it in Lemma 5.3.12. It is easy to see that $\sigma_{q^\downarrow} \leq \sigma_q \leq \sigma_{q^\uparrow}$, but $\sigma_{q^\downarrow} = \rho_{q^\downarrow}$, and $\sigma_{q^\uparrow} = \rho_{q^\uparrow}$, which implies the desired statement. □

The probability distributions p, q, r defined in the previous section are such that the three terms in (5.6.53) are all either finite or infinite at the same time. Hence, all these lemmas together imply that for $\mathbf{D}$ one of $D, \widehat{D}_\alpha$ (with $\alpha \in (1,2]$) or $D_{\max}$, $\lambda \in (0,1]$ and p, q one of these three probability distributions it holds that

$$\mathbf{D}(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_q}) = \infty \Leftrightarrow \mathbf{D}(p\|q) = \infty\,. \tag{5.6.54}$$

Additionally, since $D \leq D^{\mathrm{reg}}$ and the upper bound in Lemma 5.6.7 also includes the statement for tensor products of the channels, we get

$$D^{\mathrm{reg}}(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_q}) = \infty \Leftrightarrow D(p\|q) = \infty\,. \tag{5.6.55}$$

Specifically, for $\lambda \in (0,1]$ and

$$p_n = \frac{2^{-n}n^{-2}}{\sum_k 2^{-k}k^{-2}} \qquad q_n = \frac{2^{-n}n^{-3}}{\sum_k 2^{-k}k^{-3}} \tag{5.6.56}$$

we have $\widehat{D}_\alpha(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_q}) < \infty$ for all $\alpha \in (1,2]$, while $D_{\max}(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_q}) = \infty$. Similarly, for

$$r_n = \frac{2^{-n}2^{-2^n}n^{-2}}{\sum_k 2^{-k}2^{-2^k}k^{-2}} \tag{5.6.57}$$

we have $D^{\mathrm{reg}}(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_r}) < \infty$, while $\widehat{D}_\alpha(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_r}) = \infty$ for all $\alpha \in (1,2]$ (and hence also for all $\alpha > 1$, as the geometric Rényi divergence is easily seen to be increasing in α) and also $D_{\max}(\Lambda^\lambda_{\rho_p}\|\Lambda^\lambda_{\sigma_q}) = \infty$.

5.7 Outlook

In this final section we discuss certain limitations of our results, and ways in which they could be extended.

Recall that we stated our main theorems under the condition that our two channels satisfy $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ for some $\alpha > 1$. As discussed in Section 5.6.3, this condition is not required when restricting to classical channels or finite-dimensional quantum channels. We would hence conjecture that it is unnecessary also in the infinite-dimensional quantum case, however we do not currently know how one would be able to prove that. In particular, requiring no condition at all would imply that one can not use the well-established approach of bounding the desired quantity by a Rényi divergence for $\alpha > 1$, and then taking a limit $\alpha \downarrow 1$, as this limit in infinite dimensions converges to the quantum relative entropy only under a finiteness condition,

which is exactly what one wants to avoid. Hence, proving such a statement seems to require a fairly different approach.

In addition to maybe relaxing the restrictions, it would already be interesting to obtain simple characterizations of channels for which $\widehat{D}_\alpha(\mathcal{E}\|\mathcal{F}) < \infty$ holds for some $\alpha > 1$. In Section 5.6.3 we establish a simple condition for generalized depolarizing channels, but we leave finding such conditions for more general channels open to further work.

Finally, we have shown our results in the setting of a separable Hilbert spaces, and not the more general setting of von Neumann algebras, where quantum hypothesis testing and state discrimination have also already been studied (Hiai and Petz 1991, O. Fawzi, Gao, and Rahaman 2023, Berta, Scholz, and Tomamichel 2018, Jenčová 2018, 2021, Jakšić et al. 2012, Mosonyi 2023). The main motivation for this was that many of our technical ingredients are considerably easier to generalize to the infinite-dimensional Hilbert space setting, and proving the equivalent statements in von Neumann algebras is often not obvious. Nevertheless, we have tried to cite results from the von Neumann algebra literature wherever we were aware of them, and often employed tools and techniques which are also available in that setting to facilitate such a future generalization.

6 Bibliography

Aasi, J. et al. (Aug. 2013). "Enhanced Sensitivity of the LIGO Gravitational Wave Detector by Using Squeezed States of Light". In: *Nature Photonics* 7.8, pp. 613–619. DOI: 10.1038/nphoton.2013.177.

Alberti, Peter M. and Armin Uhlmann (Mar. 1, 1983). "Stochastic Linear Maps and Transition Probability". In: *Letters in Mathematical Physics* 7.2, pp. 107–112. DOI: 10.1007/BF00419927.

Androulakis, George and Tiju Cherian John (July 13, 2023). "Quantum F-Divergences via Nussbaum–Szkoła Distributions and Applications to f-Divergence Inequalities". In: *Reviews in Mathematical Physics*, p. 2360002. DOI: 10.1142/S0129055X23600024.

Anjos, Miguel F. and Jean B. Lasserre, eds. (2012). *Handbook on Semidefinite, Conic and Polynomial Optimization*. Vol. 166. International Series in Operations Research & Management Science. Boston, MA: Springer US. ISBN: 978-1-4614-0768-3. DOI: 10.1007/978-1-4614-0769-0.

Anshu, Anurag, Mario Berta, Rahul Jain, and Marco Tomamichel (Dec. 1, 2019). "A Minimax Approach to One-Shot Entropy Inequalities". In: *Journal of Mathematical Physics* 60.12, p. 122201. DOI: 10.1063/1.5126723. arXiv: 1906.00333v1.

Araki, Huzihiro (Dec. 31, 1975). "Relative Entropy of States of von Neumann Algebras". In: *Publications of the Research Institute for Mathematical Sciences* 11.3, pp. 809–833. DOI: 10.2977/prims/1195191148.

Araki, Huzihiro (Apr. 30, 1977). "Relative Entropy for States of von Neumann Algebras II". In: *Publications of the Research Institute for Mathematical Sciences* 13.1, pp. 173–192. DOI: 10.2977/prims/1195190105.

Araujo, Leonardo C., João P. H. Sansão, and Adriano S. Vale-Cardoso (Oct. 1, 2021). "Fast Computation of Multinomial Coefficients". In: *Numerical Algorithms* 88.2, pp. 837–851. DOI: 10.1007/s11075-020-01059-5.

Araújo, Mateus, Cyril Branciard, Fabio Costa, Adrien Feix, Christina Giarmatzi, and Časlav Brukner (Oct. 2015). "Witnessing Causal Nonseparability". In: *New Journal of Physics* 17.10, p. 102001. DOI: 10.1088/1367-2630/17/10/102001.

Audenaert, Koenraad M. R., J. Calsamiglia, R. Muñoz-Tapia, E. Bagan, Ll. Masanes, A. Acin, and F. Verstraete (Apr. 17, 2007). "Discriminating States: The Quantum Chernoff Bound". In: *Physical Review Letters* 98.16, p. 160501. DOI: 10.1103/PhysRevLett.98.160501. arXiv: quant-ph/0610027.

Audenaert, Koenraad M. R., Milán Mosonyi, and Frank Verstraete (Dec. 2012). "Quantum State Discrimination Bounds for Finite Sample Size". In: *Journal of Mathematical Physics* 53.12, p. 122205. DOI: 10.1063/1.4768252. arXiv: 1204.0711.

Audenaert, Koenraad M. R., M. Nussbaum, A. Szkola, and F. Verstraete (Apr. 2008). "Asymptotic Error Rates in Quantum Hypothesis Testing". In: *Communications in Mathematical Physics* 279.1, pp. 251–283. DOI: 10.1007/s00220-008-0417-5. arXiv: 0708.4282.

Bae, Joonwoo and Leong-Chuan Kwek (Feb. 27, 2015). "Quantum State Discrimination and Its Applications". In: *Journal of Physics A: Mathematical and Theoretical* 48.8, p. 083001. DOI: 10.1088/1751-8113/48/8/083001. arXiv: 1707.02571.

Bavaresco, Jessica, Mio Murao, and Marco Túlio Quintino (Nov. 12, 2021). "Strict Hierarchy between Parallel, Sequential, and Indefinite-Causal-Order Strategies for Channel Discrimination". In: *Physical Review Letters* 127.20, p. 200504. DOI: 10.1103/PhysRevLett.127.200504. arXiv: 2011.08300.

Beigi, Salman (Dec. 2013). "Sandwiched R\'enyi Divergence Satisfies Data Processing Inequality". In: *Journal of Mathematical Physics* 54.12, p. 122202. DOI: 10.1063/1.4838855. arXiv: 1306.5920.

Belavkin, V. P. and P. Staszewski (1982). "C*-algebraic generalization of relative entropy and entropy". In: *Annales de l'I.H.P. Physique théorique* 37.1, pp. 51–58. URL: http://www.numdam.org/item/AIHPA_1982__37_1_51_0/.

Bergh, Bjarne, Nilanjana Datta, and Robert Salzmann (Mar. 3, 2023). *Composite Classical and Quantum Channel Discrimination*. arXiv: 2303.02016. Pre-published.

Bergh, Bjarne, Nilanjana Datta, Robert Salzmann, and Mark M. Wilde (Apr. 2024). "Parallelization of Adaptive Quantum Channel Discrimination in the Non-Asymptotic Regime". In: *IEEE Transactions on Information Theory* 70.4, pp. 2617–2636. DOI: 10.1109/TIT.2024.3355929. arXiv: 2206.08350.

Bergh, Bjarne, Jan Kochanowski, Robert Salzmann, and Nilanjana Datta (Aug. 24, 2023). *Infinite Dimensional Asymmetric Quantum Channel Discrimination*. arXiv: 2308.12959. Pre-published.

Bergh, Bjarne, Robert Salzmann, and Nilanjana Datta (Sept. 1, 2021). "The $\alpha \to 1$ limit of the sharp quantum Rényi divergence". In: *Journal of Mathematical Physics* 62.9, p. 092205. DOI: 10.1063/5.0049791. arXiv: 2102.06576.

Berta, Mario, Fernando G. S. L. Brandao, and Christoph Hirche (July 2021). "On Composite Quantum Hypothesis Testing". In: *Communications in Mathematical Physics* 385.1, pp. 55–77. DOI: 10.1007/s00220-021-04133-8. arXiv: 1709.07268.

Berta, Mario, Fernando G. S. L. Brandão, Gilad Gour, Ludovico Lami, Martin B. Plenio, Bartosz Regula, and Marco Tomamichel (Sept. 7, 2023). "On a Gap in the Proof of the Generalised Quantum Stein's Lemma and Its Consequences for the Reversibility of Quantum Resources". In: *Quantum* 7, p. 1103. DOI: 10.22331/q-2023-09-07-1103. arXiv: 2205.02813.

Berta, Mario, Omar Fawzi, and Marco Tomamichel (Dec. 2017). "On Variational Expressions for Quantum Relative Entropies". In: *Letters in Mathematical Physics* 107.12, pp. 2239–2265. DOI: 10.1007/s11005-017-0990-7. arXiv: 1512.02615.

Berta, Mario, Volkher B. Scholz, and Marco Tomamichel (June 1, 2018). "Rényi Divergences as Weighted Non-commutative Vector-Valued L_p-Spaces". In: *Annales Henri Poincaré* 19.6, pp. 1843–1867. DOI: 10.1007/s00023-018-0670-x.

Berta, Mario and Marco Tomamichel (June 26, 2022). "Chain Rules for Quantum Channels". In: *2022 IEEE International Symposium on Information Theory (ISIT)*, pp. 2427–2432. DOI: 10.1109/ISIT50566.2022.9834391. arXiv: 2204.11153.

Brandao, Fernando G. S. L. and Martin B. Plenio (May 2010). "A Generalization of Quantum Stein's Lemma". In: *Communications in Mathematical Physics* 295.3, pp. 791–828. DOI: 10.1007/s00220-010-1005-z. arXiv: 0904.0281.

Brandão, Fernando G. S. L., Aram W. Harrow, James R. Lee, and Yuval Peres (Aug. 2020). "Adversarial Hypothesis Testing and a Quantum Stein's Lemma for Restricted Measurements". In: *IEEE Transactions on Information Theory* 66.8, pp. 5037–5054. DOI: 10.1109/TIT.2020.2979704. arXiv: 1308.6702.

Braunstein, Samuel L. and Peter van Loock (June 29, 2005). "Quantum Information with Continuous Variables". In: *Reviews of Modern Physics* 77.2, pp. 513–577. DOI: 10.1103/RevModPhys.77.513.

Chiribella, Giulio (Oct. 10, 2012). "Perfect Discrimination of No-Signalling Channels via Quantum Superposition of Causal Structures". In: *Physical Review A* 86.4, p. 040301. DOI: 10.1103/PhysRevA.86.040301. arXiv: 1109.5154.

Chiribella, Giulio, G. M. D'Ariano, P. Perinotti, and B. Valiron (Aug. 14, 2013). "Quantum Computations without Definite Causal Structure". In: *Physical Review A* 88.2, p. 022318. DOI: 10.1103/PhysRevA.88.022318. arXiv: 0912.0195.

Chiribella, Giulio, Giacomo M. D'Ariano, and Paolo Perinotti (Oct. 27, 2008). "Memory Effects in Quantum Channel Discrimination". In: *Physical Review Letters* 101.18, p. 180501. DOI: 10.1103/PhysRevLett.101.180501. arXiv: 0803.3237.

Chiribella, Giulio and Daniel Ebler (Sept. 29, 2016). "Optimal Quantum Networks and One-Shot Entropies". In: *New Journal of Physics* 18.9, p. 093053. DOI: 10.1088/1367-2630/18/9/093053. arXiv: 1606.02394.

Christandl, Matthias, Robert Koenig, Graeme Mitchison, and Renato Renner (June 4, 2007). "One-and-a-Half Quantum de Finetti Theorems". In: *Communications in Mathematical Physics* 273.2, pp. 473–498. DOI: 10.1007/s00220-007-0189-3. arXiv: quant-ph/0602130.

Clarke, Frank H. (1975). "Generalized Gradients and Applications". In: *Transactions of the American Mathematical Society* 205, pp. 247–262. DOI: 10.1090/S0002-9947-1975-0367131-6.

Cooney, Tom, Milán Mosonyi, and Mark M. Wilde (June 1, 2016). "Strong Converse Exponents for a Quantum Channel Discrimination Problem and Quantum-Feedback-Assisted Communication". In: *Communications in Mathematical Physics* 344.3, pp. 797–829. DOI: 10.1007/s00220-016-2645-4. arXiv: 1408.3373.

Cover, Thomas M. and Joy A. Thomas (Nov. 28, 2012). *Elements of Information Theory*. John Wiley & Sons. 788 pp. ISBN: 978-1-118-58577-1. Google Books: VWq5GG6ycxMC.

Datta, Nilanjana (June 2009). "Min- and Max- Relative Entropies and a New Entanglement Monotone". In: *IEEE Transactions on Information Theory* 55.6, pp. 2816–2826. DOI: 10.1109/TIT.2009.2018325. arXiv: 0803.2770.

Datta, Nilanjana, Milán Mosonyi, Min-Hsiu Hsieh, and Fernando G. S. L. Brandão (Dec. 2013). "A Smooth Entropy Approach to Quantum Hypothesis Testing and the Classical Capacity of Quantum Channels". In: *IEEE Transactions on Information Theory* 59.12, pp. 8014–8026. DOI: 10.1109/TIT.2013.2282160. arXiv: 1106.3089.

Davies, Edward Brian (1976). *Quantum Theory of Open Systems*. Academic Press. 192 pp. ISBN: 978-0-12-206150-9.

De Klerk, Etienne, Dmitrii V. Pasechnik, and Alexander Schrijver (Mar. 1, 2007). "Reduction of Symmetric Semidefinite Programs Using the Regular *-Representation". In: *Mathematical Programming* 109.2, pp. 613–624. DOI: 10.1007/s10107-006-0039-7.

De Klerk, Etienne (2002). *Aspects of Semidefinite Programming*. Applied Optimization. Boston, MA: Springer US. ISBN: 978-0-306-47819-2. DOI: 10.1007/b105286.

Dem'yanov, V. F. and V. N. Malozemov (June 30, 1971). "On the Theory of Non-Linear Minimax Problems". In: *Russian Mathematical Surveys* 26.3, p. 57. DOI: 10.1070/RM1971v026n03ABEH003834.

Díaz, María García, Kun Fang, Xin Wang, Matteo Rosati, Michalis Skotiniotis, John Calsamiglia, and Andreas Winter (Oct. 19, 2018). "Using and Reusing Coherence to Realize Quantum Processes". In: *Quantum* 2, p. 100. DOI: 10.22331/q-2018-10-19-100. arXiv: 1805.04045.

Douglas, R. G. (1966). "On Majorization, Factorization, and Range Inclusion of Operators on Hilbert Space". In: *Proceedings of the American Mathematical Society* 17.2, pp. 413–415. DOI: 10.1090/S0002-9939-1966-0203464-1.

Duan, Runyao, Yuan Feng, and Mingsheng Ying (Nov. 20, 2009). "Perfect Distinguishability of Quantum Operations". In: *Physical Review Letters* 103.21, p. 210501. DOI: 10.1103/PhysRevLett.103.210501.

Fang, Kun and Hamza Fawzi (June 2021). "Geometric Rényi Divergence and Its Applications in Quantum Channel Capacities". In: *Communications in Mathematical Physics* 384.3, pp. 1615–1677. DOI: 10.1007/s00220-021-04064-4. arXiv: 1909.05758.

Fang, Kun, Omar Fawzi, Renato Renner, and David Sutter (Mar. 10, 2020). "Chain Rule for the Quantum Relative Entropy". In: *Physical Review Letters* 124.10, p. 100501. DOI: 10.1103/PhysRevLett.124.100501. arXiv: 1909.05826.

Fang, Kun, Gilad Gour, and Xin Wang (Mar. 1, 2022). *Towards the Ultimate Limits of Quantum Channel Discrimination*. arXiv: 2110.14842. Pre-published.

Farkas, Balint and Szilard Gy Revesz (2006). "Potential Theoretic Approach to Rendezvous Numbers". In: *Monatshefte fur Mathematik*, 148, pp. 309–331. DOI: 10.48550/arXiv.math/0503423. arXiv: math/0503423.

Fawzi, Hamza and Omar Fawzi (Jan. 26, 2021). "Defining Quantum Divergences via Convex Optimization". In: *Quantum* 5, p. 387. DOI: 10.22331/q-2021-01-26-387. arXiv: 2007.12576.

Fawzi, Omar, Li Gao, and Mizanur Rahaman (May 16, 2023). *Asymptotic Equipartition Theorems in von Neumann Algebras*. arXiv: 2212.14700v2. Pre-published.

Fawzi, Omar, Ala Shayeghi, and Hoang Ta (Nov. 2022). "A Hierarchy of Efficient Bounds on Quantum Capacities Exploiting Symmetry". In: *IEEE Transactions on Information Theory* 68.11, pp. 7346–7360. DOI: 10.1109/TIT.2022.3182101. arXiv: 2203.02127.

Flurin, E., V. V. Ramasesh, S. Hacohen-Gourgy, L. S. Martin, N. Y. Yao, and I. Siddiqi (Aug. 3, 2017). "Observing Topological Invariants Using Quantum Walks in Superconducting Circuits". In: *Physical Review X* 7.3, p. 031023. DOI: 10.1103/PhysRevX.7.031023.

Frank, Rupert L. and Elliott H. Lieb (Dec. 2013). "Monotonicity of a Relative Renyi Entropy". In: *Journal of Mathematical Physics* 54.12, p. 122201. DOI: 10.1063/1.4838835. arXiv: 1306.5358.

Fuchs, Christopher A. and Jeroen van de Graaf (1999). "Cryptographic Distinguishability Measures for Quantum Mechanical States". In: *IEEE Transactions on Information Theory* 45.4, pp. 1216–1227. DOI: 10.1109/18.761271. arXiv: quant-ph/9712042.

Gijswijt, D. C. (2005). "Matrix Algebras and Semidefinite Programming Techniques for Codes". University of Amsterdam. arXiv: 1007.0906.

Gilchrist, Alexei, Nathan K. Langford, and Michael A. Nielsen (June 13, 2005). "Distance Measures to Compare Real and Ideal Quantum Processes". In: *Physical Review A* 71.6, p. 062310. DOI: 10.1103/PhysRevA.71.062310. arXiv: quant-ph/0408063.

Gottesman, Daniel, Alexei Kitaev, and John Preskill (June 11, 2001). "Encoding a Qubit in an Oscillator". In: *Physical Review A* 64.1, p. 012310. DOI: 10.1103/PhysRevA.64.012310.

Gour, Gilad (Sept. 2019). "Comparison of Quantum Channels by Superchannels". In: *IEEE Transactions on Information Theory* 65.9, pp. 5880–5904. DOI: 10.1109/TIT.2019.2907989. arXiv: 1808.02607.

Guha, Saikat (May 20, 2008). "Multiple-User Quantum Information Theory for Optical Communication Channels". In: URL: https://dspace.mit.edu/handle/1721.1/41840.

Guillaud, Jérémie and Mazyar Mirrahimi (Dec. 12, 2019). "Repetition Cat Qubits for Fault-Tolerant Quantum Computation". In: *Physical Review X* 9.4, p. 041053. DOI: 10.1103/PhysRevX.9.041053.

Harrow, Aram W., Avinatan Hassidim, Debbie W. Leung, and John Watrous (Mar. 31, 2010). "Adaptive versus Nonadaptive Strategies for Quantum Channel Discrimination". In: *Physical Review A* 81.3, p. 032339. DOI: 10.1103/PhysRevA.81.032339.

Hayashi, Masahito (2006). *Quantum Information: An Introduction*. Berlin: Springer. xiv+424. ISBN: 978-3-540-30265-0.

Hayashi, Masahito (Dec. 5, 2007). "Error Exponent in Asymmetric Quantum Hypothesis Testing and Its Application to Classical-Quantum Channel Coding". In: *Physical Review A* 76.6, p. 062301. DOI: 10.1103/PhysRevA.76.062301. arXiv: quant-ph/0611013.

Hayashi, Masahito (Aug. 2009). "Discrimination of Two Channels by Adaptive Methods and Its Application to Quantum System". In: *IEEE Transactions on Information Theory* 55.8, pp. 3807–3820. DOI: 10.1109/TIT.2009.2023726. arXiv: 0804.0686.

Hayashi, Masahito and Shun Watanabe (Apr. 2016). "Uniform Random Number Generation from Markov Chains: Non-Asymptotic and Asymptotic Analyses". In: *IEEE Transactions on Information Theory* 62.4, pp. 1795–1822. DOI: 10.1109/TIT.2016.2530084. arXiv: 1503.04371.

Hayashi, Masahito and Shun Watanabe (Apr. 2020). "Finite-Length Analyses for Source and Channel Coding on Markov Chains". In: *Entropy* 22.4, p. 460. DOI: 10.3390/e22040460. arXiv: 1309.7528.

Hayashi, Masahito and Hayata Yamasaki (Aug. 5, 2024). *Generalized Quantum Stein's Lemma and Second Law of Quantum Resource Theories*. arXiv: 2408.02722. Pre-published.

Helstrom, Carl W. (June 1, 1969). "Quantum Detection and Estimation Theory". In: *Journal of Statistical Physics* 1.2, pp. 231–252. DOI: 10.1007/BF01007479.

Hiai, Fumio (Sept. 26, 2018). "Quantum F-Divergences in von Neumann Algebras. I. Standard f-Divergences". In: *Journal of Mathematical Physics* 59.10, p. 102202. DOI: 10.1063/1.5039973.

Hiai, Fumio (Jan. 2019). "Quantum F-Divergences in von Neumann Algebras II. Maximal f-Divergences". In: *Journal of Mathematical Physics* 60.1, p. 012203. DOI: 10.1063/1.5051427. arXiv: 1807.03118.

Hiai, Fumio (2021). *Quantum F-Divergences in von Neumann Algebras: Reversibility of Quantum Operations*. Mathematical Physics Studies. Singapore: Springer. ISBN: 978-981-334-199-9. DOI: 10.1007/978-981-33-4199-9.

Hiai, Fumio and Milán Mosonyi (Aug. 2017). "Different Quantum F-Divergences and the Reversibility of Quantum Operations". In: *Reviews in Mathematical Physics* 29.07, p. 1750023. DOI: 10.1142/S0129055X17500234. arXiv: 1604.03089.

Hiai, Fumio and Dénes Petz (Dec. 1, 1991). "The Proper Formula for Relative Entropy and Its Asymptotics in Quantum Probability". In: *Communications in Mathematical Physics* 143.1, pp. 99–114. DOI: 10.1007/BF02100287.

Holevo, A. S. (1972). "An Analogue of Statistical Decision Theory and Noncommutative Probability Theory". In: *Trudy Moskovskogo Matematičeskogo Obščestva* 26, pp. 133–149. URL: https://mathscinet.ams.org/mathscinet-getitem?mr=0365809.

Holevo, A. S. (2001). *Statistical Structure of Quantum Theory*. Berlin ; New York : Springer. ISBN: 978-3-540-42082-8. URL: http://archive.org/details/springer_10.1007-3-540-44998-1.

Holevo, A. S. and R. F. Werner (Feb. 15, 2001). "Evaluating Capacities of Bosonic Gaussian Channels". In: *Physical Review A* 63.3, p. 032312. DOI: 10.1103/PhysRevA.63.032312.

Horn, Roger A. and Charles R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge: Cambridge University Press. ISBN: 978-0-521-46713-1. DOI: 10.1017/CBO9780511840371.

Jakšić, V., Y. Ogata, C.-A. Pillet, and R. Seiringer (July 2012). "Quantum Hypothesis Testing and Non-Equilibrium Statistical Mechanics". In: *Reviews in Mathematical Physics* 24.06, p. 1230002. DOI: 10.1142/S0129055X12300026.

Jenčová, Anna (Aug. 1, 2018). "Rényi Relative Entropies and Noncommutative L_p-Spaces". In: *Annales Henri Poincaré* 19.8, pp. 2513–2542. DOI: 10.1007/s00023-018-0683-5.

Jenčová, Anna (Oct. 1, 2021). "Rényi Relative Entropies and Noncommutative L_p-Spaces II". In: *Annales Henri Poincaré* 22.10, pp. 3235–3254. DOI: 10.1007/s00023-021-01074-9.

Ji, Kaiyuan, Hemant K. Mishra, Milán Mosonyi, and Mark M. Wilde (July 18, 2024). *Barycentric Bounds on the Error Exponents of Quantum Hypothesis Exclusion*. arXiv: 2407.13728. Pre-published.

Jiang, Haotian, Tarun Kathuria, Yin Tat Lee, Swati Padmanabhan, and Zhao Song (Nov. 2020). "A Faster Interior Point Method for Semidefinite Programming". In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), pp. 910–918. DOI: 10.1109/FOCS46700.2020.00089. arXiv: 2009.10217.

Katariya, Vishal and Mark M. Wilde (Apr. 22, 2020). "Geometric Distinguishability Measures Limit Quantum Channel Estimation and Discrimination". arXiv: 2004.10708.

Katariya, Vishal and Mark M. Wilde (Nov. 8, 2021). "Evaluating the Advantage of Adaptive Strategies for Quantum Channel Distinguishability". In: *Physical Review A* 104.5, p. 052406. DOI: 10.1103/PhysRevA.104.052406. arXiv: 2001.05376.

Khatri, Sumeet and Mark M. Wilde (Nov. 9, 2020). *Principles of Quantum Communication Theory: A Modern Approach*. Version 1. arXiv: 2011.04672v1. Pre-published.

Kholevo, A. S. (Mar. 1979). "On Asymptotically Optimal Hypothesis Testing in Quantum Statistics". In: *Theory of Probability & Its Applications* 23.2, pp. 411–415. DOI: 10.1137/1123048.

Kubo, Fumio and Tsuyoshi Ando (1979). "Means of Positive Linear Operators." In: *Mathematische Annalen* 246, pp. 205–224. URL: https://eudml.org/doc/163339.

Lami, Ludovico (Aug. 12, 2024). *A Solution of the Generalised Quantum Stein's Lemma*. arXiv: 2408.06410. Pre-published.

Lami, Ludovico and Mark M. Wilde (June 2023). "Exact Solution for the Quantum and Private Capacities of Bosonic Dephasing Channels". In: *Nature Photonics* 17.6, pp. 525–530. DOI: 10.1038/s41566-023-01190-4. arXiv: 2205.05736.

Leditzky, Felix, Eneet Kaur, Nilanjana Datta, and Mark M. Wilde (Jan. 25, 2018). "Approaches for Approximate Additivity of the Holevo Information of Quantum Channels". In: *Physical Review A* 97.1, p. 012332. DOI: 10.1103/PhysRevA.97.012332. arXiv: 1709.01111.

Li, Ke (Feb. 1, 2014). "Second-Order Asymptotics for Quantum Hypothesis Testing". In: *The Annals of Statistics* 42.1. DOI: 10.1214/13-AOS1185. arXiv: 1208.1400.

Li, Yonglong, Christoph Hirche, and Marco Tomamichel (June 26, 2022). "Sequential Quantum Channel Discrimination". In: *2022 IEEE International Symposium on Information Theory (ISIT)*. Espoo, Finland: IEEE Press, pp. 270–275. DOI: 10.1109/ISIT50566.2022.9834768.

Li, Yonglong, Vincent Y. F. Tan, and Marco Tomamichel (June 2022). "Optimal Adaptive Strategies for Sequential Quantum Hypothesis Testing". In: *Communications in Mathematical Physics* 392.3, pp. 993–1027. DOI: 10.1007/s00220-022-04362-5. arXiv: 2104.14706.

Lindblad, Göran (June 1, 1975). "Completely Positive Maps and Entropy Inequalities". In: *Communications in Mathematical Physics* 40.2, pp. 147–151. DOI: 10.1007/BF01609396.

Litjens, Bart, Sven Polak, and Alexander Schrijver (July 2017). "Semidefinite Bounds for Nonbinary Codes Based on Quadruples". In: *Designs, Codes and Cryptography* 84.1-2, pp. 87–100. DOI: 10.1007/s10623-016-0216-5. arXiv: 1602.02531.

Martínez Vargas, Esteban, Christoph Hirche, Gael Sentís, Michalis Skotiniotis, Marta Carrizo, Ramon Muñoz-Tapia, and John Calsamiglia (May 6, 2021). "Quantum Sequential Hypothesis Testing". In: *Physical Review Letters* 126.18, p. 180502. DOI: 10.1103/PhysRevLett.126.180502.

Matsumoto, Keiji (2018). "A New Quantum Version of F-Divergence". In: *Reality and Measurement in Algebraic Quantum Theory*. Ed. by Masanao Ozawa, Jeremy Butterfield, Hans Halvorson, Miklós Rédei, Yuichiro Kitajima, and Francesco Buscemi. Springer Proceedings in Mathematics & Statistics. Singapore: Springer, pp. 229–273. ISBN: 9789811324871. DOI: 10.1007/978-981-13-2487-1_10. arXiv: 1311.4722.

McCormick, Katherine C., Jonas Keller, Shaun C. Burd, David J. Wineland, Andrew C. Wilson, and Dietrich Leibfried (Aug. 2019). "Quantum-Enhanced Sensing of a Single-Ion Mechanical Oscillator". In: *Nature* 572.7767, pp. 86–90. DOI: 10.1038/s41586-019-1421-y.

Meyer, V., M. A. Rowe, D. Kielpinski, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland (June 25, 2001). "Experimental Demonstration of Entanglement-Enhanced Rotation Angle Estimation Using Trapped Ions". In: *Physical Review Letters* 86.26, pp. 5870–5873. DOI: 10.1103/PhysRevLett.86.5870.

Michael, Marios H., Matti Silveri, R. T. Brierley, Victor V. Albert, Juha Salmilehto, Liang Jiang, and S. M. Girvin (July 14, 2016). "New Class of Quantum Error-Correcting Codes for a Bosonic Mode". In: *Physical Review X* 6.3, p. 031006. DOI: 10.1103/PhysRevX.6.031006.

Mosonyi, Milán (May 1, 2023). "The Strong Converse Exponent of Discriminating Infinite-Dimensional Quantum States". In: *Communications in Mathematical Physics* 400.1, pp. 83–132. DOI: 10.1007/s00220-022-04598-1. arXiv: 2107.08036.

Mosonyi, Milán and Fumio Hiai (2024). "Some Continuity Properties of Quantum Rényi Divergences". In: *IEEE Transactions on Information Theory* 70.4, pp. 2674–2700. DOI: 10.1109/TIT.2023.3324758. arXiv: 2209.00646.

Mosonyi, Milán and Tomohiro Ogawa (Mar. 1, 2015). "Quantum Hypothesis Testing and the Operational Interpretation of the Quantum Rényi Relative Entropies". In: *Communications in Mathematical Physics* 334.3, pp. 1617–1648. DOI: 10.1007/s00220-014-2248-x. arXiv: 1309.3228.

Mosonyi, Milán, Zsombor Szilágyi, and Mihály Weiner (Feb. 2022). "On the Error Exponents of Binary State Discrimination With Composite Hypotheses". In: *IEEE Transactions on Information Theory* 68.2, pp. 1032–1067. DOI: 10.1109/TIT.2021.3125683. arXiv: 2011.04645.

Müller-Hermes, Alexander and David Reeb (May 1, 2017). "Monotonicity of the Quantum Relative Entropy Under Positive Maps". In: *Annales Henri Poincaré* 18.5, pp. 1777–1788. DOI: 10.1007/s00023-017-0550-9. arXiv: 1512.06117.

Müller-Lennert, Martin, Frédéric Dupuis, Oleg Szehr, Serge Fehr, and Marco Tomamichel (Dec. 1, 2013). "On Quantum Rényi Entropies: A New Generalization and Some Properties". In: *Journal of Mathematical Physics* 54.12, p. 122203. DOI: 10.1063/1.4838856. arXiv: 1306.3142.

Nagaoka, Hiroshi (Nov. 29, 2006). "The Converse Part of The Theorem for Quantum Hoeffding Bound". arXiv: quant-ph/0611289.

Nielsen, Michael A. and Isaac L. Chuang (2010). *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge: Cambridge University Press. 709 pp. ISBN: 978-1-107-00217-3. DOI: 10.1017/CBO9780511976667.

Nussbaum, Michael and Arleta Szkoła (Apr. 1, 2009). "The Chernoff Lower Bound for Symmetric Quantum Hypothesis Testing". In: *The Annals of Statistics* 37.2. DOI: 10.1214/08-AOS593. arXiv: quant-ph/0607216.

Ofek, Nissim, Andrei Petrenko, Reinier Heeres, Philip Reinhold, Zaki Leghtas, Brian Vlastakis, Yehan Liu, Luigi Frunzio, S. M. Girvin, L. Jiang, Mazyar Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (Aug. 2016). "Extending the Lifetime of a Quantum Bit with Error Correction in Superconducting Circuits". In: *Nature* 536.7617, pp. 441–445. DOI: 10.1038/nature18949.

Ogawa, Tomohiro and Hiroshi Nagaoka (Nov. 2000). "Strong Converse and Stein's Lemma in Quantum Hypothesis Testing". In: *IEEE Transactions on Information Theory* 46.7, pp. 2428–2433. DOI: 10.1109/18.887855. arXiv: quant-ph/9906090.

Oreshkov, Ognyan (Dec. 2, 2019). "Time-Delocalized Quantum Subsystems and Operations: On the Existence of Processes with Indefinite Causal Structure in Quantum Mechanics". In: *Quantum* 3, p. 206. DOI: 10.22331/q-2019-12-02-206.

Pedersen, Gert K. (1989). *Analysis Now*. Vol. 118. Graduate Texts in Mathematics. New York, NY: Springer. ISBN: 978-1-4612-1007-8. DOI: 10.1007/978-1-4612-1007-8.

Petz, Dénes (Feb. 1, 1986). "Quasi-Entropies for Finite Quantum Systems". In: *Reports on Mathematical Physics* 23.1, pp. 57–65. DOI: 10.1016/0034-4877(86)90067-4.

Petz, Dénes and Mary Beth Ruskai (Jan. 1, 1998). "Contraction of Generalized Relative Entropy Under Stochastic Mappings on Matrices". In: *Infinite Dimensional Analysis, Quantum Probability and Related Topics* 01.01, pp. 83–89. DOI: 10.1142/S0219025798000077.

Pirandola, Stefano, Riccardo Laurenza, Carlo Ottaviani, and Leonardo Banchi (Apr. 26, 2017). "Fundamental Limits of Repeaterless Quantum Communications". In: *Nature Communications* 8.1, p. 15043. DOI: 10.1038/ncomms15043.

Polak, Sven (Sept. 10, 2019). "New Methods in Coding Theory: Error-Correcting Codes and the Shannon Capacity". PhD thesis. University of Amsterdam. arXiv: 2005.02945.

Rastegin, A. E. (Oct. 10, 2002). "Relative Error of State-Dependent Cloning". In: *Physical Review A* 66.4, p. 042304. DOI: 10.1103/PhysRevA.66.042304.

Rastegin, A. E. (Oct. 2003). "A Lower Bound on the Relative Error of Mixed-State Cloning and Related Operations". In: *Journal of Optics B: Quantum and Semiclassical Optics* 5.6, S647–S650. DOI: 10.1088/1464-4266/5/6/017.

Rastegin, A. E. (Feb. 14, 2006). *Sine Distance for Quantum States*. arXiv: quant-ph/0602112. Pre-published.

Reed, Michael and Barry Simon (1980). *Methods of Modern Mathematical Physics: Functional Analysis*. Gulf Professional Publishing. 417 pp. ISBN: 978-0-12-585050-6. Google Books: bvuRuwuFBWwC.

Regula, Bartosz, Ludovico Lami, and Nilanjana Datta (Jan. 21, 2025). *Tight Relations and Equivalences between Smooth Relative Entropies*. arXiv: 2501.12447. Pre-published.

Renner, Renato (2005). "Security of Quantum Key Distribution". Doctoral Thesis. ETH Zurich. arXiv: quant-ph/0512258.

Rényi, Alfréd (Jan. 1, 1961). "On Measures of Entropy and Information". In: *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pp. 547–561. URL: https://projecteuclid.org/ebooks/berkeley-symposium-on-mathematical-statistics-and-probability/Proceedings-of-the-Fourth-Berkeley-Symposium-on-Mathematical-Statistics-and/chapter/On-Measures-of-Entropy-and-Information/bsmsp/1200512181.

Rosati, Matteo, Andrea Mari, and Vittorio Giovannetti (Oct. 18, 2018). "Narrow Bounds for the Quantum Capacity of Thermal Attenuators". In: *Nature Communications* 9.1, p. 4339. DOI: 10.1038/s41467-018-06848-0.

Salek, Farzin, Masahito Hayashi, and Andreas Winter (Feb. 14, 2022). "Usefulness of Adaptive Strategies in Asymptotic Quantum Channel Discrimination". In: *Physical Review A* 105.2, p. 022419. DOI: 10.1103/PhysRevA.105.022419. arXiv: 2011.06569.

Shannon, C. E. (July 1948). "A Mathematical Theory of Communication". In: *The Bell System Technical Journal* 27.3, pp. 379–423. DOI: 10.1002/j.1538-7305.1948.tb01338.x.

Takeoka, Masahiro, Saikat Guha, and Mark M. Wilde (Oct. 24, 2014). "Fundamental Rate-Loss Tradeoff for Optical Quantum Key Distribution". In: *Nature Communications* 5.1, p. 5235. DOI: 10.1038/ncomms6235.

Tomamichel, Marco (2012). "A Framework for Non-Asymptotic Quantum Information Theory". Doctoral Thesis. ETH Zurich. arXiv: 1203.2142.

Tomamichel, Marco (2016). *Quantum Information Processing with Finite Resources*. Vol. 5. SpringerBriefs in Mathematical Physics. Cham: Springer International Publishing. ISBN: 978-3-319-21891-5. DOI: 10.1007/978-3-319-21891-5.

Tomamichel, Marco, Roger Colbeck, and Renato Renner (Dec. 2009). "A Fully Quantum Asymptotic Equipartition Property". In: *IEEE Transactions on Information Theory* 55.12, pp. 5840–5847. DOI: 10.1109/TIT.2009.2032797. arXiv: 0811.1221.

Tomamichel, Marco and Masahito Hayashi (Nov. 2013). "A Hierarchy of Information Quantities for Finite Block Length Analysis of Quantum Tasks". In: *IEEE Transactions on Information Theory* 59.11, pp. 7693–7710. DOI: 10.1109/TIT.2013.2276628.

Uhlmann, Armin (Apr. 1, 1976). "The "Transition Probability" in the State Space of a $*$-Algebra". In: *Reports on Mathematical Physics* 9.2, pp. 273–279. DOI: 10.1016/0034-4877(76)90060-4.

Umegaki, Hisaharu (Jan. 1962). "Conditional Expectation in an Operator Algebra. IV. Entropy and Information". In: *Kodai Mathematical Seminar Reports* 14.2, pp. 59–85. DOI: 10.2996/kmj/1138844604.

Vandenberghe, Lieven and Stephen Boyd (Mar. 1996). "Semidefinite Programming". In: *SIAM Review* 38.1, pp. 49–95. DOI: 10.1137/1038003.

Von Neumann, John (1932). *Mathematische grundlagen der quantenmechanik*. Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen mit besonderer Berücksichtigung der Anwendungsgebiete ; Bd. 38. Berlin: Springer, J. Springer. 262 pp.

Wang, Jie (July 18, 2024). *A More Efficient Reformulation of Complex SDP as Real SDP*. arXiv: 2307.11599. Pre-published.

Wang, Ligong and Renato Renner (May 15, 2012). "One-Shot Classical-Quantum Capacity and Hypothesis Testing". In: *Physical Review Letters* 108.20, p. 200501. DOI: 10.1103/PhysRevLett.108.200501. arXiv: 1007.5456.

Wang, Xin and Mark M. Wilde (Dec. 11, 2019a). "Resource Theory of Asymmetric Distinguishability". In: *Physical Review Research* 1.3, p. 033170. DOI: 10.1103/PhysRevResearch.1.033170. arXiv: 1905.11629.

Wang, Xin and Mark M. Wilde (Dec. 11, 2019b). "Resource Theory of Asymmetric Distinguishability for Quantum Channels". In: *Physical Review Research* 1.3, p. 033169. DOI: 10.1103/PhysRevResearch.1.033169. arXiv: 1907.06306.

Watanabe, Shun and Masahito Hayashi (Apr. 2017). "Finite-Length Analysis on Tail Probability for Markov Chain and Application to Simple Hypothesis Testing". In: *The Annals of Applied Probability* 27.2, pp. 811–845. DOI: 10.1214/16-AAP1216. arXiv: 1401.3801.

Watrous, John (Apr. 26, 2018). *The Theory of Quantum Information*. 1st ed. Cambridge: Cambridge University Press. ISBN: 978-1-107-18056-7. DOI: 10.1017/9781316848142.

Wechs, Julian, Hippolyte Dourdent, Alastair A. Abbott, and Cyril Branciard (Aug. 26, 2021). "Quantum Circuits with Classical versus Quantum Control of Causal Order". In: *PRX Quantum* 2.3, p. 030335. DOI: 10.1103/PRXQuantum.2.030335. arXiv: 2101.08796.

Wilde, Mark M., Mario Berta, Christoph Hirche, and Eneet Kaur (Aug. 2020). "Amortized Channel Divergence for Asymptotic Quantum Channel Discrimination". In: *Letters in Mathematical Physics* 110.8, pp. 2277–2336. DOI: 10.1007/s11005-020-01297-7. arXiv: 1808.01498.

Wilde, Mark M., Marco Tomamichel, and Mario Berta (Mar. 2017). "Converse Bounds for Private Communication Over Quantum Channels". In: *IEEE Transactions on Information Theory* 63.3, pp. 1792–1817. DOI: 10.1109/TIT.2017.2648825. arXiv: 1602.08898.

Wilde, Mark M., Andreas Winter, and Dong Yang (Oct. 1, 2014). "Strong Converse for the Classical Capacity of Entanglement-Breaking and Hadamard Channels via a Sandwiched Rényi Relative Entropy". In: *Communications in Mathematical Physics* 331.2, pp. 593–622. DOI: 10.1007/s00220-014-2122-x.

Winter, Andreas (1999). "Coding Theorems of Quantum Information Theory". PhD thesis. Bielefeld: Bielefeld University. arXiv: quant-ph/9907077.

Wolf, Michael M., David Pérez-García, and Geza Giedke (Mar. 26, 2007). "Quantum Capacities of Bosonic Channels". In: *Physical Review Letters* 98.13, p. 130501. DOI: 10.1103/PhysRevLett.98.130501.

Zhang, Junhua, Mark Um, Dingshun Lv, Jing-Ning Zhang, Lu-Ming Duan, and Kihwan Kim (Oct. 18, 2018). "NOON States of Nine Quantized Vibrations in Two Radial Modes of a Trapped Ion". In: *Physical Review Letters* 121.16, p. 160502. DOI: 10.1103/PhysRevLett.121.160502.