



HAL
open science

An invitation to the sample complexity of quantum hypothesis testing

Hao-Chung Cheng, Nilanjana Datta, Nana Liu, Theshani Nuradha, Robert Salzmann, Mark M Wilde

► To cite this version:

Hao-Chung Cheng, Nilanjana Datta, Nana Liu, Theshani Nuradha, Robert Salzmann, et al.. An invitation to the sample complexity of quantum hypothesis testing. 2024. hal-04840569

HAL Id: hal-04840569

<https://hal.science/hal-04840569v1>

Preprint submitted on 16 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

An invitation to the sample complexity of quantum hypothesis testing

Hao-Chung Cheng,^{1,2,3} Nilanjana Datta,⁴ Nana Liu,^{5,6,7,*}
Theshani Nuradha,⁸ Robert Salzmann,^{4,9} and Mark M. Wilde⁸

¹*Department of Electrical Engineering,
National Taiwan University, Taipei 106, Taiwan (R.O.C.)*

²*Hon Hai (Foxconn) Quantum Computing Center,
New Taipei City 236, Taiwan (R.O.C.)*

³*Physics Division, National Center for Theoretical Sciences, Taipei 106, Taiwan (R.O.C.)*

⁴*Department of Applied Mathematics and Theoretical Physics,
University of Cambridge, Wilberforce Road,
Cambridge, CB3 0WA, United Kingdom*

⁵*Institute of Natural Sciences, School of Mathematical Sciences, MOE-LSC,
Shanghai Jiao Tong University, Shanghai, 200240, P. R. China*

⁶*Shanghai Artificial Intelligence Laboratory, Shanghai, China*

⁷*University of Michigan-Shanghai Jiao Tong University Joint Institute, Shanghai 200240, China*

⁸*School of Electrical and Computer Engineering,
Cornell University, Ithaca, New York 14850, USA*

⁹*Université de Lyon, Inria, ENS Lyon,
UCBL, LIP, F-69342, Lyon Cedex 07, France*

(Dated: May 17, 2024)

Quantum hypothesis testing has been traditionally studied from the information-theoretic perspective, wherein one is interested in the optimal decay rate of error probabilities as a function of the number of samples of an unknown state. In this paper, we study the sample complexity of quantum hypothesis testing, wherein the goal is to determine the minimum number of samples needed to reach a desired error probability. By making use of the wealth of knowledge that already exists in the literature on quantum hypothesis testing, we characterize the sample complexity of binary quantum hypothesis testing in the symmetric and asymmetric settings, and we provide bounds on the sample complexity of multiple quantum hypothesis testing. In more detail, we prove that the sample complexity of symmetric binary quantum hypothesis testing depends logarithmically on the inverse error probability and inversely on the negative logarithm of the fidelity. As a counterpart of the quantum Stein's lemma, we also find that the sample complexity of asymmetric binary quantum hypothesis testing depends logarithmically on the inverse type II error probability and inversely on the quantum relative entropy, provided that the type II error probability is sufficiently small. We then provide lower and upper bounds on the sample complexity of multiple quantum hypothesis testing, with it remaining an intriguing open question to improve these bounds. The final part of our paper outlines and reviews how sample complexity of quantum hypothesis testing is relevant to a broad swathe of research areas and can enhance understanding of many fundamental concepts, including quantum algorithms for simulation and search, quantum learning and classification, and foundations of quantum mechanics. As such, we view our paper as an invitation to researchers coming from different communities to study

and contribute to the problem of sample complexity of quantum hypothesis testing, and we outline a number of open directions for future research.

CONTENTS

I. Introduction	3
II. Background	4
A. Quantum information-theoretic quantities	5
B. Quantum hypothesis testing	7
1. Symmetric binary quantum hypothesis testing	7
2. Asymmetric binary quantum hypothesis testing	8
3. Multiple quantum hypothesis testing	8
4. Efficient algorithm for computing optimal error probabilities in quantum hypothesis testing	9
III. Statement of the problem	9
IV. Sample complexity results	11
A. Symmetric binary quantum hypothesis testing	11
1. Two pure states	11
2. Two general states	13
B. Asymmetric binary quantum hypothesis testing	14
C. Multiple quantum hypothesis testing	16
V. Applications and further avenues of research	17
A. Optimality of quantum algorithms	18
1. Quantum simulation of linear ordinary and partial differential equations	18
2. Quantum simulation of nonlinear ordinary and partial differential equations	20
3. Unstructured search and related applications	21
4. Hidden subgroup problem	22
B. Quantum learning and classification	22
C. Foundational quantum information and quantum mechanics	26
VI. Summary and future directions	27
Acknowledgments	27
References	28
A. Properties of and relations between quantum divergences	34
B. Efficient representation of optimal measurements and efficient calculation of optimal error probabilities in quantum hypothesis testing	36

* nana.liu@quantumlah.org

C. Proof of Equations (31)–(32)	39
D. Proof of Remark 2	41
E. Proof of Theorem 6	42
F. Proof of Theorem 7	43
G. Fuchs–Caves measurement strategy for achieving the asymptotic sample complexity of symmetric binary quantum hypothesis testing	45
H. Proof of Theorem 9	47
I. Proof of Corollary 10	49
J. Proof of Theorem 11	56
K. Proof of Remark 4	58

I. INTRODUCTION

Distinguishing between various possibilities is fundamental to the scientific method and the process of discovery and categorization. As such, mathematical methods underlying distinguishability have been studied for a long time [1] and are now highly developed and fall under the general framework of hypothesis testing [2]. There are various ways of formulating the hypothesis testing problem, and due to the core role of distinguishability in a variety of fields, it has also been influential in many domains beyond mathematical statistics.

Here we are interested in quantum hypothesis testing, which has its origins in early work in the field of quantum information theory [3–5]. Quantum hypothesis testing is an important foundational topic in the sense that it merges quantum mechanics and mathematical statistics and is applicable to basic quantum physics experiments. Moreover, in addition to applications in communication, information processing, and computation [6], it has even found use in understanding foundational aspects of quantum mechanics [7]. In the most basic version of the problem, a quantum system is prepared in one of two possible states (density operators), denoted by ρ and σ , and it is the goal of the distinguisher, who does not know *a priori* which state was prepared, to determine the identity of the unknown state. The distinguisher ideally wishes to minimize the probability of making an incorrect decision.

The task of quantum hypothesis testing becomes more interesting when multiple copies of the unknown state are provided to the distinguisher. In the case that $n \in \mathbb{N}$ copies (or samples) are provided, the states are then denoted by $\rho^{\otimes n}$ and $\sigma^{\otimes n}$ (i.e., tensor-power states). Intuitively, extra samples are helpful for decreasing the error probability when trying to determine the unknown state. In fact, it is well known in the information theory literature that the error probability decreases exponentially fast in the number of samples provided that $\rho \neq \sigma$ [8–13].

Quantum hypothesis testing can be generalized to the case in which there are multiple hypotheses [14]. In this case, a value x is selected from a finite alphabet (say, \mathcal{X}) with a prior probability distribution $\{p(x)\}_{x \in \mathcal{X}}$. and a state $\rho_x^{\otimes n}$ is prepared, where $n \in \mathbb{N}$. It is then the goal of the distinguisher to determine the value of x with as small an error

probability as possible. It is also known in this case that the error probability generally decays exponentially fast with n [15]. Alternatively, one can turn this problem around and fix a constant upper bound on the desired error probability and demand that the distinguisher determines the value of x with as few samples as possible (minimize n) while meeting the error probability constraint.

On the one hand, it is a common goal in the information theory literature to determine the optimal exponential rate of decay of the error probability, known as an error exponent. This is typically studied in the large n limit, but more recently researchers have sought out more refined statements [16–18]. Furthermore, this line of study has direct links with communication theory and allows for making statements about error exponents when communicating messages over a quantum channel [13, 19–25].

On the other hand, it is more common in the algorithms and machine learning literature to consider a notion called sample complexity [26, 27]; that is, if we fix the error probability to be no larger than a constant $\varepsilon \in [0, 1]$, what is the smallest value of n (minimum number of samples) needed to ensure that the distinguisher can figure out the unknown state with an error probability no larger than ε ? This latter notion is useful because it indicates how long one needs to wait in order to make a decision with a desired performance, and it is more compatible with the runtime of a probabilistic or quantum algorithm that solves a problem of interest.

In this paper, we provide a systematic study of the sample complexity of quantum hypothesis testing, when there are either two or more hypotheses to distinguish. More broadly, we consider our paper to be an invitation to the wider community to study more general questions regarding the sample complexity of quantum hypothesis testing, which touches the interface between statistics, quantum information, quantum algorithms, and learning theory. We divide our paper into three main parts:

1. The first part is mainly pedagogical, which we have aimed to make broadly accessible. We provide some background on quantum information-theoretic quantities of interest used throughout our paper, and we also review symmetric binary, asymmetric binary, and multiple quantum hypothesis testing. Here we also include a new finding that there is an efficient, polynomial-time algorithm for computing the optimal error probabilities in these three different settings.
2. The second part contains our main technical results. Here we precisely define various instances of the problem of sample complexity for quantum hypothesis testing, and we prove bounds for these various sample complexities. We also discuss distinctions between sample complexity in the classical and quantum cases.
3. The last part includes a selection of applications, which serve only as a springboard for open questions that the wider community can begin to study in more depth.

II. BACKGROUND

In this section, we establish some notation and recall various quantities of interest used throughout the rest of our paper.

Let $\mathbb{N} := \{1, 2, \dots\}$. Throughout our paper, we let ρ and σ denote quantum states, which are positive semi-definite operators acting on a separable Hilbert space and with trace equal to one. Note that separable Hilbert spaces are those that are infinite-dimensional and are

spanned by a countable orthonormal basis [28, Proposition 1.12]. Let I denote the identity operator. For every bounded operator A and $p \geq 1$, we define the Schatten p -norm as

$$\|A\|_p := \left(\text{Tr} \left[(A^\dagger A)^{p/2} \right] \right)^{1/p}. \quad (1)$$

Due to the variational characterization of the trace norm,

$$\|A\|_1 = \max_U \text{Re}[\text{Tr}[AU]], \quad (2)$$

where the optimization is over every unitary U .

A. Quantum information-theoretic quantities

Definition 1 (Fidelities). *Let ρ and σ be quantum states.*

1) *The quantum fidelity is defined as [29]*

$$F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1^2. \quad (3)$$

Note that $F(\rho, \sigma) \in [0, 1]$, it is equal to one if and only if $\rho = \sigma$, and it is equal to zero if and only if ρ is orthogonal to σ (i.e., $\rho\sigma = 0$).

2) *The Holevo fidelity is defined as [30]*

$$F_H(\rho, \sigma) := \left(\text{Tr}[\sqrt{\rho}\sqrt{\sigma}] \right)^2. \quad (4)$$

Note that $F_H(\rho, \sigma) \in [0, 1]$, it is equal to one if and only if $\rho = \sigma$, and it is equal to zero if and only if ρ is orthogonal to σ (i.e., $\rho\sigma = 0$).

A divergence $\mathbf{D}(\rho\|\sigma)$ is a function of two quantum states ρ and σ , and we say that it obeys the data-processing inequality if the following holds for all states ρ and σ and every channel \mathcal{N} :

$$\mathbf{D}(\rho\|\sigma) \geq \mathbf{D}(\mathcal{N}(\rho)\|\mathcal{N}(\sigma)). \quad (5)$$

Definition 2 (Distances and divergences). *Let ρ and σ be quantum states.*

1) *The normalized trace distance is defined as*

$$\frac{1}{2} \|\rho - \sigma\|_1. \quad (6)$$

2) *The Bures distance is defined as [31, 32]*

$$d_B(\rho, \sigma) := \min_U \|\sqrt{\rho} - U\sqrt{\sigma}\|_2 \quad (7)$$

$$= \min_U \sqrt{2 \left(1 - \text{Re}[\text{Tr}[\sqrt{\rho}\sqrt{\sigma}U]] \right)} \quad (8)$$

$$\stackrel{(2)}{=} \sqrt{2 \left(1 - \sqrt{F}(\rho, \sigma) \right)}, \quad (9)$$

where the optimization is over every unitary U .

3) The quantum Hellinger distance is defined as [33, 34]

$$d_{\text{H}}(\rho, \sigma) := \|\sqrt{\rho} - \sqrt{\sigma}\|_2 = \sqrt{2 \left(1 - \sqrt{F_{\text{H}}}(\rho, \sigma)\right)}. \quad (10)$$

4) The Petz–Rényi divergence of order $\alpha \in (0, 1) \cup (1, \infty)$ is defined as [35]

$$D_{\alpha}(\rho\|\sigma) := \frac{1}{\alpha - 1} \ln Q_{\alpha}(\rho\|\sigma) \quad (11)$$

$$\text{where } Q_{\alpha}(A\|B) := \lim_{\varepsilon \rightarrow 0^+} \text{Tr} \left[A^{\alpha} (B + \varepsilon I)^{1-\alpha} \right], \quad \forall A, B \geq 0. \quad (12)$$

It obeys the data-processing inequality for all $\alpha \in (0, 1) \cup (1, 2]$ [35]. Note also that $D_{1/2}(\rho\|\sigma) = -\ln F_{\text{H}}(\rho, \sigma)$.

5) The sandwiched Rényi divergence of order $\alpha \in (0, 1) \cup (1, \infty)$ is defined as [36, 37]

$$\tilde{D}_{\alpha}(\rho\|\sigma) := \frac{1}{\alpha - 1} \ln \tilde{Q}_{\alpha}(\rho\|\sigma) \quad (13)$$

$$\text{where } \tilde{Q}_{\alpha}(A\|B) := \lim_{\varepsilon \rightarrow 0^+} \text{Tr} \left[\left(A^{\frac{1}{2}} (B + \varepsilon I)^{\frac{1-\alpha}{\alpha}} A^{\frac{1}{2}} \right)^{\alpha} \right], \quad \forall A, B \geq 0. \quad (14)$$

It obeys the data-processing inequality for all $\alpha \in [1/2, 1) \cup (1, \infty)$ [38, 39]. Note also that $\tilde{D}_{1/2}(\rho\|\sigma) = -\ln F(\rho, \sigma)$.

6) The quantum relative entropy is defined as [40]

$$D(\rho\|\sigma) := \lim_{\varepsilon \rightarrow 0^+} \text{Tr} [\rho (\ln \rho - \ln(\sigma + \varepsilon I))], \quad (15)$$

and it obeys the data-processing inequality [41].

7) The Chernoff divergence is defined as [10]

$$C(\rho\|\sigma) := -\ln Q_{\min}(\rho\|\sigma), \quad (16)$$

$$\text{where } Q_{\min}(A\|B) := \min_{s \in [0, 1]} Q_s(A\|B), \quad \forall A, B \geq 0. \quad (17)$$

If the quantum states ρ and σ commute, i.e., $\rho\sigma = \sigma\rho$, then they share a common eigenbasis. Let $\{P(x)\}_x$ and $\{Q(x)\}_x$ be the sets of eigenvalues of ρ and σ , respectively. This situation is known as the *classical scenario* because the above-defined divergences reduce to their corresponding classical divergences between probability mass functions P and Q . For instance, in this case, the normalized trace distance is equal to the classical total variation distance, i.e., $\frac{1}{2} \|\rho - \sigma\|_1 = \frac{1}{2} \sum_x |P(x) - Q(x)|$. Furthermore, both the quantum fidelity and Holevo fidelity reduce to the squared Bhattacharyya coefficient, i.e., $F(\rho, \sigma) = F_{\text{H}}(\rho, \sigma) = \left(\sum_x \sqrt{P(x)Q(x)} \right)^2$. Likewise, the Bures distance and quantum Hellinger distance correspond to the classical Hellinger distance (up to a factor $\sqrt{2}$). Both the Petz–Rényi and sandwiched Rényi divergences reduce to the classical Rényi divergence [42], and the quantum relative entropy reduces to the Kullback–Leibler divergence [43].

In the general case in which ρ and σ do not commute, $F(\rho, \sigma) \neq F_{\text{H}}(\rho, \sigma)$ and $D_{\alpha}(\rho\|\sigma) \neq \tilde{D}_{\alpha}(\rho\|\sigma)$ for all $\alpha \in (0, 1) \cup (1, \infty)$ if they are finite [44, Theorem 2.1]. This explains why there are more divergences in the quantum scenario than there are in the classical scenario. We refer readers to Lemma 18 in Appendix A for the relations between those quantities.

B. Quantum hypothesis testing

Let us first review binary classical hypothesis testing, which is a special case of quantum hypothesis testing. Suppose a classical system is modeled by a random variable Y , which is distributed according to a distribution P under the null hypothesis and according to a distribution Q under the alternative hypothesis. The goal of a distinguisher is to guess which hypothesis is correct. To do so, the distinguisher may take n independent and identical samples of Y and apply a (possibly randomized) test, which is mathematically described by a conditional distribution $P_{0|Y^n}^{(n)}$ and $P_{1|Y^n}^{(n)} = 1 - P_{0|Y^n}^{(n)}$, where 0 and 1 correspond to guessing the null hypothesis P and alternative hypothesis Q , respectively. There are two kinds of errors that can occur, and performance metrics can be constructed based on the probabilities of these errors. The probability of guessing Q under null hypothesis P is then given by the expectation $\mathbb{E}_{Y^n \sim P^{\otimes n}} \left[P_{1|Y^n}^{(n)} \right]$, and the probability of guessing P under the alternative hypothesis Q is given by $\mathbb{E}_{Y^n \sim Q^{\otimes n}} \left[P_{0|Y^n}^{(n)} \right]$.

Now let us move on to the scenario of quantum hypothesis testing, in which the underlying system is quantum mechanical. Instead of being modeled by a probability distribution, a quantum system is prepared in a quantum state, which is modeled by a density operator. In the case of two hypotheses (i.e., two states ρ and σ), the system is either prepared in the state $\rho^{\otimes n}$ or $\sigma^{\otimes n}$, where $n \in \mathbb{N}$. As in the classical scenario, the goal of a distinguisher is to guess which state was prepared, doing so by means of a quantum measurement with two outcomes. Mathematically, such a measurement is described by two operators $\Lambda_\rho^{(n)}$ and $\Lambda_\sigma^{(n)}$ satisfying $\Lambda_\rho^{(n)}, \Lambda_\sigma^{(n)} \geq 0$ and $\Lambda_\rho^{(n)} + \Lambda_\sigma^{(n)} = I^{\otimes n}$, where the outcome $\Lambda_\rho^{(n)}$ is associated with guessing ρ and $\Lambda_\sigma^{(n)}$ is associated with guessing σ . The probability of guessing σ when the prepared state is ρ is equal to $\text{Tr} \left[\Lambda_\sigma^{(n)} \rho^{\otimes n} \right]$, and the probability of guessing ρ when the prepared state is σ is equal to $\text{Tr} \left[\Lambda_\rho^{(n)} \sigma^{\otimes n} \right]$. If ρ and σ commute, then without loss of generality the distinguisher can choose a measurement $(\Lambda_\rho^{(n)}, \Lambda_\sigma^{(n)})$ sharing a common eigenbasis with ρ and σ [45, Remark 3]. In such a case, quantum hypothesis testing reduces to classical hypothesis testing, comparing the distribution resulting from the eigenvalues of ρ against that resulting from the eigenvalues of σ .

There are two scenarios of interest here, called the symmetric and asymmetric settings, reviewed in more detail in Sections II B 1 and II B 2, respectively. We then move on to reviewing multiple hypothesis testing in Section II B 3.

1. Symmetric binary quantum hypothesis testing

In the setting of symmetric binary quantum hypothesis testing, we suppose that there is a prior probability $p \in (0, 1)$ associated with preparing the state ρ , and there is a prior probability $q \equiv 1 - p$ associated with preparing the state σ . Indeed, the unknown state is prepared by flipping a coin, with the probability of heads being p and the probability of tails being q . If the outcome of the coin flip is heads, then n quantum systems are prepared in the state $\rho^{\otimes n}$, and if the outcome is tails, then the n quantum systems are prepared in the state $\sigma^{\otimes n}$. Thus, the expected error probability in this experiment is as follows:

$$p_{e, \Lambda_\rho^{(n)}}(p, \rho, q, \sigma, n) := p \text{Tr} \left[\Lambda_\sigma^{(n)} \rho^{\otimes n} \right] + q \text{Tr} \left[\Lambda_\rho^{(n)} \sigma^{\otimes n} \right] \quad (18)$$

$$= p \operatorname{Tr}[(I^{\otimes n} - \Lambda_\rho^{(n)})\rho^{\otimes n}] + q \operatorname{Tr}[\Lambda_\rho^{(n)}\sigma^{\otimes n}], \quad (19)$$

where the second equality follows because $\Lambda_\sigma^{(n)} = I^{\otimes n} - \Lambda_\rho^{(n)}$ and furthermore indicates that the error probability can be written as a function of only the first measurement operator $\Lambda_\rho^{(n)}$.

Given p and descriptions of the states ρ and σ , the distinguisher can minimize the error-probability expression in (19) over all possible measurements. The Helstrom–Holevo theorem [4, 5] states that the optimal error probability $p_e(p, \rho, q, \sigma, n)$ of hypothesis testing is as follows:

$$p_e(p, \rho, q, \sigma, n) := \inf_{\Lambda^{(n)}} p_{e, \Lambda^{(n)}}(p, \rho, q, \sigma, n) \quad (20)$$

$$= \inf_{\Lambda^{(n)}} \{p \operatorname{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] + q \operatorname{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] : 0 \leq \Lambda^{(n)} \leq I^{\otimes n}\} \quad (21)$$

$$= \frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1). \quad (22)$$

2. Asymmetric binary quantum hypothesis testing

In the setting of asymmetric binary quantum hypothesis testing, there are no prior probabilities associated with preparing an unknown state—we simply assume that a state is prepared deterministically, but the identity of the prepared state is unknown to the distinguisher. The goal of the distinguisher is to minimize the probability of the second kind of error subject to a constraint on the probability of the first kind of error. Indeed, given a fixed $\varepsilon \in [0, 1]$, the scenario reduces to the following optimization problem:

$$\beta_\varepsilon(\rho^{\otimes n} \parallel \sigma^{\otimes n}) := \inf_{\Lambda^{(n)}} \left\{ \operatorname{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] : \operatorname{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] \leq \varepsilon, \right. \\ \left. 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \right\}. \quad (23)$$

This formulation of the hypothesis testing problem is more relevant in radar applications [46], where the type I error probability is called the “false alarm rate,” and the type II error probability is called the “missed detection rate.” Indeed, in such a scenario, one is willing to tolerate a fixed rate of false alarms but then wishes to minimize the rate of missed detections.

3. Multiple quantum hypothesis testing

We now review multiple quantum hypothesis testing, also known as M -ary quantum hypothesis testing because the goal is to select one among M possible hypotheses. Let $\mathcal{S} := \{(p_m, \rho_m)\}_{m=1}^M$ be an ensemble of M states with prior probabilities taking values in the set $\{p_m\}_{m=1}^M$. Without loss of generality, let us assume that $p_m > 0$ for all $m \in \{1, 2, \dots, M\}$. The minimum error probability of M -ary hypothesis testing, given n copies of the unknown state, is as follows:

$$p_e(\mathcal{S}, n) := \inf_{(\Lambda_1^{(n)}, \dots, \Lambda_M^{(n)})} \sum_{m=1}^M p_m \operatorname{Tr}[(I^{\otimes n} - \Lambda_m^{(n)})\rho_m^{\otimes n}], \quad (24)$$

where the minimization is over every positive operator-valued measure (POVM) (i.e., a tuple $(\Lambda_1^{(n)}, \dots, \Lambda_M^{(n)})$ satisfying $\Lambda_i^{(n)} \geq 0$ for all $i \in \{1, \dots, M\}$ and $\sum_{m=1}^M \Lambda_m^{(n)} = I^{\otimes n}$).

4. *Efficient algorithm for computing optimal error probabilities in quantum hypothesis testing*

In Sections II B 1, II B 2, and II B 3, we outlined several quantum hypothesis testing tasks, including symmetric binary, asymmetric binary, and multiple hypotheses. Along with these tasks are the associated optimal error probabilities in (20), (23), and (24), which involve an optimization over all possible measurements acting on n quantum systems.

Let us suppose that each quantum system has a fixed dimension d . In this case, all of these optimization problems can be cast as semi-definite programs (SDPs) [47], meaning that standard semi-definite programming solvers can be used to calculate the optimal values. However, the matrices involved in these optimization problems are of dimension $d^n \times d^n$, and thus a naive approach to solving these SDPs requires time exponential in the number n of systems. As such, it might seem as if calculating these optimal values is an intractable task.

The naive approach mentioned above neglects the fact that the optimization problems in (20), (23), and (24) possess a large amount of symmetry, due to the fact that the inputs to these problems are tensor-power states and are thus invariant under permutations of the systems. By exploiting this permutation symmetry, we show in Appendix B that the SDPs needed to compute the optimal values can be reduced to SDPs of size polynomial in n , where the polynomial degree is a function of d . Thus, based on standard results on the efficiency of SDP solvers [48–51], there is thus a polynomial-time algorithm for computing the optimal error probabilities in quantum hypothesis testing. This observation closely follows other recent advances in quantum information, having to do with computing bounds on channel capacities [52] and with computing the optimal error probability in asymmetric channel discrimination [53].

III. STATEMENT OF THE PROBLEM

When reviewing the various hypothesis testing problems in Section II B, we see that the goal is to minimize the error probability for a fixed choice of $n \in \mathbb{N}$. As stated in the introduction (Section I), the optimization task for sample complexity flips this reasoning around: indeed, the goal is to determine the minimum value of $n \in \mathbb{N}$ (i.e., minimum number of samples) needed to meet a fixed error probability constraint. As discussed in the introduction of our paper, this formulation of the hypothesis testing problem is more consistent with the notion of the runtime of a probabilistic or quantum algorithm, for which one typically finds that the runtime depends on the desired error. Indeed, if there are procedures for preparing the various states involved in a given hypothesis testing problem, with fixed runtimes, then the sample complexity corresponds to the total amount of time one must wait to prepare n samples to achieve a desired error probability. It should certainly be mentioned that sample complexity ignores the runtime of a quantum measurement procedure that actually discriminates the states, so that this notion can be understood as straddling the boundary between information theory and complexity theory.

More formally, we state the definitions of the sample complexity of symmetric binary, asymmetric binary, and multiple quantum hypothesis testing in the following Definitions 3, 4, and 5, respectively. In each case, the simplest way to state the definitions is to employ the various error-probability metrics in (20), (23), and (24) and define the sample complexity to be the minimum value of $n \in \mathbb{N}$ needed to get the error probability metric below a threshold $\varepsilon \in [0, 1]$.

Definition 3 (Sample complexity of symmetric binary hypothesis testing). *Let $p \in (0, 1)$, $q = 1 - p$, and $\varepsilon \in [0, 1]$, and let ρ and σ be states. The sample complexity $n^*(p, \rho, q, \sigma, \varepsilon)$ of symmetric binary quantum hypothesis testing is defined as follows:*

$$n^*(p, \rho, q, \sigma, \varepsilon) := \inf \{n \in \mathbb{N} : p_\varepsilon(p, \rho, q, \sigma, n) \leq \varepsilon\}. \quad (25)$$

Definition 4 (Sample complexity of asymmetric binary hypothesis testing). *Let $\varepsilon, \delta \in [0, 1]$, and let ρ and σ be states. The sample complexity $n^*(\rho, \sigma, \varepsilon, \delta)$ of asymmetric binary quantum hypothesis testing is defined as follows:*

$$n^*(\rho, \sigma, \varepsilon, \delta) := \inf \{n \in \mathbb{N} : \beta_\varepsilon(\rho^{\otimes n} \| \sigma^{\otimes n}) \leq \delta\}. \quad (26)$$

Definition 5 (Sample complexity of M -ary hypothesis testing). *Let $\varepsilon \in [0, 1]$, and let $\mathcal{S} := \{(p_m, \rho_m)\}_{m=1}^M$ be an ensemble of M states. The sample complexity $n^*(\mathcal{S}, \varepsilon)$ of M -ary quantum hypothesis testing is defined as follows:*

$$n^*(\mathcal{S}, \varepsilon) := \inf \{n \in \mathbb{N} : p_\varepsilon(\mathcal{S}, n) \leq \varepsilon\}. \quad (27)$$

Remark 1 (Equivalent expressions for sample complexities). *The sample complexity $n^*(p, \rho, q, \sigma, \varepsilon)$ of symmetric binary quantum hypothesis testing has the following equivalent expressions:*

$$n^*(p, \rho, q, \sigma, \varepsilon) = \inf_{\Lambda^{(n)}} \left\{ n \in \mathbb{N} : \begin{aligned} & p \operatorname{Tr}[(I^{\otimes n} - \Lambda^{(n)}) \rho^{\otimes n}] + q \operatorname{Tr}[\Lambda^{(n)} \sigma^{\otimes n}] \leq \varepsilon, \\ & 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{aligned} \right\} \quad (28)$$

$$= \inf \left\{ n \in \mathbb{N} : \frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1) \leq \varepsilon \right\} \quad (29)$$

$$= \inf \{n \in \mathbb{N} : 1 - 2\varepsilon \leq \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1\}, \quad (30)$$

where the equality (29) follows from the Helstrom–Holevo theorem in (20)–(22). By recalling the quantity $\beta_\varepsilon(\rho^{\otimes n} \| \sigma^{\otimes n})$ defined in (23), we can rewrite the sample complexity $n^*(\rho, \sigma, \varepsilon, \delta)$ of asymmetric binary quantum hypothesis testing in the following two ways:

$$n^*(\rho, \sigma, \varepsilon, \delta) = \inf_{\Lambda^{(n)}} \left\{ n \in \mathbb{N} : \begin{aligned} & \operatorname{Tr}[(I^{\otimes n} - \Lambda^{(n)}) \rho^{\otimes n}] \leq \varepsilon, \\ & \operatorname{Tr}[\Lambda^{(n)} \sigma^{\otimes n}] \leq \delta, \quad 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{aligned} \right\} \quad (31)$$

$$= \inf \{n \in \mathbb{N} : \beta_\delta(\sigma^{\otimes n} \| \rho^{\otimes n}) \leq \varepsilon\}. \quad (32)$$

See Appendix C for an explicit proof. The expression in (31) indicates that the sample complexity for asymmetric binary quantum hypothesis testing can be thought of as the minimum number of samples required to get the type I error probability below ε and the type II error probability below δ . Finally, the sample complexity $n^*(\mathcal{S}, \varepsilon)$ of M -ary quantum hypothesis testing can be rewritten as follows:

$$n^*(\mathcal{S}, \varepsilon) = \inf_{(\Lambda_1^{(n)}, \dots, \Lambda_M^{(n)})} \left\{ n \in \mathbb{N} : \sum_{m=1}^M p_m \operatorname{Tr}[(I^{\otimes n} - \Lambda_m^{(n)}) \rho_m^{\otimes n}] \leq \varepsilon \right\}, \quad (33)$$

where $\Lambda_1^{(n)}, \dots, \Lambda_M^{(n)} \geq 0$ and $\sum_{m=1}^M \Lambda_m^{(n)} = I^{\otimes n}$.

Before proceeding to the development of our main results in the next section, let us first identify some conditions under which the sample complexity of symmetric binary quantum hypothesis testing is trivial, i.e., such that it is equal to either one or infinity.

Remark 2 (Trivial cases). *Let $p, q, \varepsilon, \rho,$ and σ be as stated in Definition 3. If $\rho \perp \sigma$ (i.e., $\rho\sigma = 0$), $\varepsilon \in [1/2, 1]$, or $\exists s \in [0, 1]$ such that $\varepsilon \geq p^s q^{1-s}$, then the following equality holds*

$$n^*(p, \rho, q, \sigma, \varepsilon) = 1. \quad (34)$$

If $\rho = \sigma$ and $\min\{p, q\} > \varepsilon \in [0, 1/2)$, then

$$n^*(p, \rho, q, \sigma, \varepsilon) = +\infty. \quad (35)$$

Proof. See Appendix D. □

IV. SAMPLE COMPLEXITY RESULTS

Having defined various sample complexities of interest in Definitions 3, 4, and 5, it is clear that calculating the precise values of $n^*(p, \rho, q, \sigma, \varepsilon)$, $n^*(\rho, \sigma, \varepsilon, \delta)$, and $n^*(\mathcal{S}, \varepsilon)$ is not an easy computational problem. As such, we are then motivated to find lower and upper bounds on these sample complexities, which are easier to compute and ideally match in an asymptotic sense. We are able to meet this challenge for the symmetric and asymmetric binary settings, mostly by building on the vast knowledge that already exists regarding quantum hypothesis testing. For M -ary hypothesis testing, we are only able to give lower and upper bounds that differ asymptotically by a factor of $\ln M$. However, we note that this finding is consistent with the best known result in the classical case [54, Fact 2.4]. Before proceeding with stating our results, let us note here that all of our results hold for states acting on a separable (infinite-dimensional) Hilbert space, unless otherwise noted.

A. Symmetric binary quantum hypothesis testing

1. Two pure states

Let us begin by considering the sample complexity of symmetric binary quantum hypothesis testing when distinguishing two pure states (i.e., rank-1 projection operators), which is much simpler than the more general case of two arbitrary mixed states. It is also interesting from a fundamental perspective because there is no classical analog of pure states [55, 56], as all classical pure states correspond to degenerate (deterministic) probability distributions that are either perfectly distinguishable or perfectly indistinguishable. In this case, we find an exact result, which furthermore serves as a motivation for the kind of expression we wish to obtain in the general mixed-state case. This finding was essentially already reported in [57, Eq. (39)], with the main difference below being a generalization to arbitrary priors. In any case, Theorem 6 is a direct consequence of some simple algebra and the following equality [58, Proposition 21], which holds for (unnormalized) vectors $|\varphi\rangle$ and $|\zeta\rangle$:

$$\| |\varphi\rangle\langle\varphi| - |\zeta\rangle\langle\zeta| \|_1^2 = (\langle\varphi|\varphi\rangle + \langle\zeta|\zeta\rangle)^2 - 4|\langle\zeta|\varphi\rangle|^2. \quad (36)$$

Theorem 6 (Sample complexity: symmetric binary case with two pure states). *Let $p, q, \rho,$ and ε be as stated in Definition 3, and furthermore let $\rho = \psi \equiv |\psi\rangle\langle\psi|$ and $\sigma = \varphi \equiv |\varphi\rangle\langle\varphi|$ be pure states, such that the conditions in Remark 2 do not hold. Then*

$$n^*(p, \psi, q, \phi, \varepsilon) = \left\lceil \frac{\ln\left(\frac{pq}{\varepsilon(1-\varepsilon)}\right)}{-\ln F(\psi, \phi)} \right\rceil. \quad (37)$$




FIG. 1. Comparison between the functions $\frac{1}{-\ln x}$ and $\frac{1}{2(1-\sqrt{x})}$ for $x \in [0, 1]$, demonstrating that there is little difference in characterizing sample complexity of binary symmetric hypothesis testing by $[-\ln F(\rho, \sigma)]^{-1}$ instead of $[d_B(\rho, \sigma)]^{-2}$. The largest gap between these functions is equal to $\frac{1}{2}$ and occurs at $x = 0$.

Proof. See Appendix E. □

The exact result above indicates the kind of expression that we should strive for in the more general mixed-state case: logarithmic dependence on the inverse error probability and inverse dependence on the divergence $-\ln F(\psi, \phi)$. For pure states, by inspecting (11) and (13), we see that this latter divergence is equivalent to the Petz–Rényi relative entropy of order $1/2$, as well as the sandwiched Rényi relative entropy of order $1/2$. Furthermore, it is equivalent, up to a constant, to the same divergences for every order $\alpha \in (0, 1)$.

The characterization in Theorem 6 depends inversely on the negative logarithm of the fidelity, rather than the inverse of the squared Hellinger or Bures distance, the latter being more common in formulations of the sample complexity of classical hypothesis testing (see [59, Theorem 4.7] and [60]). However, we should note that there is little difference between the functions $\frac{1}{-\ln x}$ and $\frac{1}{2(1-\sqrt{x})}$ when $x \in [0, 1]$. Indeed, the largest gap between these functions is $1/2$, occurring at $x = 0$. As such, our characterization in terms of $[-\ln F(\rho, \sigma)]^{-1}$ instead of $[d_B(\rho, \sigma)]^{-2}$ makes little to no difference in terms of asymptotic sample complexity. Note, however, that this factor of $1/2$ can make a notable difference in applications in the finite or small sample regime. The two functions are plotted in Figure 1 in order to make this point visually clear.

2. Two general states

Let us now move on to the general mixed-state case. Theorem 7 below provides lower and upper bounds on the sample complexity for this case. The main tool for establishing both lower bounds is the generalized Fuchs-van-de-Graaf inequality recalled in (A11), and the main tool for establishing the upper bound is the Audenaert inequality recalled in (A10). Let us note that the upper bound is achieved by the Helstrom–Holevo test, sometimes also called the quantum Neyman–Pearson test; i.e., $\Lambda^{(n)}$ is a projection onto the positive part of $p\rho^{\otimes n} - q\sigma^{\otimes n}$.

Theorem 7 (Sample complexity: symmetric binary case with two general states). *Let p , q , ε , ρ , and σ be as stated in Definition 3 such that the conditions in Remark 2 do not hold. Then the following bounds hold*

$$\max \left\{ \frac{\ln(pq/\varepsilon)}{-\ln F(\rho, \sigma)}, \frac{pq - \varepsilon(1 - \varepsilon)}{pq [d_B(\rho, \sigma)]^2} \right\} \leq n^*(p, \rho, q, \sigma, \varepsilon) \leq \left[\inf_{s \in [0, 1]} \frac{\ln \left(\frac{p^s q^{1-s}}{\varepsilon} \right)}{-\ln \text{Tr}[\rho^s \sigma^{1-s}]} \right]. \quad (38)$$

Proof. See Appendix F. □

The statement given in Theorem 7 is sufficiently strong to lead to the following characterization of sample complexity of symmetric binary quantum hypothesis testing in the general case, given by Corollary 8 below. This finding essentially matches the exact result for the pure-state case in Theorem 6, up to constant factors and for small error probability ε . Corollary 8 below follows by using the first lower bound in (38) and by picking $s = 1/2$ in the upper bound in (38), along with relations between F and F_H recalled in (A7).

Corollary 8. *Let p, q, ε, ρ , and σ be as stated in Definition 3, such that the conditions in Remark 2 do not hold. Then the following inequalities hold:*

$$\frac{\ln \left(\frac{pq}{\varepsilon} \right)}{-\ln F_H(\rho, \sigma)} \leq n^*(p, \rho, q, \sigma, \varepsilon) \leq \left[\frac{\ln \left(\frac{\sqrt{pq}}{\varepsilon} \right)}{-\frac{1}{2} \ln F_H(\rho, \sigma)} \right]. \quad (39)$$

$$\frac{\ln \left(\frac{pq}{\varepsilon} \right)}{-\ln F(\rho, \sigma)} \leq n^*(p, \rho, q, \sigma, \varepsilon) \leq \left[\frac{\ln \left(\frac{\sqrt{pq}}{\varepsilon} \right)}{-\frac{1}{2} \ln F(\rho, \sigma)} \right]. \quad (40)$$

Thus, after fixing the priors p and q to be constants, we have characterized the sample complexity as follows:

$$n^*(p, \rho, q, \sigma, \varepsilon) = \Theta \left(\frac{\ln \left(\frac{1}{\varepsilon} \right)}{-\ln F_H(\rho, \sigma)} \right) = \Theta \left(\frac{\ln \left(\frac{1}{\varepsilon} \right)}{-\ln F(\rho, \sigma)} \right). \quad (41)$$

Proof. The first inequality in (39) follows from the first inequality in (38) and the inequalities in (A7). The second inequality in (39) follows from the second inequality in (38) by picking $s = 1/2$. The inequalities in (40) follow from similar reasoning and using the inequalities in (A7). □

Corollary 8 demonstrates that the asymptotic sample complexity in the classical case of commuting ρ and σ is uniquely characterized by the negative logarithm of the classical fidelity, because $F_{\text{H}}(\rho, \sigma) = F(\rho, \sigma)$ in such a case. The characterization in Corollary 8 is a strengthening of the existing characterizations of sample complexity in the classical case (see [59, Theorem 4.7] and [60]), as it has asymptotically matching lower and upper bounds and includes dependence on the error probability, and the underlying distributions (commuting states).

However, Corollary 8 also demonstrates that the asymptotic sample complexity in the quantum case does not have a unique characterization. Indeed, the quantities $-\ln F(\rho, \sigma)$ and $-\ln F_{\text{H}}(\rho, \sigma)$ are related by multiplicative constants, as recalled in (A9), and these constants get discarded when using O , Ω , and Θ notation. These constants can actually have dramatic ramifications for the quantum technology required to implement a given measurement strategy. On the one hand, the upper bound on sample complexity in (38) assumes the ability to perform a collective measurement on all n copies of the unknown state. This is essentially equivalent to performing a general unitary on all n systems followed by a product measurement, and so will likely need a full-scale, fault-tolerant quantum computer for its implementation. On the other hand, the upper bound on sample complexity in (40) can be achieved by product measurements and classical post-processing. This is in contrast to the upper bound in (38). Indeed, one can repeatedly apply a measurement known as the Fuchs–Caves measurement [61] on each individual copy of the unknown state and then process the measurement outcomes classically to achieve the upper bound in (40) (see Appendix G for details). Asymptotically, there is no difference between the sample complexities of these strategies, even though there are drastic differences between the technologies needed to realize them. Thus, if one is interested in merely achieving the asymptotic sample complexity, then the latter strategy using the Fuchs–Caves measurement and classical post-processing is preferred.

Continuing with the observation that the asymptotic sample complexity in the quantum case is not unique, let us note that one can even equivalently characterize it by means of the following family of z -fidelities, which are based on the α - z divergences [62] by setting $\alpha = 1/2$ and obey the data-processing inequality for all $z \geq 1/2$ [63, Theorem 1.2]:

$$F_z(\rho, \sigma) := \left(\text{Tr} \left[\left(\sigma^{1/4z} \rho^{1/2z} \sigma^{1/4z} \right)^z \right] \right)^2 = \left\| \rho^{1/4z} \sigma^{1/4z} \right\|_{2z}^{4z}. \quad (42)$$

When $z = 1/2$, we get that $F_{z=1/2}(\rho, \sigma) = F(\rho, \sigma)$, and when $z = 1$, we get that $F_{z=1}(\rho, \sigma) = F_{\text{H}}(\rho, \sigma)$. Furthermore, [64, Proposition 6] implies that the z -fidelities are monotone decreasing for all $z \geq 1/2$, so that

$$-\ln F(\rho, \sigma) \leq -\ln F_z(\rho, \sigma) \leq -\ln F_{z'}(\rho, \sigma) \leq -\ln F_{\text{H}}(\rho, \sigma) \leq -2 \ln F(\rho, \sigma), \quad (43)$$

for all z and z' satisfying $1/2 \leq z \leq z' \leq 1$. The inequalities in (43) combined with (41) then imply that all of these z -fidelities equivalently characterize the asymptotic sample complexity of symmetric binary quantum hypothesis testing.

B. Asymmetric binary quantum hypothesis testing

Theorem 9 below provides lower and upper bounds on the sample complexity of asymmetric quantum hypothesis testing, as introduced in Definition 4. Unlike the symmetric

setting considered in Section IV A 2, the lower bound is expressed in terms of the sandwiched Rényi divergence \tilde{D}_α , while the upper bound is expressed in terms of the Petz–Rényi divergence D_α . The main tool for establishing the lower bounds is the strong converse bound recalled in Lemma 19, and the main tool for establishing the upper bound is the quantum Hoeffding bound recalled in Lemma 20. Let us note that the upper bound can be achieved by the Helstrom–Holevo test, i.e., projection onto the positive part of $\rho^{\otimes n} - \lambda\sigma^{\otimes n}$, or the *pretty-good measurement* [65, 66] $\Lambda^{(n)} = (\rho^{\otimes n} + \lambda\sigma^{\otimes n})^{-1/2}\rho^{\otimes n}(\rho^{\otimes n} + \lambda\sigma^{\otimes n})^{-1/2}$ for some properly chosen parameter $\lambda > 0$.

Theorem 9. *Fix $\varepsilon, \delta \in (0, 1)$, and let ρ and σ be states. Suppose there exists $\gamma > 1$ such that $\tilde{D}_\gamma(\rho\|\sigma) < +\infty$ and $\tilde{D}_\gamma(\sigma\|\rho) < +\infty$. Then the following bounds hold for the sample complexity $n^*(\rho, \sigma, \varepsilon, \delta)$ of asymmetric binary quantum hypothesis testing:*

$$\begin{aligned} \max \left\{ \sup_{\alpha \in (1, \gamma]} \left(\frac{\ln \left(\frac{(1-\varepsilon)^{\alpha'}}{\delta} \right)}{\tilde{D}_\alpha(\rho\|\sigma)} \right), \sup_{\alpha \in (1, \gamma]} \left(\frac{\ln \left(\frac{(1-\delta)^{\alpha'}}{\varepsilon} \right)}{\tilde{D}_\alpha(\sigma\|\rho)} \right) \right\} &\leq n^*(\rho, \sigma, \varepsilon, \delta) \\ &\leq \min \left\{ \left[\inf_{\alpha \in (0, 1)} \left(\frac{\ln \left(\frac{\varepsilon^{\alpha'}}{\delta} \right)}{D_\alpha(\rho\|\sigma)} \right) \right], \left[\inf_{\alpha \in (0, 1)} \left(\frac{\ln \left(\frac{\delta^{\alpha'}}{\varepsilon} \right)}{D_\alpha(\sigma\|\rho)} \right) \right] \right\}. \end{aligned} \quad (44)$$

where $\alpha' := \frac{\alpha}{\alpha-1}$.

Proof. See Appendix H. □

Remark 3. *For the finite-dimensional case, let us note that*

$$\tilde{D}_\alpha(\rho\|\sigma) \leq D_{\max}(\rho\|\sigma) := \ln \inf \{ \lambda > 0 : \rho \leq \lambda\sigma \} = \ln \lambda_{\max}(\sigma^{-1/2}\rho\sigma^{-1/2}) < +\infty \quad (45)$$

for all $\alpha \in [0, \infty]$ as long as $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$ [36, Theorem 7], [67, Lemma 3.12] (see [68] for D_{\max}).

The bounds given in Theorem 9 lead to the following asymptotic characterization of the sample complexity of asymmetric binary quantum hypothesis testing, given by Corollary 10 below. Interestingly, the asymptotic sample complexity given in (48) below establishes, up to a multiplicative constant, the following asymptotic relationship between the number n of samples and the type II error probability δ :

$$n \simeq \frac{\ln \left(\frac{1}{\delta} \right)}{D(\rho\|\sigma)} \quad \Leftrightarrow \quad \delta \simeq e^{-nD(\rho\|\sigma)}, \quad (46)$$

whenever $\varepsilon \in (0, 1)$ is fixed to be a constant and provided that δ is sufficiently small. This characterization is consistent with that from the quantum Stein’s lemma [8, 9], which states that the largest decaying rate of the type II error probability is governed by the quantum relative entropy whenever the type I error probability is at most a fixed $\varepsilon \in (0, 1)$; i.e.,

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \ln \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) = D(\rho\|\sigma). \quad (47)$$

Corollary 10 (Bounds on sample complexity of asymmetric binary quantum hypothesis testing). *Fix $\varepsilon, \delta \in (0, 1)$, and let ρ and σ be as stated in Definition 4. For sufficiently small δ , the sample complexity of asymmetric binary quantum hypothesis testing satisfies the following bounds:*

$$\frac{2}{5} \left(\frac{\ln(\frac{1}{\delta})}{D(\rho\|\sigma)} \right) \leq n^*(\rho, \sigma, \varepsilon, \delta) \leq \left\lceil 4 \left(\frac{\ln(\frac{1}{\delta})}{D(\rho\|\sigma)} \right) \right\rceil, \quad (48)$$

provided that there exists $\gamma \in (0, 1]$ such that $D_{1+\gamma}(\rho\|\sigma) < +\infty$.

Proof. See Appendix I. Also, see (I71) and (I84) for the precise meaning of “sufficiently small δ .” \square

C. Multiple quantum hypothesis testing

Now we generalize the binary mixed-state case discussed in Section IV A 2 to an arbitrary number of states, M . The lower bound essentially follows from the lower bound for the two-state case because discriminating M states simultaneously is not easier than discriminating any pair of states. The upper bound is achieved by the *pretty-good measurement* [65, 66], which can be implemented via a quantum algorithm [69]. The main tool for showing the upper bound is by testing each state against any of the other states. Since there are $M(M-1)$ such pairs, the upper bound ends up with a logarithmic dependence on M . This idea was introduced in [70, Theorem 4], and a similar idea has been used in [71] and [72, Section 4]. Let us note that the upper bound in (49) below is a refinement of that presented in [71].

Theorem 11 (Sample complexity of M -ary quantum hypothesis testing). *Let $n^*(\mathcal{S}, \varepsilon)$ be as stated in Definition 5. Then,*

$$\max_{m \neq \bar{m}} \frac{\ln\left(\frac{p_m p_{\bar{m}}}{(p_m + p_{\bar{m}})\varepsilon}\right)}{-\ln F(\rho_m, \rho_{\bar{m}})} \leq n^*(\mathcal{S}, \varepsilon) \leq \left\lceil \max_{m \neq \bar{m}} \frac{2 \ln\left(\frac{M(M-1)\sqrt{p_m}\sqrt{p_{\bar{m}}}}{2\varepsilon}\right)}{-\ln F(\rho_m, \rho_{\bar{m}})} \right\rceil. \quad (49)$$

Proof. See Appendix J. \square

Remark 4. *Ref. [15, Eq. (36)] established a multiple quantum Chernoff bound for M -ary hypothesis testing with error exponent $\min_{m \neq \bar{m}} C(\rho_m\|\rho_{\bar{m}})$, which holds for finite-dimensional states. This result also implies*

$$n^*(\mathcal{S}, \varepsilon) \leq O\left(\max_{m \neq \bar{m}} \frac{\ln M}{-\ln F(\rho_m, \rho_{\bar{m}})}\right), \quad (50)$$

similarly to the upper bound in (49).

Proof. See Appendix K. \square

Remark 5 (Pure-state case). *For a collection of pure states, the following upper bound can be derived by using [72, Eq. (9)]:*

$$n^*(\mathcal{S}, \varepsilon) \leq \left\lceil \max_{m \neq \bar{m}} \frac{\ln\left(\frac{M(M-1)(p_m^2 + p_{\bar{m}}^2)}{2p_m p_{\bar{m}} \varepsilon}\right)}{-\ln F(\rho_m, \rho_{\bar{m}})} \right\rceil. \quad (51)$$

Remark 6 (Classical case). *As far as we know, the tightest upper bound on the sample complexity of M -ary hypothesis testing in the classical scenario is as follows [73, Theorem 15]:*

$$n^*(\mathcal{S}, \varepsilon) \leq \left[\max_{m \neq \bar{m}} \inf_{s \in [0,1]} \frac{\ln \left(\frac{M p_m^s p_{\bar{m}}^{1-s}}{\varepsilon} \right)}{-\ln \text{Tr}[\rho_m^s \rho_{\bar{m}}^{1-s}]} \right], \quad (52)$$

which has an improved logarithmic dependence on M .

By employing the same analysis as given in the proof of Corollary 8, we obtain the following results:

Corollary 12 (Asymptotic sample complexity). *Let \mathcal{S} be as stated in Definition 5. If we regard each element of $\{p_m\}_{m=1}^M$ and ε as constants, and suppose that*

$$\frac{p_m p_{\bar{m}}}{p_m + p_{\bar{m}}} \geq \varepsilon, \quad (53)$$

for $m, \bar{m} \in \{1, 2, \dots, M\}$, and $\bar{m} \neq m$, then we have characterized the sample complexity of M -ary quantum hypothesis testing as follows:

$$\Omega \left(\frac{1}{\min_{m \neq \bar{m}} [-\ln F(\rho_m, \rho_{\bar{m}})]} \right) \leq n^*(\mathcal{S}, \varepsilon) \leq O \left(\frac{\ln M}{\min_{m \neq \bar{m}} [-\ln F(\rho_m, \rho_{\bar{m}})]} \right). \quad (54)$$

The above bounds also hold when replacing F with F_{H} .

V. APPLICATIONS AND FURTHER AVENUES OF RESEARCH

We briefly discuss three broad areas in quantum science where the concept of sample complexity in quantum hypothesis testing is relevant. For those who come from diverse fields, our goal here is to illustrate how sample complexity can appear in widely different and perhaps unexpected areas – thus it should not be confined to limited areas in information theory or computer science alone. We try to distill the intuitive essence of these connections that have appeared in various literature and to encourage the reader to explore these in further detail and to identify new relationships.

In Section VA, we explore possible uses of sample complexity in the proofs of optimality of certain quantum algorithms. In Section VB, we devise ways of understanding sample complexity in the learning framework for quantum states. In Section VC, we note the relevance of sample complexity in the foundations of quantum information and quantum mechanics. Sample complexity in these contexts can serve as an improved technical tool, introduce a modified framework for an old problem, or provide new interpretations that originate from seemingly disparate areas.

Let us emphasise here again that our results for sample complexity, as presented in Section IV, hold for the discrimination not only of discrete-variable (qubits and qudits) quantum systems, but also continuous-variable quantum systems (qumodes) as well as hybrid discrete-continuous variable quantum systems, unless otherwise stated. These sample complexity bounds also hold for mixed states and multiple hypotheses M . In the applications discussed, while many previous works have mostly considered $M = 2$ and pure qubit-based systems, our findings directly extend the applicability of those results to $M > 2$, mixed states, continuous-variable quantum systems and also to hybrid discrete-continuous variable settings.

A. Optimality of quantum algorithms

1. Quantum simulation of linear ordinary and partial differential equations

Any linear system of ordinary differential equations (ODEs) and linear partial differential equations (PDEs) can be represented in the following way:

$$\frac{d\mathbf{u}(t)}{dt} = -i\mathbf{A}(t)\mathbf{u}(t), \text{ with } \mathbf{u}(0) = \mathbf{u}_0, \quad (55)$$

For a system of D linear ordinary differential equations for D scalar functions $\{u_j(t)\}_{j=1}^D$ and with $\{|j\rangle\}_j$ an orthonormal basis, then $\mathbf{u}(t) \equiv \sum_{j=1}^D u_j(t)|j\rangle$ in (55) is a D -dimensional vector and $\mathbf{A}(t)$ is a $D \times D$ matrix. We note that any inhomogeneous terms \mathbf{f} and higher-order time derivatives can always be accommodated through an appropriate dilation, e.g., $\mathbf{u} \rightarrow \mathbf{u} \otimes |0\rangle + \mathbf{f} \otimes |1\rangle$ and $\mathbf{u} \rightarrow \mathbf{u} \otimes |0\rangle + d\mathbf{u}/dt \otimes |1\rangle$, respectively (for example, see [74, 75]). When D is large enough, this system can also represent a discretised linear partial differential equation. For example, in Eulerian discretisation (e.g., finite difference methods), where n is the size of the discretisation and d is the spatial dimension of the PDE, so that $D \sim O(n^d)$. We will discuss the continuous representation of the partial differential equation later.

A pure quantum state vector is the normalised vector $|u(t)\rangle \equiv \mathbf{u}(t)/\|\mathbf{u}(t)\|$, where normalisation is through the l_2 norm $\|\cdot\|$. For Schrödinger-like equations, $\mathbf{A}(t)$ is a Hermitian matrix, so that (55) can be directly approached through quantum simulation [76, 77] with $\mathbf{A}(t)$ the corresponding Hamiltonian. This means that the transformation between $|u(0)\rangle$ and $|u(t)\rangle$ is unitary and the corresponding norms are preserved. For example, if \mathbf{A} is time-independent, we can simulate the final state $|u(t)\rangle = \exp(-i\mathbf{A}t)|u(0)\rangle$ through the unitary operator $\exp(-i\mathbf{A}t)$. However, for general ODEs, $\mathbf{A}(t)$ is *not* necessarily Hermitian. This means that in order to simulate $|u(t)\rangle$ through quantum simulation, we need to dilate the state $\mathbf{u}(t)$ so that evolution in the dilated space is unitary. Various dilation procedures are possible, including Schrödingerisation [74, 75, 78], block-encoding [79–81], or Stinespring dilation [82]. While explicit procedures are necessary for actual implementation of these methods, we will see that the usefulness of sample complexity in quantum hypothesis testing lies in providing us with a lower bound on the optimal cost of any quantum algorithm to prepare $|u(t)\rangle$ [81].

The basic suggestive idea is the following (for example, see [81]). Suppose we begin with two different initial conditions \mathbf{u}_0 and $\tilde{\mathbf{u}}_0$ for the same differential equation in (55). When \mathbf{A} is Hermitian like in the example above, then the distance is preserved; i.e.,

$$\|\mathbf{u}(t) - \tilde{\mathbf{u}}(t)\| = \|\exp(-i\mathbf{A}t)(\mathbf{u}_0 - \tilde{\mathbf{u}}_0)\| = \|\mathbf{u}_0 - \tilde{\mathbf{u}}_0\|. \quad (56)$$

Thus, the distinguishability of the two initial states and the distinguishability of the two final states do not change since unitary evolution preserves $\|\mathbf{u}(t) - \tilde{\mathbf{u}}(t)\|$ for all t . However, in more general cases, $\mathbf{A} = \mathbf{A}_1 + i\mathbf{A}_2$ is *not* Hermitian (where $\mathbf{A}_1, \mathbf{A}_2$ are Hermitian). Define the ratio R as follows:

$$R \equiv \|\mathbf{u}(t) - \tilde{\mathbf{u}}(t)\|/\|\mathbf{u}_0 - \tilde{\mathbf{u}}_0\| = \|\exp(\mathbf{A}_2 T)(\mathbf{u}_0 - \tilde{\mathbf{u}}_0)\|/\|\mathbf{u}_0 - \tilde{\mathbf{u}}_0\|. \quad (57)$$

Then if we assume the spectral norm of $\exp(\mathbf{A}_2 t)$ is large – i.e., $\max_{\mathbf{v} \neq \mathbf{0}} (\|\exp(\mathbf{A}_2 t)\mathbf{v}\|/\|\mathbf{v}\|) = (\max \text{eigenvalue of } \exp(2\mathbf{A}_2 t))^{1/2}$ is large – then R grows with t when we assume $\mathbf{v} = \mathbf{u}_0 - \tilde{\mathbf{u}}_0$.

This means that an initially hard to distinguish pair $(\mathbf{u}_0, \tilde{\mathbf{u}}_0)$ can become easily distinguishable when t is large enough or when \mathbf{A}_2 has large enough eigenvalues. The latter case occurs when the differential equation is very dissipative, like the heat equation with high diffusion coefficients. In this case, a quantum simulation algorithm for (55) with non-Hermitian \mathbf{A} can act as a state discriminator. Since the optimal sample complexity n^* in quantum hypothesis testing provides the optimal cost in performing state discrimination, this sample complexity n^* can be used to derive a lower bound on the cost for the successful quantum simulation of (55) with some maximum error probability. Note that here we have ignored many details and subtleties (like differences between distances between classical vectors and those between quantum states) to focus only on conveying the intuitive essence of the idea. The arguments above can be refined with the improved n^* lower bounds in this paper and for time-dependent $\mathbf{A}(t)$. Although better bounds than those provided by quantum hypothesis testing are possible [81], these require stronger assumptions than are usually considered in state discrimination – assuming access to not only the unitary operator creating the state, which then effectively allows amplitude amplification.

We can extend this reasoning also to linear PDEs in their continuous representation, without discretisation. In this case, Eq. (55) can be used to represent a $(d + 1)$ -dimensional PDE using the continuous-variable representation $\mathbf{u}(t) \equiv \int_{-\infty}^{\infty} u(t, x) |x\rangle dx$, where we use $x = (x_1, \dots, x_d)$ to denote the d -dimensional spatial degrees of freedom in the PDE. Here $|x_j\rangle$ is a position eigenstate with corresponding position operator \hat{x}_j , where $\hat{x}_j |x_j\rangle = x_j |x_j\rangle$. In this case \mathbf{A} is no longer a finite matrix like for ODEs, but it is rather an infinite-dimensional operator acting on $\mathbf{u}(t)$. Here $\mathbf{A}(t)$ involves both position and momentum operators \hat{x}_j, \hat{p}_j , for $j \in \{1, \dots, d\}$, obeying the commutation relations $[\hat{x}_j, \hat{p}_j] = iI$. To derive the form of $\mathbf{A}(t)$ from the original PDE, it can be shown, for example in [75], that one only needs to make the replacement $xu(t, x) \rightarrow \hat{x}\mathbf{u}(t)$ and $\partial^k u(t, x)/\partial^k x \rightarrow (i\hat{p})^k \mathbf{u}(t)$. In fact, Eq. (55) can also represent a system of linear PDEs, by using a hybrid continuous-variable and discrete variable representation [75].

It is important to note that the sample complexities, denoted by n^* , presented in this paper are also applicable to continuous variables as well as hybrid systems. Thus the basic argument for deriving lower bounds on the optimal cost in quantum simulation of PDEs and system of PDEs using the sample complexity in quantum hypothesis testing follows. However, given subtleties with infinite-dimensional systems, these ideas need to be refined and further explored.

So far we have only considered the embedding of scalar degrees of freedom into vectors $\mathbf{u}(t)$, both finite and infinite dimensional. However, there are also differential equations for matrices $\Xi(t)$, for instance

$$\frac{d\Xi(t)}{dt} = \mathcal{L}(\Xi(t)), \quad \Xi(0) = \Xi_0, \quad (58)$$

where \mathcal{L} is a linear superoperator. We can consider the embedding of this matrix into an unnormalised density matrix, and quantum simulation for density matrices can be subsequently used. Similar arguments now based on quantum hypothesis testing for mixed states could potentially then be applied in this case – in certain regimes – to identify lower bounds on the optimal cost in the quantum simulation of (58).

Linear algebra—Another class of related applications is to consider quantum algorithms for linear systems of equations, for example [83], which exploits quantum phase estimation. Alternative methods of solving these problems through a dynamical process like (55) is

also possible. This is done by mapping the discrete problem onto iterative algorithms and then taking the continuous-time limit [84]. Thus, lower bounds for optimal algorithms for quantum algorithms for linear algebra can also be approached through n^* and can be further explored.

2. Quantum simulation of nonlinear ordinary and partial differential equations

Previously we saw that the optimal sample complexity in quantum hypothesis testing is only relevant for linear differential equations when the dynamics served to increase the distance between two initially closely-separated states, thus making them easier to distinguish. This dynamics cannot include, for instance, purely unitary transformations. In the presence of *nonlinear* dynamics, on the other hand, initially closely-separated states can be driven apart very quickly. The rate of separation is related to the Lyapunov coefficient of the dynamical system and depends on the degree of nonlinearity of the differential equation. For example, see [85].

Suppose that two initial conditions \mathbf{u}_0 and $\tilde{\mathbf{u}}_0$ satisfy $\|\mathbf{u}_0 - \tilde{\mathbf{u}}_0\|^2 \sim \eta_{|\lambda|}(0)$ and they evolve in time t under the same nonlinear dynamics characterised by some parameter $|\lambda|$ (we can let $|\lambda| = 0$ denote purely linear and unitary dynamics). We assume $\eta_{|\lambda|}(0)$ to be a constant independent of $|\lambda|$. However, as time grows, $\|\mathbf{u}(t) - \tilde{\mathbf{u}}(t)\|^2 \sim \eta_{|\lambda|}(t)$, so that $\eta_{|\lambda|}(t)$ does depend on $|\lambda|$ for $t > 0$. For many nonlinear dynamics of interest, $\eta_{|\lambda|}(t)$ increases with increasing t and $|\lambda|$, thus making the states easier to distinguish. Thus the nonlinear dynamics can serve as a state discriminator. We can now approach this in a similar way to previous arguments in the linear non-unitary dynamics scenario. Assuming $\|\mathbf{u}(t)\| \sim \|\tilde{\mathbf{u}}(t)\|$ and we embed $\mathbf{u}(t)$ and $\tilde{\mathbf{u}}(t)$ into the amplitudes of pure quantum states with density matrices ρ and $\tilde{\rho}$, respectively, then $\|\mathbf{u}(t) - \tilde{\mathbf{u}}(t)\|^2 \sim d_H^2(\rho, \tilde{\rho})\|\mathbf{u}(t)\|^2$. Ignoring the normalisation constant, we therefore see that the optimal sample complexity $n^* \sim 1/d_H^2 \sim 1/\eta_{|\lambda|}(t)$ in distinguishing quantum states $|u(t)\rangle$ and $|\tilde{u}(t)\rangle$ up to some maximum failure probability ε provides a lower bound on the quantum simulation of nonlinear dynamics. It can also place an upper bound on the amount of nonlinearity in a differential equation if we require the quantum simulation to be efficient for that nonlinear dynamics. In an illustrative sense, imposing efficiency with respect to t for the optimal quantum algorithm requires $n^* \sim 1/\eta_{|\lambda|}(t) \lesssim \text{poly}(t)$, where $\eta_{|\lambda|}(t)$ can in principle be derived or approximated from the known nonlinear dynamics. This therefore puts an upper bound on $|\lambda|$. Thus for nonlinear ODEs with high enough nonlinearity, $\eta \sim \exp(-t)$ is possible, implying $n^* \sim \exp(t)$, which means any quantum algorithm is inefficient for high enough nonlinearities. For example, in [86], a system of two differential equations with quadratic nonlinearity was studied, and such an upper bound for a characterisation of nonlinearity was identified.

Given more precise bounds n^* in the current paper, previous arguments can be refined and explored further. The analysis can also be extended to continuous-variable settings, as well as embedding of the nonlinear dynamics into density matrices, instead of pure states, by exploiting our optimal sample complexity bounds in these scenarios. Thus we can also extend to applications beyond (58) to nonlinear differential matrix equations

$$\frac{d\Xi(t)}{dt} = \mathcal{N}(\Xi(t)), \quad \Xi(0) = \Xi_0, \quad (59)$$

where \mathcal{N} can be a nonlinear superoperator. Nonlinear differential equations for matrices include important classes like the Riccati differential equation, which appear in many areas

ranging from optimal control [87], estimation problems [88], and network theory [89]. It is also interesting to explore the role of the maximum error probability ε and the interplay with the characterisation of nonlinearity.

Nonlinear quantum mechanics—We can also find applications of sample complexity bounds in quantum hypothesis testing to determining the regime of validity for quantum simulation via effective nonlinear quantum mechanics [90, 91]. For example, one can find an upper bound on the time for which certain nonlinear quantum mechanical models remain valid. For instance, it is known that in the nonlinear Gross–Pitaevskii model with strength g , it is possible to distinguish two states separated by distance η in time $t \sim (1/g) \ln(1/\eta)$. We know from the optimal sample complexity that $n^* \sim 1/\eta$, where n^* can also be interpreted as the number of particles whose effective dynamics obeys the Gross–Pitaevskii equation. This means that $t \lesssim (1/g) \ln n^*$ [91]. Further exploration with more models and using more precise n^* bounds and including the maximum failure probability could give further insight on both emergent nonlinearity in quantum mechanics and the computational power of different physical models.

3. Unstructured search and related applications

In a typical search problem for an unstructured dataset, the task is to identify an unknown value r that is an integer in the set $\{1, \dots, d\}$ (for example, for a given function f , one seeks a value of r for which $f(r) = 1$). Grover’s algorithm is a well-established quantum algorithm for this task [92], and it accomplishes the desired goal by preparing an approximation of the unknown state $|r\rangle$ when we have access to an oracle that can validate whether or not we have the correct state $|r\rangle$. In a continuous-time version of Grover’s algorithm, the oracle access one assumes is a unitary generated by the Hamiltonian $H = |r\rangle\langle r| + |s\rangle\langle s|$ acting on a d -dimensional complex Hilbert space, where $|s\rangle = (1/\sqrt{d}) \sum_{j=1}^d |j\rangle$. Evolving $|s\rangle$ by $\exp(-iHt)$, it can be easily shown that the probability of the final state being $|r\rangle$ is equal to one when $t \sim \sqrt{d}$ [93].

Suppose that we modify this search problem into a binary problem, where we ask whether or not this r value even exists (e.g., whether a solution to $f(\cdot) = 1$ exists for a given f). There are then two different corresponding Hamiltonians $H_0 = |r\rangle\langle r| + |s\rangle\langle s|$ and $H_1 = |s\rangle\langle s|$. Solving the binary version of the search problem means we want to distinguish between the oracles $U_0 = \exp(-iH_0t)$ and $U_1 = \exp(-iH_1t)$, and we want to distinguish between them by using as few samples of the initial states (and thus also access to the unitaries) as possible. The analysis in [93] reduces the optimality question to a state discrimination problem between states with distance η . For example, we can use a Hadamard test (see [91] and similar ideas in [90]), which is a circuit involving either controlled- U_0 or controlled- U_1 . The overlap of the two possible final states for the Hadamard test circuit can approach a constant for an appropriate choice for t like $t \sim \sqrt{d}$. The optimal sample complexity in quantum hypothesis testing $n^* \sim \ln(1/\varepsilon)/\eta$ and therefore produces a lower bound on sample complexity for the basic discrimination problem with a maximum failure probability ε . See [57] for a similar discussion. Extensions to the continuous-variable scenario and to mixed quantum states could also be valuable.

Other applications—The modification of the unstructured search problem into the binary problem above can also be mapped onto other problems, where the optimal sample complexity in quantum hypothesis testing provides a lower bound on the optimal complexity for

those algorithms. These other problems include quantum fingerprinting [94], orthogonality testing [57], and also density matrix exponentiation [57, 95].

4. Hidden subgroup problem

The hidden subgroup problem remains one of the most prominent classes of problems for which quantum algorithms have been developed, including Shor’s factoring algorithm [96], Shor’s quantum algorithm for discrete logarithms [96], and also the earliest quantum algorithms like Simon’s algorithm [97] and the Deutsch–Jozsa algorithm [98]. Here we are given a group G , and we want to find a generating set for a subgroup $H \subset G$, where H is a so-called ‘hidden subgroup’. For a finite set S , a function $F : G \rightarrow S$ is said to ‘hide’ the subgroup H if $F(g_1) = F(g_2)$ for all $g_1, g_2 \in G$ if and only if $g_1H = g_2H$. Then if F is given via an oracle, the task is to determine H through minimal number of queries to the oracle. This problem can actually be rephrased as a multiple quantum hypothesis testing problem [71]. Here we are given a coset state defined by $\rho_H = (1/|G|) \sum_{g \in G} |gH\rangle\langle gH|$ where $|gH\rangle = (1/\sqrt{|H|}) \sum_{h \in H} |gh\rangle$. Let the number of subgroups be M . Then the goal of identifying $H \subset G$ involves using a minimal number of samples to distinguish between M different coset states. Thus the bound $n^* \lesssim \ln(M)$ in quantum hypothesis testing for mixed states can be used to upper bound the query complexity $\sim \ln(M)$ for solving the hidden subgroup problem. Using the more precise bounds for n^* obtained in this paper, the bounds for different hidden subgroup problems can then be refined.

B. Quantum learning and classification

The goal of quantum classification is to classify an unknown quantum state σ into one of M classes, where we are not given any prior classical information about σ . That is, we would like to identify conditions under which different states cannot be discriminated (i.e., belong to the same class), whereas the goal of state discrimination is to identify how to discriminate states. The task of quantum classification requires the construction of a quantum classifier, which can come in different forms: (a) deterministic or (b) probabilistic. A deterministic classifier gives a deterministic outcome for the predicted class whereas a probabilistic classifier gives a stochastic output for the predicted class. For our current purpose, we only discuss the following probabilistic quantum classifier, for illustrative purposes.

Definition 13 (Probabilistic quantum classifier). *Suppose that for any input quantum state $\sigma^{\otimes n}$ where σ is selected from a given (can be unknown) distribution \mathcal{D} , it is guaranteed that σ already belongs to one of M distinct classes, labelled $c_\sigma \in \{1, \dots, M\}$ (true label of σ). We can define a probabilistic quantum classifier for σ with respect to a given n by a set $\{\Lambda_k^{(n)}\}_{k=1}^M$ with the following properties. (i) The set forms a POVM, so that $\Lambda_k^{(n)} \geq 0$ for all k and $\sum_{k=1}^M \Lambda_k^{(n)} = I^{\otimes n}$; (ii) $\text{Tr}[\Lambda_k^{(n)} \sigma^{\otimes n}]$ corresponds to the probability that the predicted label of σ is k .*

We remark that more generally, c_σ just needs to belong to a set with cardinality M , but we choose the set $\{1, \dots, M\}$ throughout for simplicity.

In the above definition we have not defined what it means to have a successful classifier, which we will discuss later. For example, a perfect classifier with respect to a given

n is such that, for every quantum state σ selected from the distribution \mathcal{D} , the conditions $\text{Tr}[\Lambda_{k=c_\sigma}\sigma^{\otimes n}] = 1$ and $\text{Tr}[\Lambda_{k\neq c_\sigma}\sigma^{\otimes n}] = 0$ hold. However, this is usually not possible. For whatever figures of merit for success, the construction of such a POVM requires some information about the classes. In the context of supervised learning problems, we are given training data, which consists of a set of states and their corresponding known (true) labels.

Definition 14 (Training dataset). *The training dataset with N states for the classification problem with M classes is the set $\mathcal{T}_N = \{(s_i, c_i)\}_{i=1}^N$, where s_i is a state and c_i is its known (true) corresponding label where $c_i \in \{1, \dots, M\}$. There are three main categories of training data \mathcal{T} that we can consider for our quantum classification problem.*

1. s_i is the full classical description of the corresponding quantum state Σ_i . We then call \mathcal{T}_N a classical training data set.
2. s_i is the quantum state Σ_i itself and partial classical information is given. Here we call \mathcal{T}_N a partially quantum training data set.
3. s_i is the quantum state Σ_i itself with no other information provided about Σ_i . Here we call \mathcal{T}_N a fully quantum training data set.

We note that, unless otherwise stated, Σ_i can be either a finite-dimensional, infinite-dimensional (continuous-variable) quantum state, or even a hybrid (finite-dimensional and infinite dimensional) state. In our formulation, c_i is a classical label, and we do not here consider the more general case where c_i can be a quantum state itself.

Then our M -ary supervised quantum classification problem involves two steps. Given an unknown quantum state σ selected from some distribution \mathcal{D} , the task is to assign to σ one of M possible distinct labels. The first step – training (or learning) step – is the construction of a (probabilistic) quantum classifier – given access to a training data set \mathcal{T}_N . This usually involves an optimisation procedure after being given a figure of merit (loss function) to optimise. The second step – classifying step – is the computation of the predicted class of σ when given access to σ . These two steps – training and classifying – are distinct and therefore have different costs associated with them. We can state these costs in an informal way below.

Definition 15 (Training cost). *The cost in constructing (or learning) a quantum classifier is the cost required to build this classifier given access to \mathcal{T}_N , subject to a bound in precision for a given figure of merit.*

Definition 16 (Classifying cost). *This is the minimal number of copies n of the unknown quantum state σ required to make a prediction of its class subject to a given bound in precision for a given figure of merit.*

For example, for our probabilistic quantum classifier $\{\Lambda_k^{(n)}\}_{k=1}^M$, a minimum of n copies is needed.

At the moment we have not yet discussed the criteria to identify appropriate figures of merit. There are different figures depending on one's applications, requirements, and constraints. The criterion closest in spirit to error probability in quantum hypothesis testing is the notion of training error defined below. See [99] for a discussion on the training error below for $n = 1$ and the relationship to quantum hypothesis testing. There are also other important figures of merit in learning apart from training error. This includes *test error* and also *generalisation error*. We will not go into these concepts here, but interested readers can refer to [100].

Definition 17 (Training error). *Suppose we are given the fully quantum training dataset \mathcal{T}_N defined in Definition 14 with $\Sigma_i = \rho_i^{\otimes n}$. Then the training error ε of a probabilistic quantum classifier in Definition 13 is the probability that this classifier makes the wrong prediction, i.e.,*

$$\varepsilon_\Lambda := \frac{1}{N} \sum_{(c_i, \rho_i^{\otimes n}) \in \mathcal{T}_N} (1 - \text{Tr}[\Lambda_{c_i}^{(n)} \rho_i^{\otimes n}]) = 1 - \frac{1}{N} \sum_{(c_i, \rho_i^{\otimes n}) \in \mathcal{T}_N} \text{Tr}[\Lambda_{c_i}^{(n)} \rho_i^{\otimes n}]. \quad (60)$$

If there are M possible classes, then c_j can only take M possible different values. If we are given sufficient training data, whereby $N > M$, then there must exist an equivalence class of states \mathcal{S}_j containing ρ_i, ρ_j with $i \neq j$ where $c_i = c_j$. We can therefore relabel any c_i with $i > M$ by c_j with $j \leq M$ since there are only M distinct labels. The size of this equivalence class we can label as N_j , where $j \in \{1, \dots, M\}$ and $\sum_{j=1}^M N_j = N$. Furthermore, suppose we consider the partially and fully quantum training dataset such that all the quantum states in the same equivalence class are in fact identical states. This means $\rho_i = \rho_j$ for all i, j where $c_i = c_j$. Thus, one quantum state defines its own class and N_j corresponds to the number of copies of ρ_j one has in \mathcal{T}_N for $j \in \{1, \dots, M\}$. Such a training set \mathcal{T}_N is then equivalent to $\{N_j \text{ copies of } (c_j, \rho_j^{\otimes n})\}_{j=1}^M$ with $\sum_{j=1}^M N_j = N$. If we are only allowed to select one state ρ_j at a time and we are given N chances to make the selection, then this is also equivalent to $\{p_j, \rho_j^{\otimes n}\}_{j=1}^M$ where $p_j = N_j/N$ is the probability that the state ρ_j is selected. Note that in this scenario $c_j = j$ is a redundant label once ρ_j is given. Clearly, this can also be considered the input of an M -ary quantum hypothesis testing scenario if we are allowed to take n samples of the states ρ_j .

If the POVM $\{\Lambda_{c_j}^{(n)}\}_{j=1}^M$ is used for the quantum state discrimination problem, the corresponding expected error probability is

$$p_{e,\Lambda} = 1 - \sum_{c_j=j=1}^M p_j \text{Tr}[\Lambda_{c_j}^{(n)} \rho_j^{\otimes n}] = \varepsilon_\Lambda, \quad (61)$$

which coincides with the training error probability in Definition 17. This means that minimising ε_Λ over $\{\Lambda_{c_j}^{(n)}\}_{j=1}^M$ (i.e., optimising the training error) is equivalent to minimising $p_{e,\Lambda}$ above, which corresponds to the optimal error probability in state discrimination. For example, when $M = 2$ and $n = 1$, the optimal training error coincides with the result given by the Helstrom–Holevo theorem in (22) (also see [99]).

We note that if \mathcal{T}_N were instead the classical training dataset, then each equivalence class $|\mathcal{S}_j| = 1$ because each s_i is in the same class, being identical and can be tested to be identical without using more resources. Thus $p_j = 1/M$.

To see the relevance of sample complexity for symmetric hypothesis testing in the learning context, we can consider the following. The optimal sample complexity corresponds to the optimal number of copies n of σ required to make a prediction to some maximum error probability $\varepsilon_\Lambda \leq \varepsilon$. Thus we can define optimal sample complexity in learning as follows

$$N^*(\varepsilon) := \inf_{\{\Lambda_{c_1}^{(n)}, \dots, \Lambda_{c_M}^{(n)}\}} \{n \in \mathbb{N} : \varepsilon_\Lambda^{(n)} = 1 - \sum_{c_j=j=1}^M p_j \text{Tr}[\Lambda_{c_j}^{(n)} \rho_j^{\otimes n}] \leq \varepsilon\}, \quad (62)$$

by analogy with Definition 5.

We can interpret the constraint $\varepsilon_\Lambda \leq \varepsilon$ in (62) in different ways. In the context of interpreting ε_Λ as training error, N^* here is the optimal classifying cost in Definition 16 when we fix an upper bound ε on the training error. This upper bound can be motivated in different ways. Often for training data, a regularisation is required to prevent over-fitting. This happens when the training error might be very small or near perfect, but it could result in the classifier performing badly on new data not in the original training dataset. To prevent this, some non-zero training error is required. Thus there exists some ε_{reg} where $\varepsilon_{\text{reg}} < \varepsilon_\Lambda$. This means that any upper bound should be chosen to satisfy $\varepsilon \geq \varepsilon_{\text{reg}}$. It is also possible in some scenarios to constrain the maximum allowed training error directly, so that ε can be set to be a constant.

In the case of binary classification where $M = 2$, if \mathcal{T}_N defined above coincides with the true distribution of states from which σ is selected, then the $\varepsilon_\Lambda \leq \varepsilon$ condition actually corresponds to the classifier being called ε -approximately correct in the context of probably approximately correct (PAC) learning [26, 101–103]. However, unlike PAC learning, our sample complexity is defined with respect to a given distribution of states, and so it is not identical to the sample complexity defined in PAC learning.

In the case of arbitrary M classes, the corresponding optimal classifying cost is $N^*(\varepsilon) \leq O(\ln(M))$, where we use the result in Theorem 11 and we ignore dependencies on all parameters except M . For example, we see that the optimal protocol from quantum hypothesis testing is exponentially more efficient in M compared to the algorithm ‘classification via state discrimination’ in [104], where σ is pairwise compared to each ρ_j in \mathcal{T}_N . Of course, this $O(\ln M)$ scaling is also present for classical states, so this improvement is not due to the effect of exploiting quantum correlations in collective measurements. This $O(\ln(M))$ scaling is recovered, for example, for pure states in a different algorithm having access to Helstrom measurements [104].

Here we do not discuss the training cost in constructing the classifier subject to a maximum ε in training error. To identify the classifier, a semidefinite program is required, and the cost would vary depending on whether dataset is classical, partially quantum, or fully quantum. Since this is not directly related to sample complexity, we defer this discussion to future work.

Now we can consider the asymmetric hypothesis testing scenario. In symmetric hypothesis testing, the error probabilities due to a false positive (type I error) result and a false negative (type II error) result are weighted in the same way. This corresponds to the situation of binary quantum classification where, in the expected error probability in (61), we have $p_j = 1/2$ for $j \in \{1, 2\}$. However, in asymmetric hypothesis testing, these two types of errors are treated differently.

This asymmetric setting for hypothesis testing has been previously connected to robustness of quantum classifiers [105]. Intuitively, if the optimal hypothesis test performs poorly in distinguishing between the original state σ and a perturbed state σ' , then a quantum classifier will likely categorise σ, σ' into the same class and is thus robust against perturbations $\sigma \rightarrow \sigma'$. For example, it was shown in [105] that asymmetric quantum hypothesis testing provides the robust region around σ when the optimal type II error probability is greater than $1/2$. This formalism involves constraining the type I error probability while optimising the type II error probability, which is the usual setup in asymmetric hypothesis testing.

In the definition of optimal sample complexity in Definition 4, however, we must place constant upper bounds on both type I and type II errors and only optimise over the number of copies of the input quantum states put into the classifier. This corresponds to a more

complex learning scenario, which is related but not identical to the asymmetric loss function introduced in [106] (this corresponds to the scenario where $p_1 \neq p_2$). In our case, the sample complexity corresponds to the minimal classifying cost for the probabilistic quantum classifier we defined where we fix upper bounds to the error probability terms $p_1 \text{Tr} \left[\Lambda_2^{(n)} \rho_1^{\otimes n} \right]$ and $p_2 \text{Tr} \left[\Lambda_1^{(n)} \rho_2^{\otimes n} \right]$ *separately*.

Going beyond supervised learning problems to unsupervised problems – where we are not provided with training data – these can also be considered in the context of state discrimination problems. For an example, see [107]. The focus is on finding the optimal single-shot protocol, which does not fit our sample complexity paradigm.

It is an interesting general question to consider how learning protocols will change if we focus instead on optimising sample complexity. In the context of supervised problems, this is relevant to small sample learning, when we only have a limited number of training data to learning from. This is important when the available quantum data production is scarce.

C. Foundational quantum information and quantum mechanics

Optimal sample complexity of quantum hypothesis testing could also be relevant to foundational quantum information and quantum mechanics. We have already touched upon works that study the use of nonlinear quantum mechanical systems to solve unstructured search problems. The stronger the nonlinearity, the more efficient the corresponding search algorithm is. However, given that we know that such search problems and their generalisations can be NP-complete and even in $\#P$, it is not likely that this can be solved efficiently even by a quantum computer, in polynomial time. Although quantum mechanics is believed to be fundamentally linear, it is not known if there could be a more fundamental theory that is nonlinear. Thus, bounds from optimal sample complexity could potentially put fundamental bounds on nonlinear quantum mechanics at the foundational level [90].

There are many areas where the minimum error probability in quantum hypothesis testing is relevant to the foundations of quantum information theory and quantum mechanics (see [6] for many examples). Applications range from understanding no-go theorems in the interpretation of quantum states [7] and related to bounding types of ψ -epistemic theories [108], ‘reproducing’ standard quantum mechanics from the larger class of general probabilistic theories [109], the operational meaning of min-entropy [110], security proofs in quantum key distribution [111], the construction of dimension witnesses for quantum states [112], relationship to no-signalling [6] and quantum cloning [113].

It is then interesting to ask whether the optimal sample complexity setting can give rise to intriguing new questions. For example, it is known that in some instances optimal asymptotic cloning is equivalent to optimal state discrimination [113, 114]. Asymptotic cloning refers to the limit where the number of clones tend to infinity. Equivalence is only possible in the asymptotic limit because optimal state discrimination is in general imperfect, which leads to an imperfect cloning procedure. Suppose we fix a degree of imperfection allowed in a cloning process. Then focusing instead on the protocol for optimal sample complexity could provide an alternative imperfect cloning procedure.

VI. SUMMARY AND FUTURE DIRECTIONS

In this paper, we defined the sample complexity of quantum hypothesis testing for the symmetric binary setting, the asymmetric binary setting, and for multiple hypotheses. We gave a characterization of the sample complexity for the two aforementioned binary settings. We found in the symmetric binary setting that it depends logarithmically on the priors and inverse error probability and inversely on the negative logarithm of the fidelity. In the asymmetric binary setting, it depends logarithmically on the inverse type II error probability and inversely on the quantum relative entropy, provided that the type II error probability is sufficiently small. Due to non-commutativity in the quantum case, there is freedom in characterizing the sample complexity of symmetric binary quantum hypothesis testing, and we found that other measures like the Holevo fidelity or the broader family of z -fidelities also characterize it. Furthermore, collective measurement strategies or simpler product measurements followed by classical post-processing achieve the same sample complexity, indicating that strategies with radically different technological requirements achieve the same sample complexity, in contrast to previous findings in the information-theoretic setting. It is also interesting to extend quantum sample complexity bounds to the task of the discrimination of quantum channels.

We also summarized many applications of sample complexity of quantum hypothesis testing, including topics as diverse as quantum algorithms for simulation and unstructured search, quantum learning and classification, and foundations of quantum mechanics. Sample complexity plays an important role in establishing the fundamental limitations of the first two tasks. For simulation and search, the idea is that these algorithms produce distinguishable states as output when run on different input states, so that they serve as state discriminators and are ultimately subject to the limitations on such state discriminators set by quantum hypothesis testing. Throughout Section V, we illustrated a number of questions to motivate future work in this direction.

Note: While finalizing our paper, we noticed the independent and concurrent arXiv post [115], which considers sample complexity of classical hypothesis testing in the symmetric and asymmetric binary settings. We note here that all of our results apply to the classical case by substituting commuting density operators that encode probability distributions along their diagonals.

ACKNOWLEDGMENTS

We thank Clément Canonne for several helpful discussions. We are grateful to Marco Fanizza for notifying us of an error in a previous version of Corollary 10. HCC acknowledges support from National Science and Technology Council under grants 112-2636-E-002-009, 113-2119-M-007-006, 113-2119-M-001-006, NSTC 113-2124-M-002-003, and from Ministry of Education under grants NTU-112V1904-4, NTU-CC-113L891605, and NTU-113L900702. NL acknowledges funding from the Science and Technology Program of Shanghai, China (21JC1402900), the NSFC grant No. 12341104, the Shanghai Jiao Tong University 2030 Initiative, and the Fundamental Research Funds for the Central Universities. TN acknowledges support from the NSF under grant no. 2329662. RS acknowledges funding from the Cambridge Commonwealth, European and International Trust and from the European Research Council (ERC Grant AlgoQIP, Agreement No. 851716). MMW acknowledges support from

the NSF under grants 2329662, 2315398, 1907615 and is especially grateful for an international travel supplement to 1907615. This supplement supported a trip to visit the quantum information group at Cambridge University in October 2022, during which this research project was initiated.

-
- [1] J. Neyman and E. S. Pearson, On the problem of the most efficient tests of statistical hypotheses, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **231**, 289 (1933).
 - [2] G. Casella and R. L. Berger, *Statistical Inference* (Cengage Learning, 2021).
 - [3] C. W. Helstrom, Detection theory and quantum mechanics, *Information and Control* **10**, 254 (1967).
 - [4] C. W. Helstrom, Quantum detection and estimation theory, *Journal of Statistical Physics* **1**, 231 (1969).
 - [5] A. S. Holevo, Statistical decision theory for quantum systems, *Journal of Multivariate Analysis* **3**, 337 (1973).
 - [6] J. Bae and L.-C. Kwek, Quantum state discrimination and its applications, *Journal of Physics A: Mathematical and Theoretical* **48**, 083001 (2015).
 - [7] M. F. Pusey, J. Barrett, and T. Rudolph, On the reality of the quantum state, *Nature Physics* **8**, 475 (2012).
 - [8] F. Hiai and D. Petz, The proper formula for relative entropy and its asymptotics in quantum probability, *Communications in Mathematical Physics* **143**, 99 (1991).
 - [9] T. Ogawa and H. Nagaoka, Strong converse and Stein's lemma in quantum hypothesis testing, *IEEE Transaction on Information Theory* **46**, 2428 (2000).
 - [10] K. M. R. Audenaert, J. Calsamiglia, R. Muñoz-Tapia, E. Bagan, L. Masanes, A. Acin, and F. Verstraete, Discriminating states: The quantum Chernoff bound, *Physical Review Letters* **98**, 160501 (2007).
 - [11] M. Nussbaum and A. Szkoła, The Chernoff lower bound for symmetric quantum hypothesis testing, *Annals of Statistics* **37**, 1040 (2009).
 - [12] H. Nagaoka, The converse part of the theorem for quantum Hoeffding bound (2006), arXiv:quant-ph/0611289.
 - [13] M. Hayashi, Error exponent in asymmetric quantum hypothesis testing and its application to classical-quantum channel coding, *Physical Review A* **76**, 062301 (2007).
 - [14] H. P. Yuen, R. S. Kennedy, and M. Lax, Optimum testing of multiple hypotheses in quantum detection theory, *IEEE Transactions on Information Theory* **21**, 125 (1975).
 - [15] K. Li, Discriminating quantum states: The multiple Chernoff distance, *The Annals of Statistics* **44**, 1661 (2016).
 - [16] M. Tomamichel and M. Hayashi, A hierarchy of information quantities for finite block length analysis of quantum tasks, *IEEE Transactions on Information Theory* **59**, 7693 (2013).
 - [17] K. Li, Second-order asymptotics for quantum hypothesis testing, *The Annals of Statistics* **42**, 171 (2014).
 - [18] N. Datta, Y. Pautrat, and C. Rouzé, Second-order asymptotics for quantum hypothesis testing in settings beyond i.i.d.-quantum lattice systems and more, *Journal of Mathematical Physics* **57**, 062207 (2016).
 - [19] M. Tomamichel and V. Y. F. Tan, Second-order asymptotics for the classical capacity

- of image-additive quantum channels, *Communications in Mathematical Physics* **338**, 103 (2015).
- [20] N. Datta, M. Tomamichel, and M. M. Wilde, On the second-order asymptotics for entanglement-assisted communication, *Quantum Information Processing* **15**, 2569 (2016).
- [21] M. Tomamichel, M. Berta, and J. M. Renes, Quantum coding with finite resources, *Nature Communications* **7**, 11419 (2016).
- [22] H.-C. Cheng, M.-H. Hsieh, and M. Tomamichel, Quantum sphere-packing bounds with polynomial prefactors, *IEEE Transactions on Information Theory* **65**, 2872 (2019), arXiv:1704.05703.
- [23] C. T. Chubb, V. Y. F. Tan, and M. Tomamichel, Moderate deviation analysis for classical communication over quantum channels, *Communications in Mathematical Physics* **355**, 1283 (2017).
- [24] H.-C. Cheng and M.-H. Hsieh, Moderate deviation analysis for classical-quantum channels and quantum hypothesis testing, *IEEE Transactions on Information Theory* **64**, 1385 (2018), arXiv:1701.03195.
- [25] M. M. Wilde, M. Tomamichel, and M. Berta, Converse bounds for private communication over quantum channels, *IEEE Transactions on Information Theory* **63**, 1792 (2017).
- [26] S. Arunachalam and R. De Wolf, Optimal quantum sample complexity of learning algorithms, *The Journal of Machine Learning Research* **19**, 2879 (2018).
- [27] C. L. Canonne, Topics and techniques in distribution testing: A biased but representative sample, *Foundations and Trends in Communications and Information Theory* **19**, 1032 (2022).
- [28] T. Heinosaari and M. Ziman, *The Mathematical Language of Quantum Theory* (Cambridge University Press, 2011).
- [29] A. Uhlmann, The “transition probability” in the state space of a $*$ -algebra, *Reports on Mathematical Physics* **9**, 273 (1976).
- [30] A. S. Holevo, On quasiequivalence of locally normal states, *Theoretical and Mathematical Physics* **13**, 1071 (1972).
- [31] C. W. Helstrom, Minimum mean-squared error of estimates in quantum statistics, *Physics Letters A* **25**, 101 (1967).
- [32] D. Bures, An extension of Kakutani’s theorem on infinite product measures to the tensor product of semifinite ω^* -algebras, *Transactions of the American Mathematical Society* **135**, 199 (1969).
- [33] A. Jenčová, Geodesic distances on density matrices, *Journal of Mathematical Physics* **45**, 1787 (2004).
- [34] S. Luo and Q. Zhang, Informational distance on quantum-state space, *Physical Review A* **69**, 032106 (2004).
- [35] D. Petz, Quasi-entropies for finite quantum systems, *Reports on Mathematical Physics* **23**, 57 (1986).
- [36] M. Müller-Lennert, F. Dupuis, O. Szehr, S. Fehr, and M. Tomamichel, On quantum Rényi entropies: A new generalization and some properties, *Journal of Mathematical Physics* **54**, 122203 (2013).
- [37] M. M. Wilde, A. Winter, and D. Yang, Strong converse for the classical capacity of entanglement-breaking and Hadamard channels via a sandwiched Rényi relative entropy, *Communications in Mathematical Physics* **331**, 593 (2014).
- [38] R. L. Frank and E. H. Lieb, Monotonicity of a relative Rényi entropy, *Journal of Mathematical*

- Physics **54**, 122201 (2013).
- [39] M. M. Wilde, Optimized quantum f -divergences and data processing, *Journal of Physics A: Mathematical and Theoretical* **51**, 374002 (2018).
 - [40] H. Umegaki, Conditional expectation in an operator algebra. IV. Entropy and information, *Kodai Mathematical Seminar Reports* **14**, 59 (1962).
 - [41] G. Lindblad, Completely positive maps and entropy inequalities, *Communications in Mathematical Physics* **40**, 147 (1975).
 - [42] A. Rényi, On measures of entropy and information, *Proc. 4th Berkeley Symp. on Math. Statist. Probability* **1**, 547 (1962).
 - [43] S. Kullback and R. A. Leibler, On information and sufficiency, *The Annals of Mathematical Statistics* **22**, 79 (1951).
 - [44] F. Hiai, Equality cases in matrix norm inequalities of Golden–Thompson type, *Linear and Multilinear Algebra* **36**, 239 (1994).
 - [45] F. Buscemi, Comparison of quantum statistical models: Equivalent conditions for sufficiency, *Communications in Mathematical Physics* **310**, 625 (2012).
 - [46] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory* (John Wiley & Sons, 2004).
 - [47] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).
 - [48] F. A. Potra and S. J. Wright, Interior-point methods, *Journal of Computational and Applied Mathematics* **124**, 281 (2000).
 - [49] S. Arora, E. Hazan, and S. Kale, Fast algorithms for approximate semidefinite programming using the multiplicative weights update method, in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)* (2005) pp. 339–348.
 - [50] S. Arora, E. Hazan, and S. Kale, The multiplicative weights update method: A meta-algorithm and applications, *Theory of Computing* **8**, 121 (2012).
 - [51] Y. T. Lee, A. Sidford, and S. C.-W. Wong, A faster cutting plane method and its implications for combinatorial and convex optimization, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (2015) pp. 1049–1065.
 - [52] O. Fawzi, A. Shayeghi, and H. Ta, A hierarchy of efficient bounds on quantum capacities exploiting symmetry, *IEEE Transactions on Information Theory* **68**, 7346 (2022).
 - [53] B. Bergh, N. Datta, R. Salzmänn, and M. M. Wilde, Parallelization of adaptive quantum channel discrimination in the non-asymptotic regime, *IEEE Transactions on Information Theory* **70**, 2617 (2024).
 - [54] A. Pensia, V. Jog, and P.-L. Loh, Communication-constrained hypothesis testing: Optimality, robustness, and reverse data processing inequalities (2023), arXiv:2206.02765 [math.ST].
 - [55] L. Hardy, Quantum theory from five reasonable axioms (2001), arXiv:quant-ph/0101012.
 - [56] L. Hardy, Why quantum theory?, in *Non-locality and Modality*, edited by T. Placek and J. Butterfield (Springer Netherlands, Dordrecht, 2002) pp. 61–73.
 - [57] S. Kimmel, C. Y.-Y. Lin, G. H. Low, M. Ozols, and T. J. Yoder, Hamiltonian simulation with optimal sample complexity, *npj Quantum Information* **3**, 13 (2017).
 - [58] H. K. Mishra, M. Nussbaum, and M. M. Wilde, On the optimal error exponents for classical and quantum antidistinguishability (2023), arXiv:2309.03723.
 - [59] Z. Bar-Yossef, *The complexity of massive data set computations*, Ph.D. thesis, University of California, Berkeley (2002), <https://www.proquest.com/docview/304791145>.
 - [60] C. L. Canonne, A short note on distinguishing discrete distributions (2017), <https://github.com/ccanonne/probabilitydistributiontoolbox/blob/master/testing.pdf>.

- [61] C. A. Fuchs and C. M. Caves, Mathematical techniques for quantum communication theory, *Open Systems & Information Dynamics* **3**, 345 (1995).
- [62] K. M. R. Audenaert and N. Datta, α - z -Rényi relative entropies, *Journal of Mathematical Physics* **56**, 022202 (2015).
- [63] H. Zhang, From Wigner-Yanase-Dyson conjecture to Carlen-Frank-Lieb conjecture, *Advances in Mathematics* **365**, 107053 (2020), arXiv:1811.01205.
- [64] S. M. Lin and M. Tomamichel, Investigating properties of a family of quantum Rényi divergences, *Quantum Information Processing* **14**, 1501 (2015).
- [65] V. P. Belavkin, Optimal multiple quantum statistical hypothesis testing, *Stochastics* **1**, 315 (1975).
- [66] P. Hausladen and W. K. Wootters, A ‘pretty good’ measurement for distinguishing quantum states, *Journal of Modern Optics* **41**, 2385 (1994).
- [67] M. Mosonyi and T. Ogawa, Strong converse exponent for classical-quantum channel coding, *Communications in Mathematical Physics* **355**, 373 (2017).
- [68] N. Datta, Min- and max-relative entropies and a new entanglement monotone, *IEEE Transactions on Information Theory* **55**, 2816 (2009).
- [69] A. Gilyén, S. Lloyd, I. Marvian, Y. Quek, and M. M. Wilde, Quantum algorithm for Petz recovery channels and pretty good measurements, *Physical Review Letters* **128**, 220502 (2022).
- [70] H. Barnum and E. Knill, Reversing quantum dynamics with near-optimal quantum and classical fidelity, *Journal of Mathematical Physics* **43**, 2097 (2002).
- [71] A. W. Harrow and A. Winter, How many copies are needed for state discrimination?, *IEEE Transactions on Information Theory* **58**, 1 (2012).
- [72] K. M. R. Audenaert and M. Mosonyi, Upper bounds on the error probabilities and asymptotic error exponents in quantum multiple state discrimination, *Journal of Mathematical Physics* **55**, 102201 (2014).
- [73] I. Sason and S. Verdú, Arimoto-Rényi conditional entropy and Bayesian M -ary hypothesis testing, *IEEE Transactions on Information Theory* **64**, 4 (2018).
- [74] S. Jin, N. Liu, and Y. Yu, Quantum simulation of partial differential equations: Applications and detailed analysis, *Physical Review A* **108**, 032603 (2023).
- [75] S. Jin and N. Liu, Analog quantum simulation of partial differential equations (2023), arXiv:2308.00646.
- [76] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [77] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Toward the first quantum simulation with quantum speedup, *Proceedings of the National Academy of Sciences* **115**, 9456 (2018).
- [78] S. Jin, N. Liu, and Y. Yu, Quantum simulation of partial differential equations via schrodin-gerisation (2022), arXiv:2212.13969.
- [79] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019) p. 193–204.
- [80] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, Grand unification of quantum algorithms, *PRX Quantum* **2**, 040203 (2021).
- [81] D. An, J.-P. Liu, D. Wang, and Q. Zhao, A theory of quantum differential equation solvers: limitations and fast-forwarding (2022), arXiv:2211.05246.
- [82] P. Busch, P. Lahti, J.-P. Pellonpää, and K. Ylínen, *Dilation theory*, in *Quantum Measurement* (Springer International Publishing, 2016) pp. 137–162.
- [83] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations,

- Physical Review Letters **103**, 150502 (2009).
- [84] S. Jin and N. Liu, Quantum simulation of discrete linear dynamical systems and simple iterative methods in linear algebra via Schrödingerisation (2023), arXiv:2304.02865.
 - [85] A. Vulpiani, F. Cecconi, and M. Cencini, *Chaos: From Simple Models to Complex Systems*, Series on Advances in Statistical Mechanics, Vol. 17 (World Scientific, 2009).
 - [86] J.-P. Liu, H. Ø. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa, and A. M. Childs, Efficient quantum algorithm for dissipative nonlinear differential equations, Proceedings of the National Academy of Sciences **118**, e2026805118 (2021).
 - [87] H. Abou-Kandil, G. Freiling, V. Ionescu, and G. Jank, *Matrix Riccati Equations in Control and Systems Theory*, Systems & Control: Foundations & Applications (Birkhäuser, 2012).
 - [88] R. Nikoukhah, A. S. Willsky, and B. C. Levy, Kalman filtering and Riccati equations for descriptor systems, IEEE Transactions on Automatic Control **37**, 1325 (1992).
 - [89] B. D. O. Anderson, Riccati equations, network theory and Brune synthesis: Old solutions for contemporary problems, in *Dynamical Systems, Control, Coding, Computer Vision: New Trends, Interfaces, and Interplay*, Vol. 25 (Springer, 1999) pp. 1–25.
 - [90] D. S. Abrams and S. Lloyd, Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and #P problems, Physical Review Letters **81**, 3992 (1998).
 - [91] A. M. Childs and J. Young, Optimal state discrimination and unstructured search in nonlinear quantum mechanics, Physical Review A **93**, 022314 (2016).
 - [92] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (1996) p. 212–219.
 - [93] E. Farhi and S. Gutmann, Analog analogue of a digital quantum computation, Physical Review A **57**, 2403 (1998).
 - [94] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, Quantum fingerprinting, Physical Review Letters **87**, 167902 (2001).
 - [95] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, Nature Physics **10**, 631 (2014).
 - [96] P. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 124–134.
 - [97] D. R. Simon, On the power of quantum computation, SIAM Journal on Computing **26**, 1474 (1997).
 - [98] D. Deutsch and R. Jozsa, Rapid solution of problems by quantum computation, Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**, 553 (1992).
 - [99] L. Bianchi, J. Pereira, and S. Pirandola, Generalization in quantum machine learning: A quantum information standpoint, PRX Quantum **2**, 040321 (2021).
 - [100] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Vol. 2 (Springer, 2009).
 - [101] S. Aaronson, The learnability of quantum states, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **463**, 3089 (2007).
 - [102] H.-C. Cheng, M.-H. Hsieh, and P.-C. Yeh, The learnability of unknown quantum measurements, Quantum Information and Computation **16**, 615 (2016).
 - [103] A. Rocchetto, Stabiliser states are efficiently PAC-learnable, Quantum Information and Computation **18**, 541 (2018).
 - [104] S. Gambs, Quantum classification (2008), arXiv:0809.0444.
 - [105] M. Weber, N. Liu, B. Li, C. Zhang, and Z. Zhao, Optimal provable robustness of quantum

- classification via quantum hypothesis testing, *npj Quantum Information* **7**, 76 (2021).
- [106] F. R. Bach, D. Heckerman, and E. Horvitz, Considering cost asymmetry in learning classifiers, *The Journal of Machine Learning Research* **7**, 1713 (2006).
- [107] G. Sentís, A. Monras, R. Muñoz-Tapia, J. Calsamiglia, and E. Bagan, Unsupervised classification of quantum data, *Physical Review X* **9**, 041029 (2019).
- [108] M. S. Leifer and O. J. E. Maroney, Maximally epistemic interpretations of the quantum state and contextuality, *Physical Review Letters* **110**, 120401 (2013).
- [109] H. Arai and M. Hayashi, Derivation of standard quantum theory via state discrimination (2023), arXiv:2307.11271.
- [110] R. König, R. Renner, and C. Schaffner, The operational meaning of min-and max-entropy, *IEEE Transactions on Information Theory* **55**, 4337 (2009).
- [111] A. Leverrier and P. Grangier, Unconditional security proof of long-distance continuous-variable quantum key distribution with discrete modulation, *Physical Review Letters* **102**, 180504 (2009).
- [112] N. Brunner, M. Navascués, and T. Vértesi, Dimension witnesses and quantum state discrimination, *Physical Review Letters* **110**, 150501 (2013).
- [113] J. Bae and A. Acín, Asymptotic quantum cloning is state estimation, *Physical Review Letters* **97**, 030402 (2006).
- [114] D. Bruß, M. Cinchetti, G. M. D’Ariano, and C. Macchiavello, Phase-covariant quantum cloning, *Physical Review A* **62**, 012302 (2000).
- [115] A. Pensia, V. Jog, and P.-L. Loh, The sample complexity of simple binary hypothesis testing (2024), arXiv:2403.16981.
- [116] H. Araki, On an inequality of Lieb and Thirring, *Letters in Mathematical Physics* **19**, 167 (1990).
- [117] E. H. Lieb and W. E. Thirring, Inequalities for the moments of the eigenvalues of the Schrodinger Hamiltonian and their relation to Sobolev inequalities (Walter de Gruyter GmbH, 1976).
- [118] K. M. R. Audenaert, M. Nussbaum, A. Szkoła, and F. Verstraete, Asymptotic error rates in quantum hypothesis testing, *Communications in Mathematical Physics* **279**, 251 (2008).
- [119] C. Fuchs and J. van de Graaf, Cryptographic distinguishability measures for quantum-mechanical states, *IEEE Transactions on Information Theory* **45**, 1216 (1999).
- [120] K. M. Audenaert, Comparisons between quantum state distinguishability measures, *Quantum Information and Computation* **14**, 31 (2014).
- [121] H.-C. Cheng, Simple and tighter derivation of achievability for classical communication over quantum channels, *PRX Quantum* **4**, 040330 (2023).
- [122] R. Iten, J. M. Renes, and D. Sutter, Pretty good measures in quantum information theory, *IEEE Transactions on Information Theory* **63**, 1270 (2017).
- [123] R. T. Powers and E. Størmer, Free states of the canonical anticommutation relations, *Communications in Mathematical Physics* **16**, 1 (1970).
- [124] L. Zhang and J. Wu, A lower bound of quantum conditional mutual information, *Journal of Physics A: Mathematical and Theoretical* **47**, 415303 (2014).
- [125] P. J. Coles, J. Kaniewski, and S. Wehner, Equivalence of wave-particle duality to entropic uncertainty, *Nature Communications* **5**, 5814 (2014).
- [126] M. M. Wilde, M. Berta, C. Hirche, and E. Kaur, Amortized channel divergence for asymptotic quantum channel discrimination, *Letters in Mathematical Physics* **110**, 2277 (2020).
- [127] Y.-C. Shen, L. Gao, and H.-C. Cheng, Strong converse for privacy amplification against

- quantum side information (2022), arXiv:2202.10263 [quant-ph].
- [128] M. Mosonyi and T. Ogawa, Quantum hypothesis testing and the operational interpretation of the quantum Rényi relative entropies, *Communications in Mathematical Physics* **334**, 1617 (2014).
- [129] S. Khatri and M. M. Wilde, Principles of quantum communication theory: A modern approach (2024), arXiv:2011.04672v2.
- [130] H.-C. Cheng and L. Gao, On strong converse theorems for quantum hypothesis testing and channel coding (2024), arXiv:2403.13584.
- [131] K. M. R. Audenaert, M. Mosonyi, and F. Verstraete, Quantum state discrimination bounds for finite sample size, *Journal of Mathematical Physics* **53**, 122205 (2012).
- [132] B. Litjens, S. Polak, and A. Schrijver, Semidefinite bounds for nonbinary codes based on quadruples, *Designs, Codes and Cryptography* **84**, 87 (2017).
- [133] H. Jiang, T. Kathuria, Y. T. Lee, S. Padmanabhan, and Z. Song, A faster interior point method for semidefinite programming, in *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)* (2020) pp. 910–918.
- [134] M. M. Wilde, O. Landon-Cardinal, and P. Hayden, Towards efficient decoding of classical-quantum polar codes, in *8th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2013)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 22 (2013) pp. 157–177.
- [135] B. Bergh, J. Kochanowski, R. Salzmann, and N. Datta, Infinite dimensional asymmetric quantum channel discrimination (2023), arXiv:2308.12959 [quant-ph].
- [136] M. Tomamichel, R. Colbeck, and R. Renner, A fully quantum asymptotic equipartition property, *IEEE Transactions on Information Theory* **55**, 5840 (2009).
- [137] M. Tomamichel, A framework for non-asymptotic quantum information theory (2012), arXiv:1203.2142 [quant-ph].

Appendix A: Properties of and relations between quantum divergences

Lemma 18. *Let A , B , and C be arbitrary positive semi-definite operators, and let ρ and σ be states. The following hold.*

(i) *Multiplicativity and Additivity: For every $n \in \mathbb{N}$,*

$$F(\rho^{\otimes n}, \sigma^{\otimes n}) = [F(\rho, \sigma)]^n, \quad (\text{A1})$$

$$F_{\text{H}}(\rho^{\otimes n}, \sigma^{\otimes n}) = [F_{\text{H}}(\rho, \sigma)]^n, \quad (\text{A2})$$

$$Q_{\min}(\rho^{\otimes n}, \sigma^{\otimes n}) = [Q_{\min}(\rho, \sigma)]^n, \quad (\text{A3})$$

$$1 - \frac{1}{2} [d_{\text{B}}(\rho^{\otimes n}, \sigma^{\otimes n})]^2 = \left(1 - \frac{1}{2} [d_{\text{B}}(\rho, \sigma)]^2\right)^n, \quad (\text{A4})$$

$$1 - \frac{1}{2} [d_{\text{H}}(\rho^{\otimes n}, \sigma^{\otimes n})]^2 = \left(1 - \frac{1}{2} [d_{\text{H}}(\rho, \sigma)]^2\right)^n, \quad (\text{A5})$$

$$C(\rho^{\otimes n} \parallel \sigma^{\otimes n}) = nC(\rho \parallel \sigma). \quad (\text{A6})$$

(ii) *Relations between distances [116–118]:*

$$F_{\text{H}}(\rho, \sigma) \leq F(\rho, \sigma) \leq Q_{\min}(\rho \parallel \sigma) \leq \sqrt{F_{\text{H}}(\rho, \sigma)}, \quad (\text{A7})$$

$$d_{\text{B}}(\rho, \sigma) \leq d_{\text{H}}(\rho, \sigma) \leq \sqrt{2}d_{\text{B}}(\rho, \sigma), \quad (\text{A8})$$

$$[d_{\text{B}}(\rho, \sigma)]^2 \leq -\ln F(\rho, \sigma) \leq -\ln F_{\text{H}}(\rho, \sigma) \leq 2C(\rho\|\sigma) \leq -2\ln F(\rho, \sigma). \quad (\text{A9})$$

(iii) *Quantum Chernoff bound* [10]:

$$\frac{1}{2}(\text{Tr}[A + B] - \|A - B\|_1) \leq Q_{\min}(A\|B). \quad (\text{A10})$$

(iv) *Generalized Fuchs–van de Graaf inequality* [30, 119, 120]:

$$\text{Tr}[A + B] - 2\text{Tr}[\sqrt{A}\sqrt{B}] \leq \|A - B\|_1 \leq \sqrt{(\text{Tr}[A + B])^2 - 4\|\sqrt{A}\sqrt{B}\|_1^2}. \quad (\text{A11})$$

(v) *A relation for the pretty-good test* [121]:

$$\text{Tr}\left[A(A + B)^{-1/2}B(A + B)^{-1/2}\right] \leq \frac{1}{2}\text{Tr}[A + B] - \frac{1}{2}\|A - B\|_1. \quad (\text{A12})$$

(vi) *Subadditivity* [72]:

$$\|\sqrt{A}\sqrt{B + C}\|_1 \leq \|\sqrt{A}\sqrt{B}\|_1 + \|\sqrt{A}\sqrt{C}\|_1. \quad (\text{A13})$$

Proof. (i): The multiplicativity of F , F_{H} , and Q_{\min} follow by inspection. Eqs. (A4) and (A5) follow from the identities $1 - \frac{1}{2}[d_{\text{B}}(\rho, \sigma)]^2 = \sqrt{F}(\rho, \sigma)$ and $1 - \frac{1}{2}[d_{\text{H}}(\rho, \sigma)]^2 = \sqrt{F_{\text{H}}}(\rho, \sigma)$.

(ii): The first inequality of (A7) follows from the Araki–Lieb–Thirring inequality [116, 117], and the second from a [118, Theorem 6] with $A = \rho$, $B = \sigma$, and $s = 1/2$:

$$\|\sqrt{A}\sqrt{B}\|_1 \leq (\text{Tr}[A^s B^{1-s}])^{1/2} (\text{Tr}[A])^{(1-s)/2} (\text{Tr}[B])^{s/2}, \quad \forall s \in [0, 1]. \quad (\text{A14})$$

The last inequality of (A7) follows from definitions. We note that the relation $F(\rho, \sigma) \leq \sqrt{F_{\text{H}}}(\rho, \sigma)$ was also shown in [122, Eq. (38)].

The first inequality of (A8) follows from (A7), and the second follows from

$$[d_{\text{H}}(\rho, \sigma)]^2 = 2\left(1 - \sqrt{F_{\text{H}}}(\rho, \sigma)\right) \quad (\text{A15})$$

$$\leq 2(1 - F(\rho, \sigma)) \quad (\text{A16})$$

$$= 2\left(1 - \sqrt{F}(\rho, \sigma)\right)\left(1 + \sqrt{F}(\rho, \sigma)\right) \quad (\text{A17})$$

$$\leq 4\left(1 - \sqrt{F}(\rho, \sigma)\right) \quad (\text{A18})$$

$$= 2[d_{\text{B}}(\rho, \sigma)]^2. \quad (\text{A19})$$

The first inequality of (A9) follows from $[d_{\text{B}}(\rho, \sigma)]^2 = 2(1 - \sqrt{F}(\rho, \sigma)) \geq -2\ln\sqrt{F}(\rho, \sigma)$ because of $\ln x \leq x - 1$ for all $x > 0$. The rest of inequalities in (A9) follow from (A7).

(iii): The inequality was proved in [10, Theorem 1].

(iv): The first inequality follows from (A10) by choosing $s = 1/2$ as a feasible solution in the definition of Q_{\min} (see also [123] and [124, Eq. (2.6)]). The second inequality follows from [72, Lemma 2.4] (see also [120, Theorem 5], [125, Supplementary Lemma 3] and [126, Eq. (156)]).

(v): The inequality was proved in [121, Lemma 1] (see also [127, Lemma 3]).

(vi): The inequality was proved in [72, Lemma 4.9].

□

Recalling $\beta_\varepsilon(\rho\|\sigma)$ defined in (23), we have the following bounds relating it to quantum Rényi relative entropies.

Lemma 19 ([128, Eq. (75)], [129, Proposition 7.71], [130, Theorem 1]). *Let ρ and σ be states. For all $\alpha \in (1, \infty)$ and $\varepsilon \in [0, 1)$, the following inequality holds*

$$-\ln \beta_\varepsilon(\rho\|\sigma) \leq \tilde{D}_\alpha(\rho\|\sigma) + \frac{\alpha}{\alpha-1} \ln\left(\frac{1}{1-\varepsilon}\right). \quad (\text{A20})$$

Lemma 20 ([118, Theorem 5], [131, Proposition 3.2], [129, Proposition 7.72], [121, Section IV.1]). *Let ρ and σ be states. For all $\alpha \in (0, 1)$ and $\varepsilon \in (0, 1]$, the following inequality holds*

$$-\ln \beta_\varepsilon(\rho\|\sigma) \geq D_\alpha(\rho\|\sigma) + \frac{\alpha}{\alpha-1} \ln\left(\frac{1}{\varepsilon}\right). \quad (\text{A21})$$

Appendix B: Efficient representation of optimal measurements and efficient calculation of optimal error probabilities in quantum hypothesis testing

In this appendix, we begin by showing that the optimal measurement in symmetric hypothesis testing, i.e., the optimiser in

$$p_e(p, \rho, q, \sigma, n) = \min_{\Lambda^{(n)}} \{p \operatorname{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] + q \operatorname{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] : 0 \leq \Lambda^{(n)} \leq I^{\otimes n}\} \quad (\text{B1})$$

$$= p - \max_{\Lambda^{(n)}} \{\operatorname{Tr}[\Lambda^{(n)}(p\rho^{\otimes n} - q\sigma^{\otimes n})] : 0 \leq \Lambda^{(n)} \leq I^{\otimes n}\} \quad (\text{B2})$$

$$= \frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1), \quad (\text{B3})$$

can be efficiently represented by $\text{poly}(n)$ variables that can be computed in $\text{poly}(n)$ time (see Lemma 21 below). For that we essentially follow argument of [52] and [53, Lemma 11]. After doing so, we then remark that the same reasoning can be applied to asymmetric binary and multiple quantum hypothesis testing (see Remark 7).

Let us first introduce some notation. We denote the symmetric group by \mathfrak{S}_n and denote for every permutation $\pi \in \mathfrak{S}_n$ the unitary $P_{\mathcal{H}}(\pi)$ on $\mathcal{H}^{\otimes n}$ which permutes the n copies of $\mathcal{H}^{\otimes n}$ according to π . For any operator X on $\mathcal{H}^{\otimes n}$ we denote the group average by

$$\bar{X} := \frac{1}{|\mathfrak{S}_n|} \sum_{\pi \in \mathfrak{S}_n} P_{\mathcal{H}}(\pi) X P_{\mathcal{H}}(\pi)^\dagger, \quad (\text{B4})$$

and the set of all permutation invariant operators by

$$\text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) := \{A \in \mathcal{B}(\mathcal{H}^{\otimes n}) | P_{\mathcal{H}}(\pi) A P_{\mathcal{H}}(\pi)^\dagger = A, \quad \forall \pi \in \mathfrak{S}_n\}. \quad (\text{B5})$$

where we denoted the set of bounded operators on $\mathcal{H}^{\otimes n}$ by $\mathcal{B}(\mathcal{H}^{\otimes n})$. From standard representation theory of the symmetric group, we know that

$$\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) \leq (n+1)^{d^2}.$$

See [52, Section 4.1] and [132, Section 2.1]. The main idea is to note that the optimiser $\Lambda_\star^{(n)}$ of (B1) is permutation invariant itself, i.e.

$$\Lambda_\star^{(n)} = \{p\rho^{\otimes n} - q\sigma^{\otimes n}\} \in \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}), \quad (\text{B6})$$

which shows that it can be described by $\text{poly}(n)$ many classical variables. To find these variables efficiently, i.e. in $\text{poly}(n)$ time, we use the techniques of [52, 53] to show that the SDP in (B1) is equivalent to one with $\text{poly}(n)$ many variables and constraints whose coefficients can be found from (B1) in $\text{poly}(n)$ time. This is the content of the following lemma:

Lemma 21. *The SDP*

$$\begin{aligned} & \text{maximize} && \text{Tr}[\Lambda^{(n)} (p\rho^{\otimes n} - q\sigma^{\otimes n})] \\ & \Lambda^{(n)} \in \mathcal{B}(\mathcal{H}^{\otimes n}) && \\ & \text{subject to} && 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{aligned} \quad (\text{B7})$$

can be reduced to an equivalent SDP with $\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) \leq (n+1)^{d^2}$ many variables and $(n+1)^d$ many constraints. Moreover, the coefficients of the reduced SDP can be efficiently computed in $\text{poly}(n)$ time from the original SDP (B7).

Proof. Note that, for any feasible $\Lambda^{(n)}$ in (B7), also $\overline{\Lambda^{(n)}}$ is feasible, as the group average (B4) is a positive map. Furthermore, as by definition we have $p\rho^{\otimes n} - q\sigma^{\otimes n} \in \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$, we see

$$\text{Tr}[\Lambda^{(n)} (p\rho^{\otimes n} - q\sigma^{\otimes n})] = \text{Tr}[\overline{\Lambda^{(n)}} (p\rho^{\otimes n} - q\sigma^{\otimes n})] \quad (\text{B8})$$

and hence the SDP (B7) can equivalently be expressed by simply maximising over the smaller space $\text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})$, i.e. by

$$\begin{aligned} & \text{maximize} && \text{Tr}[\Lambda^{(n)} (p\rho^{\otimes n} - q\sigma^{\otimes n})] \\ & \Lambda^{(n)} \in \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) && \\ & \text{subject to} && 0 \leq \Lambda^{(n)} \leq I^{\otimes n}. \end{aligned} \quad (\text{B9})$$

As pointed out above, this reduces number of variables over which we optimize over to $\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n}) \leq (n+1)^{d^2} = \text{poly}(n)$. However, as in the current form the SDP (B9) is phrased in terms of exponentially large operators, this does not yet show computability in $\text{poly}(n)$ time. To close that remaining gap we use the approach of [52, p. 7353] or [53, Lemma 11]. There, the authors explicitly define an orthogonal basis (with respect to the Hilbert–Schmidt scalar product) denoted by $(C_r^{\mathcal{H}})_{r=1}^{\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})}$. For that particular basis and for all $A \in \mathcal{B}(\mathcal{H})$, the authors provide an explicit formula for the coefficients $(\gamma_r)_{r=1}^{\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})}$ of the expansion

$$A^{\otimes n} = \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})} \gamma_r C_r^{\mathcal{H}} \quad (\text{B10})$$

In particular it is shown that $(\gamma_r)_{r=1}^{\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})}$ can be computed in $\text{poly}(n)$ time. Hence, by linearity also the coefficients $(\kappa_r)_{r=1}^{\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})}$ of the operator $p\rho^{\otimes n} - q\sigma^{\otimes n}$ in this basis, i.e., satisfying

$$p\rho^{\otimes n} - q\sigma^{\otimes n} = \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n}(\mathcal{H}^{\otimes n})} \kappa_r C_r^{\mathcal{H}} \quad (\text{B11})$$

can be computed in $\text{poly}(n)$ time. Furthermore, as explained in [52, 53], the square of the norms of the basis operators, i.e. $\text{Tr}[(C_r^{\mathcal{H}})^\dagger C_r^{\mathcal{H}}]$, can be computed efficiently in $O(d^2)$, meaning constant in n , i.e. $O(1)$, time.

As $(C_r^{\mathcal{H}})_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}}$ is a basis, we can also expand the variable $\Lambda^{(n)} \in \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}$ in the SDP (B9) as

$$\Lambda^{(n)} = \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} x_r C_r^{\mathcal{H}}. \quad (\text{B12})$$

As shown in [52, 53], the constraint

$$0 \leq \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} x_r C_r^{\mathcal{H}} \leq I^{\otimes n} \quad (\text{B13})$$

in (B9) can be checked in $\text{poly}(n)$ time as well: To see this, note that there exists a $*$ -algebra isomorphism from $\text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}$ to block-diagonal matrices

$$\phi_{\mathcal{H}} : \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})} \rightarrow \bigoplus_{i=1}^{t_{\mathcal{H}}} \mathbb{C}^{m_i \times m_i}, \quad (\text{B14})$$

where

$$t_{\mathcal{H}} \leq (n+1)^d \quad \text{and} \quad \sum_{i=1}^{t_{\mathcal{H}}} m_i^2 = \dim(\text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}) \leq (n+1)^{d^2}. \quad (\text{B15})$$

For $A \in \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}$ we have $\phi_{\mathcal{H}}(A) \in \bigoplus_{i=1}^{t_{\mathcal{H}}} \mathbb{C}^{m_i \times m_i}$ and write $[\phi_{\mathcal{H}}(A)]_i$ for the i -th block of $\phi_{\mathcal{H}}(A)$. By [52, Lemma 3.3] we have

$$A \geq 0 \Leftrightarrow \phi_{\mathcal{H}}(A) \geq 0 \Leftrightarrow [\phi_{\mathcal{H}}(A)]_i \geq 0 \quad \forall i \in \{1, \dots, t_{\mathcal{H}}\}. \quad (\text{B16})$$

Hence,

$$0 \leq \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} x_r C_r^{\mathcal{H}} \leq I^{\otimes n} \iff 0 \leq \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} x_r [\phi_{\mathcal{H}}(C_r^{\mathcal{H}})]_i \leq I_{\mathbb{C}^{m_i}} \quad \forall i \in \{1, \dots, t_{\mathcal{H}}\}, \quad (\text{B17})$$

where we denoted the identity on \mathbb{C}^{m_i} by $I_{\mathbb{C}^{m_i}}$. Furthermore, as shown in [52, pp. 7352–735], for every basis element $C_r^{\mathcal{H}}$ the corresponding block diagonal matrix $\phi_{\mathcal{H}}(C_r^{\mathcal{H}})$ can be computed in $\text{poly}(n)$ time. Overall, this shows that the constraint (B17) can be checked in $\text{poly}(n)$ time.

Hence, the above shows that the SDP (B7) can equivalently be written as

$$\begin{aligned} & \underset{(x_r)_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} \in \mathbb{C}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}}}{\text{maximize}} && \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} x_r (\kappa_r)^* \text{Tr}[(C_r^{\mathcal{H}})^\dagger C_r^{\mathcal{H}}] \\ & \text{subject to} && 0 \leq \sum_{r=1}^{\dim \text{End}^{\mathfrak{S}_n(\mathcal{H}^{\otimes n})}} x_r [\phi_{\mathcal{H}}(C_r^{\mathcal{H}})]_i \leq I_{\mathbb{C}^{m_i}}, \\ & && \forall i \in \{1, \dots, t_{\mathcal{H}}\} \end{aligned} \quad (\text{B18})$$

where all coefficients can be found in $\text{poly}(n)$ from the original SDP (B7). As the SDP (B18) has $\text{poly}(n)$ variables and $\text{poly}(n)$ constraints, it can be solved with an additive error ε in $O(\text{poly}(n) \log(1/\varepsilon))$ time; see, e.g., [133]. \square

Remark 7. *By exact same line of argument as in the proof Lemma 21, we can also reduce the SDP*

$$\begin{aligned} & \underset{\Lambda^{(n)} \in \mathcal{B}(\mathcal{H}^{\otimes n})}{\text{minimize}} && \text{Tr}[\Lambda^{(n)} \sigma^{\otimes n}] \\ & \text{subject to} && 0 \leq \Lambda^{(n)} \leq I^{\otimes n}, \\ & && \text{Tr}[\Lambda^{(n)} \rho^{\otimes n}] \geq 1 - \varepsilon, \end{aligned} \tag{B19}$$

which was defined in (23) in the context of asymmetric binary hypothesis testing, and

$$\begin{aligned} & \underset{\Lambda_1^{(n)}, \dots, \Lambda_m^{(n)} \in \mathcal{B}(\mathcal{H}^{\otimes n})}{\text{minimize}} && \sum_{m=1}^M p_m \text{Tr}[(I^{\otimes n} - \Lambda^{(n)}) \rho_m^{\otimes n}] \\ & \text{subject to} && \Lambda_m^{(n)} \geq 0, \\ & && \sum_{m=1}^M \Lambda_m^{(n)} = I^{\otimes n}, \end{aligned} \tag{B20}$$

which was defined in (24) in the context of M -ary hypothesis testing to equivalent SDPs with $\text{poly}(n)$ classical variables and $\text{poly}(n)$ constraints. Moreover, the coefficients of the reduced SDPs can again be efficiently computed in $\text{poly}(n)$ time from the original SDPs.

This shows how also the optimal measurements of asymmetric binary hypothesis testing as well as the optimal measurements in M -ary hypothesis testing can be written in terms of $\text{poly}(n)$ many classical variables that can be found efficiently in $\text{poly}(n)$ time.

Appendix C: Proof of Equations (31)–(32)

To prove Equations (31)–(32), we use that the optimal error in asymmetric binary quantum hypothesis testing is attained. In finite dimensions this follows simply by compactness. In infinite dimensions the same proof goes through by noticing that we still have the right notion of compactness, i.e. weak-* compactness, due to the Banach–Alaoglu theorem. We first formally state and prove this fact in the following lemma and then continue to give the proof of Equations (31)–(32).

Lemma 22. *Let \mathcal{H} be a separable Hilbert space, and let ρ and σ be quantum states on \mathcal{H} . Fix $\varepsilon \geq 0$. Then the infimum in the optimal error in asymmetric binary quantum hypothesis testing is attained, i.e.*

$$\beta_\varepsilon(\rho \parallel \sigma) := \inf_{\substack{0 \leq \Lambda \leq I \\ \text{Tr}((I - \Lambda)\rho) \leq \varepsilon}} \text{Tr}(\Lambda\sigma) = \min_{\substack{0 \leq \Lambda \leq I \\ \text{Tr}((I - \Lambda)\rho) \leq \varepsilon}} \text{Tr}(\Lambda\sigma).$$

Proof. By the definition of $\beta_\varepsilon(\rho \parallel \sigma)$ there exists a sequence of bounded operators $(\Lambda_k)_{k \in \mathbb{N}}$ such that $0 \leq \Lambda_k \leq I$ and $\text{Tr}((I - \Lambda_k)\rho) \leq \varepsilon$ for all $k \in \mathbb{N}$ and

$$\lim_{k \rightarrow \infty} \text{Tr}(\Lambda_k \sigma) = \beta_\varepsilon(\rho \parallel \sigma). \tag{C1}$$

By the Banach–Alaoglu theorem there exists a subsequence $(\Lambda_{k_l})_{l \in \mathbb{N}}$ that converges weakly-^{*} to some bounded operator Λ ; i.e., using that the dual of the trace class operators is the space of bounded operators we have for all trace class operators T that

$$\lim_{l \rightarrow \infty} \text{Tr}(\Lambda_{k_l} T) = \text{Tr}(\Lambda T). \quad (\text{C2})$$

From that we immediately see that $0 \leq \Lambda \leq I$ and $\text{Tr}((I - \Lambda)\rho) \leq \varepsilon$ and furthermore by (C1) also

$$\text{Tr}(\Lambda\sigma) = \lim_{l \rightarrow \infty} \text{Tr}(\Lambda_{k_l}\sigma) = \beta_\varepsilon(\rho\|\sigma),$$

which finishes the proof. \square

Proof of Equations (31)–(32). Let us define

$$\beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) := \inf_{\Lambda^{(n)}} \left\{ \begin{array}{l} \text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] : \text{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] \leq \varepsilon, \\ 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{array} \right\}. \quad (\text{C3})$$

Our goal here is to prove that

$$\inf \{n \in \mathbb{N} : \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq \delta\} = \inf \left\{ \begin{array}{l} n \in \mathbb{N} : \text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] \leq \delta, \\ \text{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] \leq \varepsilon, \\ 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{array} \right\}. \quad (\text{C4})$$

Let n be a feasible choice for the optimization on the left-hand side, meaning that for this choice of n , an optimal choice $\Lambda^{(n)}$ (which exists due to Lemma 22) for the optimization $\beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n})$ satisfies $\text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] = \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq \delta$. Then, by the definition in (C3) and the constraint $\text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] \leq \delta$, this means that n and $\Lambda^{(n)}$ are feasible for the optimization on the right-hand side. So we conclude that

$$n \geq \inf \left\{ \begin{array}{l} n \in \mathbb{N} : \text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] \leq \delta, \\ \text{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] \leq \varepsilon, \\ 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{array} \right\}. \quad (\text{C5})$$

Since this argument holds for every feasible choice of n on the left-hand side, we conclude that

$$\inf \{n \in \mathbb{N} : \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq \delta\} \geq \inf \left\{ \begin{array}{l} n \in \mathbb{N} : \text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] \leq \delta, \\ \text{Tr}[(I^{\otimes n} - \Lambda^{(n)})\rho^{\otimes n}] \leq \varepsilon, \\ 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{array} \right\}. \quad (\text{C6})$$

Now let n and $\Lambda^{(n)}$ be feasible choices for the optimization on the right-hand side. Then it follows that $\Lambda^{(n)}$ is feasible for $\beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n})$. Since $\text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] \leq \delta$, we conclude that

$$\delta \geq \text{Tr}[\Lambda^{(n)}\sigma^{\otimes n}] \geq \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}). \quad (\text{C7})$$

Thus, the choice of n is feasible for the optimization on the left-hand side. We then conclude that

$$\inf \{n \in \mathbb{N} : \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq \delta\} \leq n, \quad (\text{C8})$$

and since this inequality holds for every feasible choice of n on the right-hand side, we conclude the opposite inequality:

$$\inf \{n \in \mathbb{N} : \beta_\varepsilon(\rho^{\otimes n} \|\sigma^{\otimes n}) \leq \delta\} \leq \inf \left\{ \begin{array}{l} n \in \mathbb{N} : \text{Tr}[\Lambda^{(n)} \sigma^{\otimes n}] \leq \delta, \\ \text{Tr}[(I^{\otimes n} - \Lambda^{(n)}) \rho^{\otimes n}] \leq \varepsilon, \\ 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{array} \right\}. \quad (\text{C9})$$

By similar reasoning, we conclude that

$$\inf \{n \in \mathbb{N} : \beta_\delta(\sigma^{\otimes n} \|\rho^{\otimes n}) \leq \varepsilon\} = \inf \left\{ \begin{array}{l} n \in \mathbb{N} : \text{Tr}[\Lambda^{(n)} \sigma^{\otimes n}] \leq \delta, \\ \text{Tr}[(I^{\otimes n} - \Lambda^{(n)}) \rho^{\otimes n}] \leq \varepsilon, \\ 0 \leq \Lambda^{(n)} \leq I^{\otimes n} \end{array} \right\}. \quad (\text{C10})$$

□

Appendix D: Proof of Remark 2

We begin by establishing (34) under the assumption that $\varepsilon \in [1/2, 1]$. The trivial strategy of guessing the state uniformly at random achieves an error probability of $1/2$ (i.e., with $n = 1$ and $\Lambda^{(1)} = I/2$ in (20)). Thus, if $\varepsilon \in [1/2, 1]$, then the trivial strategy accomplishes the task with just a single sample. Alternatively, by considering the expression in (30) and noting that the norm is always non-negative, if $\varepsilon \in [1/2, 1]$, then the left-hand side of the inequality is ≤ 0 while the right-hand side is ≥ 0 , so that the constraint is trivially met and the smallest possible feasible $n \in \mathbb{N}$ is $n = 1$.

If $\rho\sigma = 0$, then we set the test $\Lambda^{(1)}$ to be the projection onto the support of ρ . In this case, $p_e = 0$ and hence $n = 1$ suffices.

Now we establish (34) under the assumption that there exists $s \in [0, 1]$ such that $\varepsilon \geq p^s q^{1-s}$. With a single sample, we can achieve the error probability $\frac{1}{2}(1 - \|p\rho - q\sigma\|_1)$, and by invoking Lemma 18-(iii) with $A = p\rho$ and $B = q\sigma$, we conclude that

$$\frac{1}{2}(1 - \|p\rho - q\sigma\|_1) \leq p^s q^{1-s} \text{Tr}[\rho^s \sigma^{1-s}] \quad (\text{D1})$$

$$\leq p^s q^{1-s} \quad (\text{D2})$$

$$\leq \varepsilon. \quad (\text{D3})$$

In the above, we used the Hölder inequality $|\text{Tr}[AB]| \leq \|A\|_r \|B\|_t$, which holds for $r, t \geq 1$, such that $r^{-1} + t^{-1} = 1$, to conclude that

$$\text{Tr}[\rho^s \sigma^{1-s}] \leq 1. \quad (\text{D4})$$

That is, set $r = s^{-1}, t = (1 - s)^{-1}$, $A = \rho$, and $B = \sigma$. The last inequality follows by assumption. So we conclude that only a single sample is needed in this case also.

We finally prove (35), which is the statement that if $\rho = \sigma$ and $\min\{p, q\} > \varepsilon \in [0, 1/2]$, then $n^*(p, \rho, q, \sigma, \varepsilon) = +\infty$. This follows because it is impossible to distinguish the states and the desired inequality $\frac{1}{2}(1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1) \leq \varepsilon$ cannot be satisfied for any possible value of n . Indeed, in this case,

$$\frac{1}{2}(1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1) = \frac{1}{2}(1 - |p - q| \|\rho^{\otimes n}\|_1) \quad (\text{D5})$$

$$= \frac{1}{2} (1 - |p - (1 - p)|) \quad (\text{D6})$$

$$= \frac{1}{2} (1 - |1 - 2p|). \quad (\text{D7})$$

When $p \leq 1/2$, we have that $\frac{1}{2}(1 - |1 - 2p|) = \frac{1}{2}(1 - (1 - 2p)) = p$, while if $p > 1/2$, we have that $\frac{1}{2}(1 - |1 - 2p|) = \frac{1}{2}(1 - (2p - 1)) = 1 - p = q$. Thus,

$$\frac{1}{2}(1 - |1 - 2p|) = \min\{p, q\}, \quad (\text{D8})$$

and so

$$\frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1) = \min\{p, q\} \quad (\text{D9})$$

in this case. Thus, (35) follows.

Appendix E: Proof of Theorem 6

First recall the identity in (36). Applying this to our case of interest, and setting $|\varphi\rangle = \sqrt{p}|\psi\rangle^{\otimes n}$ and $|\zeta\rangle = \sqrt{q}|\phi\rangle^{\otimes n}$, we find that

$$1 - 2\varepsilon \leq \|p\psi^{\otimes n} - q\phi^{\otimes n}\|_1 \quad (\text{E1})$$

$$= \sqrt{(p+q)^2 - 4pq|\langle\psi|\phi\rangle|^{2n}} \quad (\text{E2})$$

$$= \sqrt{1 - 4pq|\langle\psi|\phi\rangle|^{2n}} \quad (\text{E3})$$

$$= \sqrt{1 - 4pqF(\psi, \phi)^n}, \quad (\text{E4})$$

which we can rewrite as

$$(1 - 2\varepsilon)^2 \leq 1 - 4pqF(\psi, \phi)^n$$

$$\Leftrightarrow F(\psi, \phi)^n \leq \frac{1 - (1 - 2\varepsilon)^2}{4pq} = \frac{4\varepsilon(1 - \varepsilon)}{4pq} \quad (\text{E5})$$

$$\Leftrightarrow n \ln F(\psi, \phi) \leq \ln\left(\frac{\varepsilon(1 - \varepsilon)}{pq}\right) \quad (\text{E6})$$

$$\Leftrightarrow -n \ln F(\psi, \phi) \geq -\ln\left(\frac{\varepsilon(1 - \varepsilon)}{pq}\right) \quad (\text{E7})$$

$$\Leftrightarrow n \geq \frac{\ln\left(\frac{pq}{\varepsilon(1 - \varepsilon)}\right)}{-\ln F(\psi, \phi)}. \quad (\text{E8})$$

This finally implies that

$$n = \left\lceil \frac{\ln\left(\frac{pq}{\varepsilon(1 - \varepsilon)}\right)}{-\ln F(\psi, \phi)} \right\rceil, \quad (\text{E9})$$

concluding the proof.

Appendix F: Proof of Theorem 7

We begin by establishing the first inequality in (38). That is, we first show that

$$n^*(p, \rho, q, \sigma, \varepsilon) \geq \frac{\ln(pq/\varepsilon)}{-\ln F(\rho, \sigma)}. \quad (\text{F1})$$

Let us rewrite the constraint in the definition of sample complexity as in (29):

$$\varepsilon \geq \frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1). \quad (\text{F2})$$

Applying the right-most inequality in Lemma 18-(iv) with $A = p\rho$ and $B = q\sigma$ leads to $\|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1^2 \leq 1 - 4pqF(\rho^{\otimes n}, \sigma^{\otimes n})$, which in turn leads to the following lower bound:

$$\varepsilon \geq \frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1) \quad (\text{F3})$$

$$\geq \frac{2pqF(\rho^{\otimes n}, \sigma^{\otimes n})}{1 + \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1} \quad (\text{F4})$$

$$\geq pqF(\rho^{\otimes n}, \sigma^{\otimes n}) \quad (\text{F5})$$

$$= pq [F(\rho, \sigma)]^n \quad (\text{F6})$$

This arrives at the desired (F1). Here we also used the fact that $\|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1 \leq \|p\rho^{\otimes n}\|_1 + \|q\sigma^{\otimes n}\|_1 = p + q = 1$.

Next we prove the following inequality stated in (38):

$$n^*(p, \rho, q, \sigma, \varepsilon) \geq \frac{pq - \varepsilon(1 - \varepsilon)}{pq [d_B(\rho, \sigma)]^2}. \quad (\text{F7})$$

Let us rewrite the constraint in the definition of sample complexity as in (30):

$$1 - 2\varepsilon \leq \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1. \quad (\text{F8})$$

Applying Lemma 18-(iv) with $A = p\rho$ and $B = q\sigma$, consider that

$$\|p\rho - q\sigma\|_1^2 \leq 1 - 4pqF(\rho, \sigma) \quad (\text{F9})$$

$$= 1 - 4pq + 4pq - 4pqF(\rho, \sigma) \quad (\text{F10})$$

$$= 1 - 4pq + 4pq(1 - F(\rho, \sigma)) \quad (\text{F11})$$

$$= 1 - 4pq + 4pq \left(1 + \sqrt{F(\rho, \sigma)}\right) \left(1 - \sqrt{F(\rho, \sigma)}\right) \quad (\text{F12})$$

$$\leq 1 - 4pq + 4pq [d_B(\rho, \sigma)]^2. \quad (\text{F13})$$

Applying this, we find that

$$(1 - 2\varepsilon)^2 \leq \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1^2 \quad (\text{F14})$$

$$\leq 1 - 4pq + 4pq [d_B(\rho^{\otimes n}, \sigma^{\otimes n})]^2 \quad (\text{F15})$$

$$= 1 - 4pq + 8pq \frac{[d_B(\rho^{\otimes n}, \sigma^{\otimes n})]^2}{2} \quad (\text{F16})$$

$$= 1 - 4pq + 8pq \left(1 - \left(1 - \frac{[d_{\text{B}}(\rho^{\otimes n}, \sigma^{\otimes n})]^2}{2} \right) \right) \quad (\text{F17})$$

$$= 1 - 4pq + 8pq \left(1 - \left(1 - \frac{[d_{\text{B}}(\rho, \sigma)]^2}{2} \right)^n \right) \quad (\text{F18})$$

$$\leq 1 - 4pq + 8pq n \left(\frac{[d_{\text{B}}(\rho, \sigma)]^2}{2} \right) \quad (\text{F19})$$

$$= 1 - 4pq + 4pq n [d_{\text{B}}(\rho, \sigma)]^2. \quad (\text{F20})$$

This implies that

$$n \geq \frac{(1 - 2\varepsilon)^2 - (1 - 4pq)}{4pq [d_{\text{B}}(\rho, \sigma)]^2} \quad (\text{F21})$$

$$= \frac{4(pq - \varepsilon(1 - \varepsilon))}{4pq [d_{\text{B}}(\rho, \sigma)]^2}, \quad (\text{F22})$$

thus establishing (F7).

Let us now establish the rightmost inequality in (38). Set

$$n := \left\lceil \inf_{s \in [0,1]} \frac{\ln \left(\frac{p^s q^{1-s}}{\varepsilon} \right)}{-\ln \text{Tr}[\rho^s \sigma^{1-s}]} \right\rceil, \quad (\text{F23})$$

and let $s^* \in [0, 1]$ be the optimal value of s above. Noting that $-\ln \text{Tr}[\rho^s \sigma^{1-s}] \geq 0$ due to (D4), then

$$n \geq \frac{\ln \left(\frac{p^{s^*} q^{1-s^*}}{\varepsilon} \right)}{-\ln \text{Tr}[\rho^{s^*} \sigma^{1-s^*}]} \quad (\text{F24})$$

$$\Leftrightarrow -n \ln \text{Tr}[\rho^{s^*} \sigma^{1-s^*}] \geq \ln(p^{s^*} q^{1-s^*}) - \ln \varepsilon \quad (\text{F25})$$

$$\Leftrightarrow -\ln \text{Tr} \left[(\rho^{\otimes n})^{s^*} (\sigma^{\otimes n})^{1-s^*} \right] - \ln(p^{s^*} q^{1-s^*}) \geq -\ln \varepsilon \quad (\text{F26})$$

$$\Leftrightarrow \text{Tr} \left[(p\rho^{\otimes n})^{s^*} (q\sigma^{\otimes n})^{1-s^*} \right] \leq \varepsilon. \quad (\text{F27})$$

Now applying Lemma 18-(iii) with $A = p\rho^{\otimes n}$ and $B = q\sigma^{\otimes n}$, we conclude that

$$\varepsilon \geq \text{Tr} \left[(p\rho^{\otimes n})^{s^*} (q\sigma^{\otimes n})^{1-s^*} \right] \quad (\text{F28})$$

$$\geq \frac{1}{2} (1 - \|p\rho^{\otimes n} - q\sigma^{\otimes n}\|_1). \quad (\text{F29})$$

The choice of n in (F23) thus satisfies the constraint (29) in Definition 3, and since the optimal sample complexity cannot exceed this choice, this concludes our proof of the rightmost inequality in (38).

Appendix G: Fuchs–Caves measurement strategy for achieving the asymptotic sample complexity of symmetric binary quantum hypothesis testing

In this appendix, we show how the upper bound on sample complexity in (40) can be achieved by repeatedly applying the Fuchs–Caves measurement [61] along with classical post-processing. The development here generalizes that in [134, Lemma 2]. For simplicity, here we focus on the case in which σ is a positive definite state. Let us first recall the operator geometric mean of two positive definite operators A and B :

$$A\#B := B^{1/2} (B^{-1/2}AB^{-1/2})^{1/2} B^{1/2}, \quad (\text{G1})$$

which reduces to the usual geometric mean in the case that A and B are scalars. It is defined as $\lim_{\varepsilon \rightarrow 0^+} (A + \varepsilon I)\#(B + \varepsilon I)$ in the more general case that A and B are positive semi-definite. The Fuchs–Caves observable for distinguishing $\rho^{\otimes n}$ from $\sigma^{\otimes n}$, with respective priors p and q , is then $p\rho^{\otimes n}\#(q\sigma^{\otimes n})^{-1}$. Observe that

$$p\rho^{\otimes n}\#(q\sigma^{\otimes n})^{-1} = \sqrt{\frac{p}{q}} (\rho\#\sigma^{-1})^{\otimes n}, \quad (\text{G2})$$

which implies that a diagonalizing basis for the Fuchs–Caves observable can be realized as a tensor product of the diagonalizing basis for $\rho\#\sigma^{-1}$. That is, let us suppose that $\rho\#\sigma^{-1}$ has the following spectral decomposition:

$$\rho\#\sigma^{-1} = \sum_y \lambda_y |y\rangle\langle y|. \quad (\text{G3})$$

It was observed in [61] that the eigenvalues of $\rho\#\sigma^{-1}$ take the following form:

$$\lambda_y = \left(\frac{\langle y|\rho|y\rangle}{\langle y|\sigma|y\rangle} \right)^{1/2}. \quad (\text{G4})$$

The so-called Fuchs–Caves measurement for distinguishing between the states ρ and σ is a projective measurement defined by the set of rank-one projections $\{|y\rangle\langle y|\}$. Performing this measurement leads to the outcome probability distributions $p := \{p(y)\}$, with $p(y) = \langle y|\rho|y\rangle$, and $q := \{q(y)\}$, with $q(y) = \langle y|\sigma|y\rangle$. It is known that the fidelity $F(\rho, \sigma)$, which is a measure of distinguishability between the two states ρ and σ , is exactly equal to the classical fidelity $F(p, q)$ of the distributions p and q . In other words, the quantum fidelity is achieved by the Fuchs–Caves measurement:

$$F(\rho, \sigma) = \left(\sum_y \sqrt{\langle y|\rho|y\rangle\langle y|\sigma|y\rangle} \right)^2. \quad (\text{G5})$$

Then applying (G2) and (G3), we find that

$$p\rho^{\otimes n}\#(q\sigma^{\otimes n})^{-1} = \sqrt{\frac{p}{q}} \sum_{y^n} \lambda_{y^n} |y^n\rangle\langle y^n|, \quad (\text{G6})$$

where

$$\lambda_{y^n} := \lambda_{y_1} \times \lambda_{y_2} \times \cdots \times \lambda_{y_n}, \quad (\text{G7})$$

$$|y^n\rangle := |y_1\rangle \otimes |y_2\rangle \otimes \cdots \otimes |y_n\rangle. \quad (\text{G8})$$

Furthermore, observe that

$$[F(\rho, \sigma)]^n = F(\rho^{\otimes n}, \sigma^{\otimes n}) = \left(\sum_{y^n} \sqrt{\langle y^n | \rho^{\otimes n} | y^n \rangle \langle y^n | \sigma^{\otimes n} | y^n \rangle} \right)^2. \quad (\text{G9})$$

The idea of the Fuchs–Caves measurement in this case is to perform the measurement $\{|y\rangle\langle y|\}_y$ on each individual system, leading to a measurement outcome sequence y^n , and then decide ρ if $\lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1$ and σ if $\lambda_{y^n} \sqrt{\frac{p}{q}} < 1$. In this sense, the test being performed is a “quantum likelihood ratio test,” as discussed in [61]. So indeed the strategy consists of a product measurement followed by classical post-processing. The error probability of this measurement strategy is then given by:

$$p \operatorname{Tr}[\Lambda_\sigma^{(n)} \rho^{\otimes n}] + q \operatorname{Tr}[\Lambda_\rho^{(n)} \sigma^{\otimes n}], \quad (\text{G10})$$

where

$$\Lambda_\rho^{(n)} := \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1} |y^n\rangle\langle y^n|, \quad (\text{G11})$$

$$\Lambda_\sigma^{(n)} := \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} < 1} |y^n\rangle\langle y^n|. \quad (\text{G12})$$

Now generalizing the error analysis in the proof of [134, Lemma 2] and observing from (G2), (G3), and (G4) that

$$\lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1 \quad \Leftrightarrow \quad \left(p \prod_{i=1}^n \langle y_i | \rho | y_i \rangle \right)^{1/2} \geq \left(q \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \right)^{1/2}, \quad (\text{G13})$$

we find the following upper bound on the error probability of this Fuchs–Caves strategy:

$$\begin{aligned} & p \operatorname{Tr}[\Lambda_\sigma^{(n)} \rho^{\otimes n}] + q \operatorname{Tr}[\Lambda_\rho^{(n)} \sigma^{\otimes n}] \\ &= p \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} < 1} \langle y^n | \rho^{\otimes n} | y^n \rangle + q \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1} \langle y^n | \sigma^{\otimes n} | y^n \rangle \end{aligned} \quad (\text{G14})$$

$$= p \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} < 1} \prod_{i=1}^n \langle y_i | \rho | y_i \rangle + q \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1} \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \quad (\text{G15})$$

$$\begin{aligned} &= \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} < 1} \left(p \prod_{i=1}^n \langle y_i | \rho | y_i \rangle \right)^{1/2} \left(p \prod_{i=1}^n \langle y_i | \rho | y_i \rangle \right)^{1/2} \\ &\quad + \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1} \left(q \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \right)^{1/2} \left(q \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \right)^{1/2} \end{aligned} \quad (\text{G16})$$

$$\begin{aligned} &\leq \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} < 1} \left(p \prod_{i=1}^n \langle y_i | \rho | y_i \rangle \right)^{1/2} \left(q \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \right)^{1/2} \\ &\quad + \sum_{y^n: \lambda_{y^n} \sqrt{\frac{p}{q}} \geq 1} \left(p \prod_{i=1}^n \langle y_i | \rho | y_i \rangle \right)^{1/2} \left(q \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \right)^{1/2} \end{aligned} \quad (\text{G17})$$

$$= \sum_{y^n} \left(p \prod_{i=1}^n \langle y_i | \rho | y_i \rangle \right)^{1/2} \left(q \prod_{i=1}^n \langle y_i | \sigma | y_i \rangle \right)^{1/2} \quad (\text{G18})$$

$$= \sqrt{pq} \sum_{y^n} \sqrt{\langle y^n | \rho^{\otimes n} | y^n \rangle \langle y^n | \sigma^{\otimes n} | y^n \rangle} \quad (\text{G19})$$

$$= \sqrt{pq} F(\rho, \sigma)^{n/2}. \quad (\text{G20})$$

Turning this bound around for the purposes of sample complexity, by setting

$$n := \left\lceil \frac{\ln\left(\frac{\sqrt{pq}}{\varepsilon}\right)}{-\frac{1}{2} \ln F(\rho, \sigma)} \right\rceil, \quad (\text{G21})$$

where ε is the desired threshold in (25), we find that

$$n \geq \frac{\ln\left(\frac{\sqrt{pq}}{\varepsilon}\right)}{-\frac{1}{2} \ln F(\rho, \sigma)} \quad (\text{G22})$$

$$\Leftrightarrow n \ln\left(\frac{1}{\sqrt{F(\rho, \sigma)}}\right) \geq \ln\left(\frac{\sqrt{pq}}{\varepsilon}\right) \quad (\text{G23})$$

$$\Leftrightarrow \frac{1}{\left[\sqrt{F(\rho, \sigma)}\right]^n} \geq \frac{\sqrt{pq}}{\varepsilon} \quad (\text{G24})$$

$$\Leftrightarrow \varepsilon \geq \sqrt{pq} \left[\sqrt{F(\rho, \sigma)}\right]^n, \quad (\text{G25})$$

which, when combined with (G14)–(G20), establishes that the Fuchs–Caves measurement strategy achieves the sample complexity upper bound in (40).

Appendix H: Proof of Theorem 9

Let us begin by proving the lower bound in (44). Here we assume the existence of a scheme such that the constraint $\beta_\varepsilon(\rho^{\otimes n} \| \sigma^{\otimes n}) \leq \delta$ holds. Then we apply Lemma 19 to conclude that the following holds for all $\alpha \in (1, \gamma]$:

$$\ln\left(\frac{1}{\delta}\right) \leq -\ln \beta_\varepsilon(\rho^{\otimes n} \| \sigma^{\otimes n}) \quad (\text{H1})$$

$$\leq \tilde{D}_\alpha(\rho^{\otimes n} \| \sigma^{\otimes n}) + \frac{\alpha}{\alpha - 1} \ln\left(\frac{1}{1 - \varepsilon}\right) \quad (\text{H2})$$

$$= n\tilde{D}_\alpha(\rho\|\sigma) + \frac{\alpha}{\alpha-1} \ln\left(\frac{1}{1-\varepsilon}\right). \quad (\text{H3})$$

Rewriting this gives the lower bound:

$$n^*(\rho, \sigma, \varepsilon, \delta) \geq \sup_{\alpha \in (1, \gamma]} \left(\frac{\ln\left(\frac{(1-\varepsilon)^{\alpha'}}{\delta}\right)}{\tilde{D}_\alpha(\rho\|\sigma)} \right). \quad (\text{H4})$$

Now applying the same reasoning, but instead applying the expression in (32), we conclude the lower bound

$$n^*(\rho, \sigma, \varepsilon, \delta) \geq \sup_{\alpha \in (1, \gamma]} \left(\frac{\ln\left(\frac{(1-\delta)^{\alpha'}}{\varepsilon}\right)}{\tilde{D}_\alpha(\sigma\|\rho)} \right), \quad (\text{H5})$$

thus finishing the proof of the first inequality in (44).

Now we prove the upper bound in (44). For every $n \in \mathbb{N}$, all $\alpha \in (0, 1)$, and all $\varepsilon \in (0, 1]$, we apply Lemma 20 and additivity of the Petz–Rényi relative entropy to obtain

$$-\ln \beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \geq D_\alpha(\rho^{\otimes n}\|\sigma^{\otimes n}) + \alpha' \ln\left(\frac{1}{\varepsilon}\right) \quad (\text{H6})$$

$$= nD_\alpha(\rho\|\sigma) + \alpha' \ln\left(\frac{1}{\varepsilon}\right). \quad (\text{H7})$$

Note here that $D_\alpha(\rho\|\sigma)$ is strictly positive and finite for every $\alpha \in (0, 1)$, provided that $\rho \not\ll \sigma$. Now setting

$$n := \left\lceil \frac{\ln\left(\frac{\varepsilon^{\alpha'}}{\delta}\right)}{D_\alpha(\rho\|\sigma)} \right\rceil, \quad (\text{H8})$$

we conclude that

$$n \geq \frac{\ln\left(\frac{\varepsilon^{\alpha'}}{\delta}\right)}{D_\alpha(\rho\|\sigma)}, \quad (\text{H9})$$

which can be rewritten as

$$nD_\alpha(\rho\|\sigma) + \alpha' \ln\left(\frac{1}{\varepsilon}\right) \geq \ln\left(\frac{1}{\delta}\right). \quad (\text{H10})$$

By applying (H6)–(H7) and (H10), we then conclude that

$$\beta_\varepsilon(\rho^{\otimes n}\|\sigma^{\otimes n}) \leq \delta, \quad (\text{H11})$$

which means that the choice of n in (H8) leads to the desired constraint in (26) being satisfied. Since the sample complexity is the minimum value of $n \in \mathbb{N}$ such that (26) holds, we thus conclude the following upper bound on sample complexity:

$$n^*(\rho, \sigma, \varepsilon, \delta) \leq \left\lceil \frac{\ln\left(\frac{\varepsilon^{\alpha'}}{\delta}\right)}{D_\alpha(\rho\|\sigma)} \right\rceil. \quad (\text{H12})$$

Since the above bound holds for every $\alpha \in (0, 1)$, we conclude

$$n^*(\rho, \sigma, \varepsilon, \delta) \leq \inf_{\alpha \in (0,1)} \left[\frac{\ln\left(\frac{\varepsilon^{\alpha'}}{\delta}\right)}{D_\alpha(\rho\|\sigma)} \right] = \left[\inf_{\alpha \in (0,1)} \left(\frac{\ln\left(\frac{\varepsilon^{\alpha'}}{\delta}\right)}{D_\alpha(\rho\|\sigma)} \right) \right]. \quad (\text{H13})$$

Now applying the same reasoning, but instead considering the expression in (32), we conclude the following upper bound:

$$n^*(\rho, \sigma, \varepsilon, \delta) \leq \left[\inf_{\alpha \in (0,1)} \left(\frac{\ln\left(\frac{\delta^{\alpha'}}{\varepsilon}\right)}{D_\alpha(\sigma\|\rho)} \right) \right], \quad (\text{H14})$$

thus finishing the proof of the second inequality in (44).

Appendix I: Proof of Corollary 10

Before proving Corollary 10, we state Lemma 23 below. Several aspects of this lemma were already proven as [53, Lemma 13] and extended to the infinite-dimensional case in [135, Lemma 4.13] (see also [136, 137] for earlier statements along these lines). Here, we restate the lemma, since our convention with all relative entropies in our paper involves the natural logarithm instead of the binary logarithm. Furthermore, since it is rather cumbersome to check and translate all the details of the proof of [53, Lemma 13] using our convention, we have opted to provide a complete proof below, which follows that of [53, Lemma 13] quite closely. It should be noted also that the proof below holds in the more general case that σ is a positive semi-definite operator, which can be useful in other applications. It can be extended to the infinite-dimensional case by employing the techniques of [135].

Lemma 23. *Let ρ be a quantum state and σ a positive semi-definite operator. For $\gamma \in (0, 1]$, define*

$$c_\gamma(\rho\|\sigma) := \frac{1}{\gamma} \ln(e^{\gamma D_{1+\gamma}(\rho\|\sigma)} + e^{-\gamma D_{1-\gamma}(\rho\|\sigma)} + 1). \quad (\text{I1})$$

The following bounds hold for all $\gamma \in (0, 1]$:

$$D(\rho\|\sigma) \leq c_\gamma(\rho\|\sigma), \quad (\text{I2})$$

$$D(\rho\|\sigma) \leq \gamma [c_\gamma(\rho\|\sigma)]^2. \quad (\text{I3})$$

In the case that $D(\rho\|\sigma) \geq 0$, for all $\gamma \in (0, 1]$ and $\zeta \in (0, \frac{\gamma}{2}]$, the following bound holds:

$$D_{1+\zeta}(\rho\|\sigma) \leq D(\rho\|\sigma) + \zeta [c_\gamma(\rho\|\sigma)]^2. \quad (\text{I4})$$

For all $\gamma \in (0, 1]$ and $\zeta \in (0, \frac{\ln 3}{2c_\gamma(\rho\|\sigma)}]$, the following bound holds:

$$D_{1+\zeta}(\rho\|\sigma) \leq D(\rho\|\sigma) + \zeta K [c_\gamma(\rho\|\sigma)]^2, \quad (\text{I5})$$

where

$$K := \cosh\left(\frac{\ln 3}{2}\right) \approx 1.155. \quad (\text{I6})$$

If there exists $\gamma' \in (0, 1]$ such that $D_{1+\gamma'}(\rho\|\sigma) < +\infty$, then for all $\gamma \in (0, \gamma']$ and $\zeta \in (0, \frac{\ln 3}{2c_\gamma(\rho\|\sigma)}]$, the following bound holds:

$$D(\rho\|\sigma) \leq D_{1-\zeta}(\rho\|\sigma) + \zeta K [c_\gamma(\rho\|\sigma)]^2, \quad (\text{I7})$$

Proof. First, suppose that $D(\rho\|\sigma) = +\infty$. Then it follows that $c_\gamma(\rho\|\sigma) = +\infty$, so that (I2), (I3), and (I4) hold trivially in this case.

So let us now suppose that $D(\rho\|\sigma) < +\infty$ and prove (I2) and (I3). Observe that

$$D(\rho\|\sigma) \leq D_{1+\gamma}(\rho\|\sigma) = \frac{1}{\gamma} \ln(e^{\gamma D_{1+\gamma}(\rho\|\sigma)}) \leq c_\gamma(\rho\|\sigma), \quad (\text{I8})$$

thus establishing (I2). Also, under this assumption, we necessarily have that $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$. So then we can restrict the Hilbert space to the support of σ , where σ is an invertible (positive definite) operator. Now define

$$X := \rho \otimes (\sigma^{-1})^T, \quad (\text{I9})$$

$$|\varphi\rangle := (\sqrt{\rho} \otimes I) \sum_i |i\rangle|i\rangle, \quad (\text{I10})$$

so that $|\varphi\rangle$ is a normalized state vector, and observe that

$$D(\rho\|\sigma) = \langle \varphi | \ln X | \varphi \rangle, \quad (\text{I11})$$

and for all $\zeta \in [-1, 0) \cup (0, 1]$,

$$D_{1+\zeta}(\rho\|\sigma) = \frac{1}{\zeta} \ln \langle \varphi | X^\zeta | \varphi \rangle. \quad (\text{I12})$$

Additionally, observe that for all $\gamma \in (0, 1]$,

$$c_\gamma(\rho\|\sigma) = \frac{1}{\gamma} \ln \langle \varphi | (X^\gamma + X^{-\gamma} + I) | \varphi \rangle, \quad (\text{I13})$$

Since the function $x^\gamma + x^{-\gamma} + 1$ has a minimum at $x = 1$, it follows that $x^\gamma + x^{-\gamma} + 1 \geq 3$ for all $x > 0$. So the following inequality holds for all $\gamma \in (0, 1]$:

$$\gamma c_\gamma(\rho\|\sigma) \geq \ln 3 \geq 1. \quad (\text{I14})$$

Thus, from (I8) and (I14), we conclude that $\gamma D(\rho\|\sigma) \leq \gamma c_\gamma(\rho\|\sigma) \leq [\gamma c_\gamma(\rho\|\sigma)]^2$, which implies (I3).

We now prove (I4). For all $t > 0$ and $\zeta \in \mathbb{R}$, the first term in the Taylor expansion of t^ζ around $\zeta = 0$ is $\zeta \ln t$. Then we can write $t^\zeta = 1 + \zeta \ln t + r_\zeta(t)$, where

$$r_\zeta(t) := t^\zeta - \zeta \ln t - 1. \quad (\text{I15})$$

Given that $1 + x \leq e^x$ for all $x \in \mathbb{R}$, we find that

$$r_\zeta(t) = t^\zeta - \zeta \ln t - 1 \quad (\text{I16})$$

$$= t^\zeta + 1 - \zeta \ln t - 2 \quad (\text{I17})$$

$$\leq t^\zeta + e^{-\zeta \ln t} - 2 \quad (\text{I18})$$

$$= e^{\zeta \ln t} + e^{-\zeta \ln t} - 2 \quad (\text{I19})$$

$$= 2 [\cosh(\zeta \ln t) - 1] =: s_\zeta(t). \quad (\text{I20})$$

Observe that

$$s_{-\zeta}(t) = s_{\zeta}(t), \quad (\text{I21})$$

$$s_{\zeta}(t) = s_{\gamma\zeta}(t^{1/\gamma}) \quad \forall \gamma \in \mathbb{R}. \quad (\text{I22})$$

One can also verify that $s_{\zeta}(t)$ is monotonically increasing in t for $t \geq 1$ and concave in t if $\zeta \leq 1/2$ and $t \geq 3$. For all $t \geq 0$, either t or $1/t$ is ≥ 1 , so that we can use the monotonicity of $s_{\zeta}(t)$ to conclude that

$$s_{\zeta}(t) = s_{\zeta}\left(\frac{1}{t}\right) \leq s_{\zeta}\left(t + \frac{1}{t}\right). \quad (\text{I23})$$

Using this, we conclude for all $t \geq 0$ and $\gamma \in (0, 1]$ that

$$s_{\zeta}(t) = s_{\zeta/\gamma}(t^{\gamma}) \quad (\text{I24})$$

$$\leq s_{\zeta/\gamma}(t^{\gamma} + t^{-\gamma}) \quad (\text{I25})$$

$$\leq s_{\zeta/\gamma}(t^{\gamma} + t^{-\gamma} + 1). \quad (\text{I26})$$

As mentioned above, the function $t^{\gamma} + t^{-\gamma} + 1$ has a minimum at $t = 1$, and so it follows that $t^{\gamma} + t^{-\gamma} + 1 \geq 3$ for all $t > 0$. Then we can use the inequality above, i.e., $s_{\zeta}(t) \leq s_{\zeta/\gamma}(t^{\gamma} + t^{-\gamma} + 1)$, and concavity of $s_{\zeta/\gamma}$ for $\zeta \leq \frac{\gamma}{2}$ to conclude that

$$\langle \varphi | s_{\zeta}(X) | \varphi \rangle \leq \langle \varphi | s_{\zeta/\gamma}(X^{\gamma} + X^{-\gamma} + I) | \varphi \rangle \quad (\text{I27})$$

$$\leq s_{\zeta/\gamma}(\langle \varphi | (X^{\gamma} + X^{-\gamma} + I) | \varphi \rangle) \quad (\text{I28})$$

$$= s_{\zeta/\gamma}(e^{\gamma c_{\gamma}(\rho \| \sigma)}). \quad (\text{I29})$$

The second inequality follows from applying Jensen's inequality. The equality follows from (I13). We finally use Taylor's theorem in the Lagrange form of the remainder to conclude that the following holds for all $t \geq 0$ and some $\eta \in (0, \zeta)$:

$$s_{\zeta}(t) = s_0(t) + \frac{d}{d\zeta} s_{\zeta}(t) \Big|_{\zeta=0} \zeta + \frac{1}{2} \frac{d^2}{d\zeta^2} s_{\zeta}(t) \Big|_{\zeta=\eta} \zeta^2 \quad (\text{I30})$$

$$= \zeta^2 (\ln t)^2 \cosh(\eta \ln t) \quad (\text{I31})$$

$$\leq \zeta^2 (\ln t)^2 \cosh(\zeta \ln t). \quad (\text{I32})$$

For the second equality, we used the facts that

$$s_0(t) = \frac{d}{d\zeta} s_{\zeta}(t) \Big|_{\zeta=0} = 0, \quad (\text{I33})$$

$$\frac{d^2}{d\zeta^2} s_{\zeta}(t) \Big|_{\zeta=\eta} = 2 (\ln t)^2 \cosh(\eta \ln t). \quad (\text{I34})$$

For the inequality in (I32), we used the fact that $\eta \mapsto \cosh(\eta \ln t)$ is monotone increasing on $\eta \in (0, \zeta]$. Then we conclude that

$$\langle \varphi | s_{\zeta}(X) | \varphi \rangle \leq s_{\zeta/\gamma}(e^{\gamma c_{\gamma}(\rho \| \sigma)}) \quad (\text{I35})$$

$$\leq \left(\frac{\zeta}{\gamma}\right)^2 (\ln e^{\gamma c_{\gamma}(\rho \| \sigma)})^2 \cosh\left(\frac{\zeta}{\gamma} \ln e^{\gamma c_{\gamma}(\rho \| \sigma)}\right) \quad (\text{I36})$$

$$= \zeta^2 [c_\gamma(\rho\|\sigma)]^2 \cosh(\zeta c_\gamma(\rho\|\sigma)). \quad (\text{I37})$$

Then for $\zeta \in (0, \frac{\gamma}{2}]$ and when $D(\rho\|\sigma) \geq 0$, consider that

$$D_{1+\zeta}(\rho\|\sigma) = \frac{1}{\zeta} \ln \langle \varphi | X^\zeta | \varphi \rangle \quad (\text{I38})$$

$$= \frac{1}{\zeta} \ln [1 + \zeta \langle \varphi | \ln X | \varphi \rangle + \langle \varphi | r_\zeta(X) | \varphi \rangle] \quad (\text{I39})$$

$$= \frac{1}{\zeta} \ln [1 + \zeta D(\rho\|\sigma) + \langle \varphi | r_\zeta(X) | \varphi \rangle] \quad (\text{I40})$$

$$\leq \frac{1}{\zeta} \ln [1 + \zeta D(\rho\|\sigma) + \langle \varphi | s_\zeta(X) | \varphi \rangle] \quad (\text{I41})$$

$$= \frac{1}{\zeta} \ln [1 + \zeta D(\rho\|\sigma)] + \frac{1}{\zeta} \ln \left[1 + \frac{\langle \varphi | s_\zeta(X) | \varphi \rangle}{1 + \zeta D(\rho\|\sigma)} \right] \quad (\text{I42})$$

$$\leq D(\rho\|\sigma) + \frac{1}{\zeta} \ln \left[1 + \frac{\langle \varphi | s_\zeta(X) | \varphi \rangle}{1 + \zeta D(\rho\|\sigma)} \right] \quad (\text{I43})$$

$$\leq D(\rho\|\sigma) + \frac{1}{\zeta} \ln [1 + \langle \varphi | s_\zeta(X) | \varphi \rangle] \quad (\text{I44})$$

$$\leq D(\rho\|\sigma) + \frac{1}{\zeta} \ln [1 + \zeta^2 [c_\gamma(\rho\|\sigma)]^2 \cosh(\zeta c_\gamma(\rho\|\sigma))] \quad (\text{I45})$$

$$\leq D(\rho\|\sigma) + \zeta [c_\gamma(\rho\|\sigma)]^2. \quad (\text{I46})$$

The first equality follows from (I12), the second equality from (I15), and the third equality from (I11). The first inequality follows from (I16)–(I20). The second inequality follows because $\ln(1+x) \leq x$ for $x \geq 0$. The third inequality follows because $\zeta D(\rho\|\sigma) \geq 0$. The penultimate inequality follows from (I35)–(I37). The final inequality follows because the function $k \rightarrow k^2 - \ln(1 + k^2 \cosh(k))$ is monotonically increasing for $k \geq 0$, so that it is non-negative.

We now prove (I5). Suppose that $\zeta \in (0, \frac{\ln 3}{2c_\gamma(\rho\|\sigma)}]$, while making no constraint on the sign of $D(\rho\|\sigma)$. Observe from (I14) that

$$\frac{\ln 3}{2c_\gamma(\rho\|\sigma)} \leq \frac{\gamma c_\gamma(\rho\|\sigma)}{2c_\gamma(\rho\|\sigma)} = \frac{\gamma}{2}, \quad (\text{I47})$$

ensuring that the needed concavity of $s_{\zeta/\gamma}(X)$ in (I28) holds. Starting from the first few steps above, consider that

$$D_{1+\zeta}(\rho\|\sigma) = \frac{1}{\zeta} \ln [1 + \zeta D(\rho\|\sigma) + \langle \varphi | r_\zeta(X) | \varphi \rangle] \quad (\text{I48})$$

$$\leq \frac{1}{\zeta} \ln [1 + \zeta D(\rho\|\sigma) + \langle \varphi | s_\zeta(X) | \varphi \rangle] \quad (\text{I49})$$

$$\leq D(\rho\|\sigma) + \frac{1}{\zeta} \langle \varphi | s_\zeta(X) | \varphi \rangle \quad (\text{I50})$$

$$\leq D(\rho\|\sigma) + \zeta [c_\gamma(\rho\|\sigma)]^2 \cosh(\zeta c_\gamma(\rho\|\sigma)) \quad (\text{I51})$$

$$\leq D(\rho\|\sigma) + \zeta [c_\gamma(\rho\|\sigma)]^2 \cosh\left(\frac{\ln 3}{2}\right) \quad (\text{I52})$$

$$= D(\rho\|\sigma) + \zeta K [c_\gamma(\rho\|\sigma)]^2. \quad (\text{I53})$$

The first inequality follows from (I12). The second inequality follows because $\ln(1+x) \leq x$ for $x \geq -1$, the latter condition ensured because $\zeta D(\rho\|\sigma) + \langle \varphi | r_\zeta(X) | \varphi \rangle \geq -1$. The third inequality follows from (I35)–(I37). The penultimate inequality follows from the assumption that $\zeta \in (0, \frac{\ln 3}{2c_\gamma(\rho\|\sigma)}]$ and the fact that \cosh is monotone increasing on this interval. The final equality follows from the definition of K in (I6).

Finally, we prove (I7). Suppose that $\zeta \in (0, \frac{\ln 3}{2c_\gamma(\rho\|\sigma)}]$. Similar to the above, consider that

$$D_{1-\zeta}(\rho\|\sigma) = -\frac{1}{\zeta} \ln[1 - \zeta D(\rho\|\sigma) + \langle \varphi | r_{-\zeta}(X) | \varphi \rangle] \quad (\text{I54})$$

$$\geq -\frac{1}{\zeta} \ln[1 - \zeta D(\rho\|\sigma) + \langle \varphi | s_{-\zeta}(X) | \varphi \rangle] \quad (\text{I55})$$

$$= -\frac{1}{\zeta} \ln[1 - \zeta D(\rho\|\sigma) + \langle \varphi | s_\zeta(X) | \varphi \rangle] \quad (\text{I56})$$

$$\geq D(\rho\|\sigma) - \frac{1}{\zeta} \langle \varphi | s_\zeta(X) | \varphi \rangle \quad (\text{I57})$$

$$\geq D(\rho\|\sigma) - \zeta [c_\gamma(\rho\|\sigma)]^2 \cosh(\zeta c_\gamma(\rho\|\sigma)) \quad (\text{I58})$$

$$\geq D(\rho\|\sigma) - \zeta [c_\gamma(\rho\|\sigma)]^2 \cosh\left(\frac{\ln 3}{2}\right) \quad (\text{I59})$$

$$= D(\rho\|\sigma) - \zeta K [c_\gamma(\rho\|\sigma)]^2. \quad (\text{I60})$$

The first inequality follows from (I12). The second equality follows from (I21). The second inequality follows because $\ln(1+x) \leq x$ for $x \geq -1$. The final steps follow from similar reasons as above. \square

With Lemma 23 in hand, we now move on to the proof of Corollary 10.

Proof of Corollary 10. Let $\gamma \in (0, 1]$ be such that $D_{1+\gamma}(\rho\|\sigma) < +\infty$. Define

$$c_\gamma(\rho\|\sigma) := \frac{1}{\gamma} \ln(e^{\gamma D_{1+\gamma}(\rho\|\sigma)} + e^{-\gamma D_{1-\gamma}(\rho\|\sigma)} + 1). \quad (\text{I61})$$

For all ξ' satisfying $0 < \xi' < \gamma/2$, the following inequality holds as a consequence of (I4):

$$D_{1+\xi'}(\rho\|\sigma) \leq D(\rho\|\sigma) + \xi' [c_\gamma(\rho\|\sigma)]^2. \quad (\text{I62})$$

Setting $\xi := \xi' [c_\gamma(\rho\|\sigma)]^2$, we can rewrite this as follows: for all ξ satisfying

$$0 < \xi < \frac{\gamma [c_\gamma(\rho\|\sigma)]^2}{2}, \quad (\text{I63})$$

the following inequality holds:

$$D_{1+\xi'}(\rho\|\sigma) \leq D(\rho\|\sigma) + \xi. \quad (\text{I64})$$

Now applying Theorem 9, while noting that the sandwiched Rényi relative entropies therein can be replaced with the Petz–Rényi relative entropies (due to $\tilde{D}_\alpha \leq D_\alpha$ holding for all $\alpha > 0$), we conclude that

$$n^*(\rho, \sigma, \varepsilon, \delta) \geq \frac{\ln\left(\frac{(1-\varepsilon)^{\frac{1+\xi'}{\xi'}}}{\delta}\right)}{D_{1+\xi'}(\rho\|\sigma)} \quad (\text{I65})$$

$$= \frac{\ln\left(\frac{1}{\delta}\right) - \left(\frac{1}{\xi'} + 1\right) \ln\left(\frac{1}{1-\varepsilon}\right)}{D_{1+\xi'}(\rho\|\sigma)} \quad (\text{I66})$$

$$\geq \frac{\ln\left(\frac{1}{\delta}\right) - \left(\frac{1}{\xi'} + 1\right) \ln\left(\frac{1}{1-\varepsilon}\right)}{D(\rho\|\sigma) + \xi} \quad (\text{I67})$$

$$= \frac{\ln\left(\frac{1}{\delta}\right) - \left(\frac{[c_\gamma(\rho\|\sigma)]^2}{\xi} + 1\right) \ln\left(\frac{1}{1-\varepsilon}\right)}{D(\rho\|\sigma) + \xi}. \quad (\text{I68})$$

Let us now pick $\xi = \frac{1}{4}D(\rho\|\sigma)$. Then the inequality in (I63) follows as a consequence of (I3). Substituting above, we find that

$$n^*(\rho, \sigma, \varepsilon, \delta) \geq \frac{\ln\left(\frac{1}{\delta}\right) - \left(\frac{4[c_\gamma(\rho\|\sigma)]^2}{D(\rho\|\sigma)} + 1\right) \ln\left(\frac{1}{1-\varepsilon}\right)}{\frac{5}{4}D(\rho\|\sigma)}. \quad (\text{I69})$$

Thus, if

$$\ln\left(\frac{1}{\delta}\right) \geq 2 \left(\frac{4[c_\gamma(\rho\|\sigma)]^2}{D(\rho\|\sigma)} + 1 \right) \ln\left(\frac{1}{1-\varepsilon}\right) \quad (\text{I70})$$

$$\iff \delta \leq \exp \left[-2 \left(\frac{4[c_\gamma(\rho\|\sigma)]^2}{D(\rho\|\sigma)} + 1 \right) \ln\left(\frac{1}{1-\varepsilon}\right) \right], \quad (\text{I71})$$

then

$$n^*(\rho, \sigma, \varepsilon, \delta) \geq \frac{2}{5} \left(\frac{\ln\left(\frac{1}{\delta}\right)}{D(\rho\|\sigma)} \right), \quad (\text{I72})$$

proving the lower bound in (48).

For all ξ' satisfying $0 < \xi' < \frac{\ln 3}{2c_\gamma(\rho\|\sigma)}$, the following inequality holds, as a consequence of (I7):

$$D(\rho\|\sigma) \leq D_{1-\xi'}(\rho\|\sigma) + \xi' K [c_\gamma(\rho\|\sigma)]^2, \quad (\text{I73})$$

where

$$K := \cosh\left(\frac{\ln 3}{2}\right). \quad (\text{I74})$$

By setting $\xi := \xi' K [c_\gamma(\rho\|\sigma)]^2$, we can rewrite this as follows: For all ξ satisfying

$$0 < \xi < \frac{K \ln 3}{2} c_\gamma(\rho\|\sigma), \quad (\text{I75})$$

the following inequality holds:

$$D(\rho\|\sigma) \leq D_{1-\xi'}(\rho\|\sigma) + \xi. \quad (\text{I76})$$

Proceeding similarly as before, and applying Theorem 9, consider that

$$n^*(\rho, \sigma, \varepsilon, \delta) \leq \left\lceil \frac{\ln\left(\frac{\varepsilon^{-(1-\xi')/\xi'}}{\delta}\right)}{D_{1-\xi'}(\rho\|\sigma)} \right\rceil \quad (\text{I77})$$

$$= \left\lceil \frac{\ln\left(\frac{1}{\delta}\right) + \left(\frac{1}{\xi'} - 1\right) \ln\left(\frac{1}{\varepsilon}\right)}{D_{1-\xi'}(\rho\|\sigma)} \right\rceil \quad (\text{I78})$$

$$\leq \left\lceil \frac{\ln\left(\frac{1}{\delta}\right) + \left(\frac{1}{\xi'} - 1\right) \ln\left(\frac{1}{\varepsilon}\right)}{D(\rho\|\sigma) - \xi} \right\rceil \quad (\text{I79})$$

$$= \left\lceil \frac{\ln\left(\frac{1}{\delta}\right) + \left(\frac{K[c_\gamma(\rho\|\sigma)]^2}{\xi} - 1\right) \ln\left(\frac{1}{\varepsilon}\right)}{D(\rho\|\sigma) - \xi} \right\rceil. \quad (\text{I80})$$

Observing that $\frac{K \ln 3}{2} \approx 0.634$, now pick

$$\xi = \frac{1}{2} D(\rho\|\sigma), \quad (\text{I81})$$

so that the inequality in (I75) is satisfied, as a consequence of (I2). Then

$$n^*(\rho, \sigma, \varepsilon, \delta) \leq \left\lceil \frac{\ln\left(\frac{1}{\delta}\right) + \left(\frac{2K[c_\gamma(\rho\|\sigma)]^2}{D(\rho\|\sigma)} - 1\right) \ln\left(\frac{1}{\varepsilon}\right)}{\frac{1}{2} D(\rho\|\sigma)} \right\rceil. \quad (\text{I82})$$

Thus, if

$$\ln\left(\frac{1}{\delta}\right) \geq \left(\frac{2K[c_\gamma(\rho\|\sigma)]^2}{D(\rho\|\sigma)} - 1\right) \ln\left(\frac{1}{\varepsilon}\right) \quad (\text{I83})$$

$$\iff \delta \leq \exp\left[-\left(\frac{2K[c_\gamma(\rho\|\sigma)]^2}{D(\rho\|\sigma)} - 1\right) \ln\left(\frac{1}{\varepsilon}\right)\right], \quad (\text{I84})$$

then

$$n^*(\rho, \sigma, \varepsilon, \delta) \leq \left\lceil 4 \left(\frac{\ln\left(\frac{1}{\delta}\right)}{D(\rho\|\sigma)}\right) \right\rceil, \quad (\text{I85})$$

proving the upper bound in (48). \square

Appendix J: Proof of Theorem 11

The achievability part (i.e., the upper bound in (49)) can be deduced from [72, Theorem 7.2]. Here, we provide an alternative simple proof for that. In addition, we provide a bound with improved constants, by using [72, Eq. (8)].

We use the following pretty-good measurement [65, 66] with respect to the ensemble \mathcal{S} ; i.e., for each $m \in \{1, \dots, M\}$,

$$\Lambda_m := \left(\sum_{\bar{m}=1}^M A_{\bar{m}} \right)^{-\frac{1}{2}} A_m \left(\sum_{\bar{m}=1}^M A_{\bar{m}} \right)^{-\frac{1}{2}}, \quad (\text{J1})$$

$$A_m := p_m \rho_m. \quad (\text{J2})$$

Then,

$$p_e(\mathcal{S}) \leq \sum_{m=1}^M \text{Tr}[A_m (I - \Lambda_m)] \quad (\text{J3})$$

$$= \sum_{m=1}^M \text{Tr} \left[A_m \left(A_m + \sum_{\bar{m} \neq m} A_{\bar{m}} \right)^{-\frac{1}{2}} \left(\sum_{\bar{m} \neq m} A_{\bar{m}} \right) \left(A_m + \sum_{\bar{m} \neq m} A_{\bar{m}} \right)^{-\frac{1}{2}} \right] \quad (\text{J4})$$

$$\stackrel{\text{(a)}}{\leq} \sum_{m=1}^M \left\| \sqrt{A_m} \sqrt{\sum_{\bar{m} \neq m} A_{\bar{m}}} \right\|_1 \quad (\text{J5})$$

$$\stackrel{\text{(b)}}{\leq} \sum_{m=1}^M \sum_{\bar{m} \neq m} \left\| \sqrt{A_m} \sqrt{A_{\bar{m}}} \right\|_1 \quad (\text{J6})$$

$$= \sum_{m=1}^M \sum_{\bar{m} \neq m} \sqrt{p_m} \sqrt{p_{\bar{m}}} \sqrt{F}(\rho_m, \rho_{\bar{m}}) \quad (\text{J7})$$

$$\leq M(M-1) \max_{\bar{m} \neq m} \sqrt{p_m} \sqrt{p_{\bar{m}}} \sqrt{F}(\rho_m, \rho_{\bar{m}}), \quad (\text{J8})$$

where (a) follows from Lemma 18-(v), the first inequality of Lemma 18-(iv), and the first inequality of Lemma 18-(ii); (b) follows from the subadditivity property stated in Fact (vi).

Using the multiplicativity of the fidelity F , i.e., Lemma 18-(i), and following the same reasoning used in (F23)–(F29), we obtain the following upper bound on the sample complexity:

$$n^*(\mathcal{S}, \varepsilon) \leq \left\lceil \max_{m \neq \bar{m}} \frac{2 \ln \left(\frac{M(M-1) \sqrt{p_m} \sqrt{p_{\bar{m}}}}{\varepsilon} \right)}{-\ln F(\rho_m, \rho_{\bar{m}})} \right\rceil. \quad (\text{J9})$$

Next, we provide an improved bound compared to (J9) by adapting [72, Eq. (8)]. Using this inequality, we find that

$$p_e(\mathcal{S}) \leq \frac{1}{2} \sum_{m=1}^M \sum_{\bar{m} \neq m} \sqrt{p_m} \sqrt{p_{\bar{m}}} \sqrt{F}(\rho_m, \rho_{\bar{m}}) \quad (\text{J10})$$

$$\leq \frac{1}{2}M(M-1) \max_{\tilde{m} \neq m} \sqrt{p_m} \sqrt{p_{\tilde{m}}} \sqrt{F}(\rho_m, \rho_{\tilde{m}}). \quad (\text{J11})$$

Similarly, by using the multiplicativity of the fidelity F , and following the same reasoning used in (F23)–(F29), we obtain the following upper bound on the sample complexity:

$$n^*(\mathcal{S}, \varepsilon) \leq \left\lceil \max_{m \neq \tilde{m}} \frac{2 \ln \left(\frac{M(M-1) \sqrt{p_m} \sqrt{p_{\tilde{m}}}}{2\varepsilon} \right)}{-\ln F(\rho_m, \rho_{\tilde{m}})} \right\rceil, \quad (\text{J12})$$

thus establishing the second inequality in (49).

The converse (i.e., the lower bound in (49)) follows from classical reasoning; namely, any instance of pairwise simple hypothesis testing is easier than the M -ary hypothesis testing. Let $\{\Lambda_1^*, \dots, \Lambda_M^*\}$ be an optimum measurement achieving the minimum error probability, i.e.

$$p_e(\mathcal{S}) = \sum_{m=1}^M \text{Tr}[A_m (I - \Lambda_m^*)]. \quad (\text{J13})$$

Then, for every $m \neq \tilde{m}$,

$$p_e(\mathcal{S}) \geq \text{Tr}[A_m (I - \Lambda_m^*)] + \text{Tr}[A_{\tilde{m}} (I - \Lambda_{\tilde{m}}^*)] \quad (\text{J14})$$

$$\geq \inf_{\Lambda_m, \Lambda_{\tilde{m}} \geq 0} \left\{ \text{Tr}[A_m (I - \Lambda_m)] + \text{Tr}[A_{\tilde{m}} (I - \Lambda_{\tilde{m}})] : \Lambda_m + \Lambda_{\tilde{m}} = I - \sum_{\tilde{m} \neq m, \tilde{m}} \Lambda_{\tilde{m}} \right\} \quad (\text{J15})$$

$$\geq \inf_{\Lambda_m, \Lambda_{\tilde{m}} \geq 0} \left\{ \text{Tr}[A_m (I - \Lambda_m)] + \text{Tr}[A_{\tilde{m}} (I - \Lambda_{\tilde{m}})] : \Lambda_m + \Lambda_{\tilde{m}} = I \right\} \quad (\text{J16})$$

$$= (p_m + p_{\tilde{m}}) \inf_{\substack{\Lambda_m, \Lambda_{\tilde{m}} \geq 0 \\ \Lambda_m + \Lambda_{\tilde{m}} = I}} \text{Tr} \left[\frac{p_m}{p_m + p_{\tilde{m}}} \rho_m (I - \Lambda_m) \right] + \text{Tr} \left[\frac{p_{\tilde{m}}}{p_m + p_{\tilde{m}}} \rho_{\tilde{m}} (I - \Lambda_{\tilde{m}}) \right] \quad (\text{J17})$$

$$\geq (p_m + p_{\tilde{m}}) \frac{p_m p_{\tilde{m}}}{(p_m + p_{\tilde{m}})^2} F(\rho_m, \rho_{\tilde{m}}). \quad (\text{J18})$$

To see the third inequality, let $\Lambda_m, \Lambda_{\tilde{m}} \geq 0$ be arbitrary operators satisfying $\Lambda_m + \Lambda_{\tilde{m}} = I - \sum_{\tilde{m} \neq m, \tilde{m}} \Lambda_{\tilde{m}}$. Then choose $X, Y \geq 0$ such that $X + Y = I - (\Lambda_m + \Lambda_{\tilde{m}})$. Now set $\Lambda'_m := \Lambda_m + X$ and $\Lambda'_{\tilde{m}} := \Lambda_{\tilde{m}} + Y$. So then $\Lambda'_m, \Lambda'_{\tilde{m}} \geq 0$ and $\Lambda'_m + \Lambda'_{\tilde{m}} = I$, and thus the choices Λ'_m and $\Lambda'_{\tilde{m}}$ are feasible for the optimization in the third line. Furthermore, we have that $\text{Tr}[(I - \Lambda_m)\omega] \geq \text{Tr}[(I - \Lambda'_m)\omega]$ and $\text{Tr}[(I - \Lambda_{\tilde{m}})\omega] \geq \text{Tr}[(I - \Lambda'_{\tilde{m}})\omega]$ for every $\omega \geq 0$. Thus, the objective function in the third line is never larger than that in the second line, thus justifying the third inequality. The last line follows from the converse (i.e., the lower bound of the sample complexity) of binary hypothesis testing, as shown in (F6) in the proof of Theorem 7.

Hence, by using the multiplicativity of the fidelity again, we obtain

$$n^*(\mathcal{S}, \varepsilon) \geq \max_{m \neq \tilde{m}} \frac{\ln \left(\frac{p_m p_{\tilde{m}}}{(p_m + p_{\tilde{m}})\varepsilon} \right)}{-\ln F(\rho_m, \rho_{\tilde{m}})}, \quad (\text{J19})$$

concluding the proof.

Remark 8. *The key ingredient of proving the upper bound in Theorem 11 is the subadditivity of the square-root of the fidelity (Lemma 18-(vi)). However, the Holevo fidelity generally does not satisfy such a property. This is why we could not directly achieve the Holevo fidelity in (J8). On the other hand, the result mentioned in Remark 4 directly yields the Holevo fidelity. Note that one could still achieve the Holevo fidelity by losing a factor of 2 (Lemma 18-(i)), i.e.,*

$$\frac{1}{-\ln F(\rho, \sigma)} \leq \frac{2}{-\ln F_{\text{H}}(\rho, \sigma)}. \quad (\text{J20})$$

Appendix K: Proof of Remark 4

Let $\mathcal{S}^{(n)} := \{(p_m, \rho_m^{\otimes n})\}_{m=1}^M$. The following bound was established as [15, Eq. (36)], as a step in the proof of the multiple quantum Chernoff exponent:

$$p_e(\mathcal{S}^{(n)}) \leq 5M(M-1)^3(n+1)^{2d} \max_{1 \leq m \leq M} \{p_m\} \max_{m \neq \bar{m}} e^{-nC(\rho_m \parallel \rho_{\bar{m}})}, \quad (\text{K1})$$

where d is the dimension of the underlying Hilbert space. Denote

$$c := 5M(M-1)^3 \max_{1 \leq m \leq M} \{p_m\} \quad \text{and} \quad B := \min_{m \neq \bar{m}} \frac{C(\rho_m \parallel \rho_{\bar{m}})}{2}. \quad (\text{K2})$$

Note that the function $f(x) = C(x+1)^{2d}e^{-Bx}$ is maximised at $x_0 = \frac{2d}{B} - 1$ with maximal value

$$A := f(x_0) = c \left(\frac{2d}{B} \right)^d e^{B-2d} \quad (\text{K3})$$

$$= 5M(M-1)^3 \max_{1 \leq m \leq M} \{p_m\} \left(\frac{4d}{\min_{m \neq \bar{m}} C(\rho_m \parallel \rho_{\bar{m}})} \right)^d e^{\min_{m \neq \bar{m}} \frac{C(\rho_m \parallel \rho_{\bar{m}})}{2} - 2d} \quad (\text{K4})$$

From (K1) we therefore see

$$p_e(\mathcal{S}^{(n)}) \leq c(n+1)^{2d}e^{-2Bn} \leq Ae^{-Bn} \quad (\text{K5})$$

which hence gives

$$n^*(\mathcal{S}, \varepsilon) \leq \left\lceil \frac{\ln(A/\varepsilon)}{\min_{m \neq \bar{m}} \frac{C(\rho_m \parallel \rho_{\bar{m}})}{2}} \right\rceil. \quad (\text{K6})$$

Invoking $C(\rho_m \parallel \rho_{\bar{m}}) \geq -\frac{1}{2} \ln F(\rho_m, \rho_{\bar{m}})$ (Lemma 18-(ii)), the above bound also implies

$$n^*(\mathcal{S}, \varepsilon) \leq O\left(\max_{m \neq \bar{m}} \frac{\ln M}{-\ln F(\rho_m, \rho_{\bar{m}})} \right). \quad (\text{K7})$$