

Improved quantum data analysis

Costin Bădescu*

Ryan O’Donnell*

November 24, 2020

Abstract

We provide more sample-efficient versions of some basic routines in quantum data analysis, along with simpler proofs. Particularly, we give a quantum “Threshold Search” algorithm that requires only $O((\log^2 m)/\epsilon^2)$ samples of a d -dimensional state ρ . That is, given observables $0 \leq A_1, A_2, \dots, A_m \leq \mathbb{1}$ such that $\text{tr}(\rho A_i) \geq 1/2$ for at least one i , the algorithm finds j with $\text{tr}(\rho A_j) \geq 1/2 - \epsilon$. As a consequence, we obtain a Shadow Tomography algorithm requiring only $\tilde{O}((\log^2 m)(\log d)/\epsilon^4)$ samples, which simultaneously achieves the best known dependence on each parameter m, d, ϵ . This yields the same sample complexity for quantum Hypothesis Selection among m states; we also give an alternative Hypothesis Selection method using $\tilde{O}((\log^3 m)/\epsilon^2)$ samples.

1 Introduction

Some of the most basic problems in statistics, unsupervised learning, and property testing involve the following scenario: One can observe data that are assumed to be drawn independently from an unknown probability distribution p ; say that p is discrete and supported on $[d] = \{1, 2, \dots, d\}$. The task is to learn, test, or estimate some properties of p . Completely estimating p up to error ϵ (in, say, total variation distance) requires $\Theta(d/\epsilon^2)$ samples, so when d is very large one may seek to only learn or test *partial* aspects of p . For example, one might only want to estimate the means of some known, fixed random variables $a_1, \dots, a_m : [d] \rightarrow [0, 1]$ (sometimes called “statistical queries” in the learning/privacy literature). Or, one might want to perform Hypothesis Selection over some set of two or more hypothesis distributions q_1, \dots, q_m on $[d]$. It is generally fairly straightforward to determine the optimal sample complexity needed for these tasks. For example, it’s easy to show that one can simultaneously estimate all expectations $\mathbf{E}_p[a_1], \dots, \mathbf{E}_p[a_m]$ to accuracy $\pm\epsilon$ using a batch of $n = O((\log m)/\epsilon^2)$ samples (independent of d): one simply computes the empirical mean for each a_i , reusing the batch of samples in each computation.

*Computer Science Department, Carnegie Mellon University. Supported by NSF grant FET-1909310 and ARO grant W911NF2110001. This material is based upon work supported by the National Science Foundation under grant numbers listed above. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation (NSF). {cbadescu,odonnell}@cs.cmu.edu

These kinds of questions become much more difficult to analyze when the classical source of randomness p is replaced by a *quantum* source of randomness, namely a d -dimensional quantum state $\rho \in \mathbb{C}^{d \times d}$ (satisfying $\rho \geq 0$, $\text{tr}(\rho) = 1$). The difficulties here are that: (i) one cannot directly observe “outcomes” for ρ , one can only measure it; (ii) measuring the state ρ inherently alters it, hence reusing samples (i.e., copies of ρ) is problematic. For example, suppose we now have some known, fixed observables $A_1, \dots, A_m \in \mathbb{C}^{d \times d}$ with $0 \leq A_i \leq \mathbb{1}$ and we wish to estimate each expectation $\mathbf{E}_\rho[A_i] := \text{tr}(\rho A_i)$ to within $\pm \epsilon$. This is the “Shadow Tomography” problem introduced by Aaronson in [Aar16] (see [Aar20] for applications to, e.g., quantum money). We do not know if this is similarly possible using $n = O((\log m)/\epsilon^2)$ copies of ρ ; indeed, prior to this work the best known upper bound was

$$n = \min \left\{ \tilde{O}((\log^4 m)(\log d)/\epsilon^4), \quad \tilde{O}((\log^2 m)(\log^2 d)/\epsilon^8) \right\}.$$

Here the sample complexity on the left is from [Aar20], combining a “Gentle Search” routine with an online learning algorithm for quantum states from [ACH⁺19]. The sample complexity on the right was obtained by Aaronson and Rothblum [AR19] by drawing inspiration and techniques from the field of Differential Privacy.¹

In fact, we propose that — rather than Differential Privacy — a closer classical match for the Shadow Tomography problem is the task known as *Adaptive Data Analysis*, introduced by [DFH⁺16]. In this problem, the random variables (“statistical queries”) a_1, \dots, a_m are not fixed in advance for the learner, but are rather received one at a time, with the crucial feature that each a_t may adaptively depend on the preceding estimates of $\mathbf{E}_p[a_1], \dots, \mathbf{E}_p[a_{t-1}]$ output by the learner. In this case, conditioning on these output estimates skews the underlying i.i.d. product distribution $p^{\otimes n}$ — reminiscent of the way measuring a quantum state affects it — and this prevents naive reuse of the sample data. Indeed, it’s far from obvious that the Adaptive Data Analysis task is doable with $\text{poly}(\log m, \log d, 1/\epsilon)$ samples; however this was shown by [DFH⁺16], who achieved complexity $n = \tilde{O}((\log m)^{3/2}(\log d)^{1/2}/\epsilon^{7/2})$, and this was later improved by [BNS⁺15] to $n = \tilde{O}((\log m)(\log d)^{1/2}/\epsilon^3)$. While Differential Privacy tools have been an ingredient in some Adaptive Data Analysis routines, the topics are not inherently linked; e.g., a viewpoint based on “KL-stability” is emphasized in [BNS⁺15].

1.1 Our work

1.1.1 Threshold Search

The first main result in our work concerns what we will call the quantum “Threshold Search” problem.² We state the problem here in a general form (recalling our notation $\mathbf{E}_\rho[A_i] = \text{tr}(\rho A_i)$):

¹See also [HKP20] for sample complexity bounds that can be better for special kinds of A_i ’s.

²Originally called the “Secret Acceptor” problem when it was introduced by Aaronson [Aar16]. Later he called it “Gentle Search” [Aar20], but we find this name unsatisfactory as it is not necessary that a successful algorithm be “gentle”. In the Differential Privacy literature, it is sometimes called “Report Noisy Max” (offline case) or “Above Threshold” (online case) [DR13].

Quantum Threshold Search problem: *Given:*

1. Parameters $0 < \epsilon, \delta < \frac{1}{2}$.
2. Access to unentangled copies of an unknown d -dimensional quantum state ρ .
3. A list of d -dimensional observables $0 \leq A_1, \dots, A_m \leq \mathbb{1}$.
4. A list of thresholds $0 \leq \theta_1, \dots, \theta_m \leq 1$.

The algorithm should either output:

- “ $\mathbf{E}_\rho[A_j] > \theta_j - \epsilon$ ” for some particular j ; or else,
- “ $\mathbf{E}_\rho[A_i] \leq \theta_i$ for all i ”.

The task is to minimize the number n of copies that are used, while ensuring the probability of a false output statement is at most δ .

We remark that all of the difficulty of the problem is contained in the case where $\epsilon = \delta = \frac{1}{4}$ and $\theta_j = \frac{3}{4}$ for all j (see [Section 4.1](#)). In this case, Aaronson [[Aar16](#)] originally showed that the Threshold Search problem can be solved using $n = \widetilde{O}(\log^4 m)$ copies of ρ . In the present paper, we improve this result quadratically:

Theorem 1.1. *The quantum Threshold Search problem can be solved using*

$$n = n_{\text{TS}}(m, \epsilon, \delta) = \frac{\log^2 m + L}{\epsilon^2} \cdot O(L) \quad (L = \log(1/\delta))$$

copies of ρ . Furthermore, this solution is online in the sense that:

- The algorithm is initially given only m, ϵ, δ . It then selects n and obtains $\rho^{\otimes n}$.
- Next, observable/threshold pairs $(A_1, \theta_1), (A_2, \theta_2), \dots$ are presented to the algorithm in sequence. When each (A_t, θ_t) is presented, the algorithm must either “pass”, or else halt and output “ $\mathbf{E}_\rho[A_t] > \theta_t - \epsilon$ ”.
- If the algorithm passes on all (A_t, θ_t) pairs, then it ends by outputting “ $\mathbf{E}_\rho[A_i] \leq \theta_i$ for all i ”.

Incidentally, the (offline) quantum Threshold *Decision* problem, where the algorithm only needs to report “ $\exists j : \mathbf{E}_\rho[A_j] > \theta_j - \epsilon$ ” without actually specifying j , is known to be solvable using just $n = O(\log(m) \log(1/\delta)/\epsilon^2)$ copies [[Aar20](#)]. We review the proof in [Appendix A](#), tightening/simplifying some quantitative aspects of the underlying theorem of Harrow, Lin, and Montanaro [[HLM17](#)]. In particular, our tightenings let us slightly improve the copy complexity to $n = O(\log(m/\delta)/\epsilon^2)$.

1.1.2 χ^2 -stable threshold reporting

The most important technical ingredient going into our proof of [Theorem 1.1](#) is a new, purely classical statistical result fitting into the Adaptive Data Analysis framework (see, e.g., [\[Smi17\]](#) for some background). In that setting one might describe our result as follows: “adding exponential noise provides a (composably) χ^2 -stable mechanism for reporting if a distribution’s mean is above a given threshold”. In more detail, the result says that given a Sample \mathbf{S} consisting of the sum of n draws from a Bernoulli(p) distribution (i.e., $\mathbf{S} \sim \text{Binomial}(n, p)$), if we add independent eXponential noise \mathbf{X} and then check the event B that $\mathbf{S} + \mathbf{X}$ exceeds some large threshold θn , then conditioning on B *not* occurring hardly changes the distribution of \mathbf{S} , provided $\mathbf{E}[\mathbf{X}] \gg \text{stddev}[\mathbf{S}]$. Here the phrase “hardly changes” is in two very strong senses: (i) we show the random variables $\mathbf{S} \mid \overline{B}$ and \mathbf{S} are close even in χ^2 -divergence, which is a more stringent measure than KL-divergence (or Hellinger distance, or total variation distance) — that is, the test is “ χ^2 -stable”; (ii) the χ^2 -divergence is not just absolutely small, but is even a small fraction of $\mathbf{P}[B]^2$ itself (hence the total variation closeness is a small fraction of $\mathbf{P}[B]$). This allows a kind of composition (as in the “Sparse Vector” mechanism [\[DR13\]](#) from the Differential Privacy literature) in which the same quantum sample can be reused for repeated “above threshold” tests, up until the point where having at least one “above threshold” outcome becomes likely. Precisely, our result is the following:

Theorem 1.2. *Let $\mathbf{S} \sim \text{Binomial}(n, p)$. Assume that \mathbf{X} is an independent Exponential random variable with mean at least $\text{stddev}[\mathbf{S}] = \sqrt{p(1-p)n}$ (and also at least 1). Let B be the event that $\mathbf{S} + \mathbf{X} > \theta n$, and assume that $\mathbf{P}[B] < \frac{1}{4}$. Then*

$$d_{\chi^2}(\mathbf{S} \mid \overline{B}, \mathbf{S}) \lesssim \left(\mathbf{P}[B] \cdot \frac{\text{stddev}[\mathbf{S}]}{\mathbf{E}[\mathbf{X}]} \right)^2 \leq \mathbf{P}[B]^2 \cdot (n / \mathbf{E}[\mathbf{X}]^2).$$

(Above we are using the notation $Y \lesssim Z$ to mean $Y \leq C \cdot Z$ for some universal constant C . We are also abusing notation by writing the χ^2 -divergence between two random variables to mean the χ^2 -divergence between their underlying distributions.)

Corollary 1.3. *Writing \mathbf{S}' for $\mathbf{S} \mid \overline{B}$, standard inequalities for f -divergences [\[GS02\]](#) imply*

$$d_{\text{TV}}(\mathbf{S}', \mathbf{S}) \leq d_{\text{H}}(\mathbf{S}', \mathbf{S}) \leq \sqrt{d_{\text{KL}}(\mathbf{S}', \mathbf{S})} \leq \sqrt{d_{\chi^2}(\mathbf{S}', \mathbf{S})} \lesssim \mathbf{P}[B] \cdot \frac{\text{stddev}[\mathbf{S}]}{\mathbf{E}[\mathbf{X}]} \leq \mathbf{P}[B] \cdot \frac{\sqrt{n}}{\mathbf{E}[\mathbf{X}]}.$$

Let us remark that our [Theorem 1.2](#) is similar to results appearing previously in the Differential Privacy/Adaptive Data Analysis literature; in particular, it is quite similar to (and inspired by) a theorem (“Claim 41”) of Aaronson and Rothblum [\[AR19\]](#). Although this Claim 41 is presented in a quantum context, the essence of it is a theorem comparable to our [Theorem 1.2](#), with the following main differences: (i) it bounds the weaker KL-divergence (though for our applications, this is acceptable); (ii) the proof is significantly more involved. (Minor differences include: (i) their result uses two-sided exponential noise for a two-sided threshold event; (ii) our bound has the stronger factor $\text{stddev}[\mathbf{S}]$ instead of just \sqrt{n} .)

1.1.3 Applications: Shadow Tomography and Hypothesis Selection

Given our improved Threshold Search algorithm, we present two applications in quantum data analysis. The first is to the aforementioned Shadow Tomography problem, where we obtain a sample complexity that simultaneously achieves the best known dependence on all three parameters m, d, ϵ . Furthermore, our algorithm is *online*, as in the Adaptive Data Analysis setting.

Theorem 1.4. *There is a quantum algorithm that, given parameters $m \in \mathbb{N}$, $0 < \epsilon < \frac{1}{2}$, and access to unentangled copies of a state $\rho \in \mathbb{C}^{d \times d}$, obtains a sample of*

$$n = \frac{(\log^2 m + L)(\log d)}{\epsilon^4} \cdot O(L) \quad (L = \log(\frac{\log d}{\delta \epsilon}))$$

copies of ρ and then has the following behavior: When any (adversarially/adaptively chosen) sequence of observables $A_1, A_2, \dots, A_m \in \mathbb{C}^{d \times d}$ with $0 \leq A_i \leq \mathbb{1}$ is presented to the algorithm one-by-one, once A_t is presented the algorithm responds with an estimate $\hat{\mu}_i$ of $\mathbf{E}_\rho[A_t] = \text{tr}(\rho A_t)$. Except with probability at most δ (over the algorithm’s measurements), all m estimates satisfy $|\hat{\mu}_i - \mathbf{E}_\rho[A_t]| \leq \epsilon$.

The proof of this theorem is almost immediate from our Threshold Search algorithm, using a known [Aar20] black-box reduction to the online quantum state learning algorithm of Aaronson, Chen, Hazan, Kale, and Nayak [ACH⁺19].

Let us philosophically remark that we believe the importance of the parameters, in increasing order, is d , then ϵ , then m . Regarding d , “in practice” one may expect that $\log d$, the number of qubits in the unknown state, is not likely to be particularly large. Indeed, many problems in quantum learning/tomography/statistics [HM02, CM06, MdW16, OW15, OW16, HHJ⁺17, OW17, AISW19, BOW19, Yu20a, Yu20b, BCL20, Yu20c] have polynomial dependence on d , so factors of polylog d seem of lesser importance. Regarding ϵ , “in practice” this might be the most important parameter, as even with a very mild value like $\epsilon = .1$, a dependence of $1/\epsilon^4$ is challenging. It’s peculiar that all works on Shadow Tomography have achieved atypical ϵ -dependence like $1/\epsilon^4$, $1/\epsilon^5$, and $1/\epsilon^8$, instead of the “expected” $1/\epsilon^2$; on the other hand, this peculiarity also seems to occur in the Adaptive Data Analysis literature. Finally, we feel that the dependence on m is of the most interest (theoretical interest, at least), and it would be extremely compelling if we could reduce the dependence from $\log^2 m$ to $\log m$. Our reason is related to quantum Hypothesis Selection, which we now discuss.

Hypothesis Selection. The classical (multiple) Hypothesis Selection problem [Yat85, DL96, DL97] is as follows: Given are m fixed “hypothesis” probability distributions q_1, \dots, q_m on $[d]$, as well as a parameter ϵ and access to samples from an unknown distribution p on $[d]$. The task is to find (with probability at least $1 - \delta$) a q_j which is, roughly, closest to p , while minimizing the number of samples drawn from p . More precisely, if $\eta = \min_i \{d_{\text{TV}}(p, q_i)\}$, the algorithm should output a hypothesis q_j with $d_{\text{TV}}(p, q_j) \leq C\eta + \epsilon$ for some fixed small constant C . There are a variety of solutions known to this problem, with standard ones [DL01, Chap. 6] achieving $n = O((\log m)/\epsilon^2)$ (and best constant $C = 3$). There are also numerous variations, including handling different distance measures besides d_{TV} [BV08], the easier (“realizable/non-robust”) case when $\eta = 0$, and the case when there is a unique answer (as

when the hypotheses q_j are pairwise far apart). We emphasize that our focus is on the *non-asymptotic regime*, where we would like an explicit sample bound $n = n(m, d, \epsilon, \delta)$ holding for all values of m, d, ϵ, δ .³ One particularly useful application of Hypothesis Selection is to *learning* an unknown probability distribution p from a class \mathcal{C} (even “agnostically”). Roughly speaking, if \mathcal{C} has an ϵ -cover of size $m = m(\epsilon)$, then one can learn p to accuracy $O(\epsilon)$ using a Hypothesis Selection over m hypotheses; i.e., with $O((\log m)/\epsilon^2)$ samples in the classical case. For further discussion of the problem, see e.g. [MŠ07]; for Differentially Private Hypothesis Selection, see [BKS20, GKK+20]; for fast classical Hypothesis Selection with a quantum computer, see [QCR20].

The *quantum* Hypothesis Selection problem is the natural analogue in which probability distributions are replaced by quantum states, and total variation distance is replaced by trace distance. As with Shadow Tomography (and Differentially Private Hypothesis Selection), it is nontrivial to upgrade classical algorithms due to the fact that samples cannot be naively reused. We show that one can use Shadow Tomography as a black box to solve quantum Hypothesis Testing. We also give a different method based on Threshold Search that achieves an incomparable copy complexity, with a better dependence on ϵ but a worse dependence on m : roughly $(\log^3 m)/\epsilon^2$, versus the $(\log^2 m)/\epsilon^4$ of Shadow Tomography. Finally, we show that if the hypothesis states are pairwise far apart, we can match the optimal bound from the classical case.

Theorem 1.5. *There is a quantum algorithm that, given m fixed hypothesis states $\sigma_1, \dots, \sigma_m \in \mathbb{C}^{d \times d}$, parameters $0 < \epsilon, \delta < \frac{1}{2}$, and access to unentangled copies of a state $\rho \in \mathbb{C}^{d \times d}$, uses a sample of*

$$n = \min \left\{ \frac{(\log^2 m + L_1)(\log d)}{\epsilon^4} \cdot O(L_1), \frac{\log^3 m + L_2 \cdot \log m}{\epsilon^2} \cdot O(L_2) \right\}$$

copies of ρ (where $L_1 = \log(\frac{\log d}{\delta \epsilon})$ and $L_2 = \log(\log(1/\max\{\eta, \epsilon\})/\delta)$) and has the following guarantee: except with probability at most δ , it outputs k such that

$$d_{\text{tr}}(\rho, \sigma_k) \leq 3.01\eta + \epsilon, \quad \text{where } \eta = \min_i \{d_{\text{tr}}(\rho, \sigma_i)\}.$$

Further, assuming $\eta < \frac{1}{2}(\min_{i \neq j} \{d_{\text{tr}}(\sigma_i, \sigma_j)\} - \epsilon)$ (so there is a unique σ_i near to ρ), one can find the σ_k achieving $d_{\text{tr}}(\rho, \sigma_k) = \eta$ (except with probability at most δ) using only $n = O(\log(m/\delta)/\epsilon^2)$ copies of ρ .

The fact that quantum Hypothesis Selection black-box reduces to Shadow Tomography provides significant motivation for trying to prove (or disprove) that Shadow Tomography can be done with $O(\log m) \cdot \text{poly}((\log d)/\epsilon)$ copies; i.e., that the power on $\log m$ can be reduced to 1. If this were possible, then as in the classical case we would be able to *learn* a quantum state $\rho \in \mathbb{C}^{d \times d}$ in a class \mathcal{C} (to constant trace distance accuracy, say) using $\log(|\text{cover}(\mathcal{C})|) \cdot \text{polylog}(d)$ copies, where $\text{cover}(\mathcal{C})$ denotes a set of states that form a (trace-distance) cover for \mathcal{C} . It’s easy to see that the class \mathcal{C} of *all* states has a cover of size

³This is as opposed to the *asymptotic regime*. There, one focuses on achieving $\delta \leq \exp(-C(m, d, \epsilon)n)$ for all $n \geq n_0(m, d, \epsilon)$, where the rate function $C(m, d, \epsilon)$ should be as large as possible, but where n_0 may be a completely uncontrolled function of m, d, ϵ . See, e.g., [NS11].

at most $O(d)^{d^2}$, and hence Shadow Tomography with $\log m$ dependence would yield a full quantum tomography algorithm with copy complexity $\tilde{O}(d^2)$, bypassing the sophisticated representation-theory methods of [OW16, HHJ⁺17, OW17]. One might also hope for more efficient learning of other interesting subclasses of states; e.g., the class *separable* states.

2 Preliminaries

2.1 Classical probability distributions and distances

Let $p = (p_1, \dots, p_d)$ denote a probability distribution on $[d] = \{1, \dots, d\}$. We consider $A : [d] \rightarrow \mathbb{R}$ to be a random variable on $[d]$, and write

$$\mathbf{E}_p[A] = \mathbf{E}_{\mathbf{i} \sim p}[A(\mathbf{i})] = \sum_{i=1}^d p_i A(i).$$

In particular, if $A : [d] \rightarrow \{0, 1\}$ we may think of it as an *event* $A \subseteq [d]$.

Given another probability distribution q on $[d]$, there are a variety of important distances/divergences between p and q . We now recall all those appearing in [Theorem 1.2](#) and [Corollary 1.3](#).

The *total variation* distance $d_{\text{TV}}(p, q)$ between p and q is defined by

$$d_{\text{TV}}(p, q) = \frac{1}{2} \sum_{i=1}^d |p_i - q_i| = \max_{A \subseteq [d]} \left| \mathbf{E}_p[A] - \mathbf{E}_q[A] \right|.$$

The *Bhattacharyya coefficient* $\text{BC}(p, q)$ (an affinity between p and q , rather than a distance) is defined by

$$\text{BC}(p, q) = \sum_{i=1}^d \sqrt{p_i q_i}.$$

This can be used to define *squared Hellinger distance* $d_{\text{H}}(p, q)^2 = d_{\text{H}^2}(p, q)$, viz.,

$$d_{\text{H}^2}(p, q) = 2(1 - \text{BC}(p, q)) = \sum_{i=1}^d (\sqrt{p_i} - \sqrt{q_i})^2.$$

The *KL-divergence* $d_{\text{KL}}(p, q)$ between p and q is defined by

$$d_{\text{KL}}(p, q) = \sum_{i=1}^d p_i \ln(p_i/q_i) = \mathbf{E}_{\mathbf{i} \sim p} \ln(p_i/q_i).$$

Finally, the χ^2 -divergence $d_{\chi^2}(p, q)$ between p and q is defined by

$$d_{\chi^2}(p, q) = \sum_{i=1}^d q_i \left(1 - \frac{p_i}{q_i}\right)^2 = \mathbf{E}_{\mathbf{i} \sim q} \left[\left(1 - \frac{p_i}{q_i}\right)^2 \right].$$

2.2 Quantum states and measurements

A matrix $A \in \mathbb{C}^{d \times d}$ is said to be *Hermitian*, or *self-adjoint*, if $A^\dagger = A$; here A^\dagger denotes the conjugate transpose of A . We write $A \geq 0$ to denote that A is self-adjoint and positive semidefinite; e.g., $B^\dagger B \geq 0$ always. In general, we write $A \geq B$ to mean $A - B \geq 0$. Recall that a positive semidefinite matrix $A \geq 0$ has a unique positive semidefinite square root $\sqrt{A} \geq 0$. We write $\mathbb{1}$ for the identity matrix (where the dimension is understood from context).

A d -dimensional *quantum state* is any $\rho \in \mathbb{C}^{d \times d}$ satisfying $\rho \geq 0$ and $\text{tr } \rho = 1$; physically speaking, this is the state of a d -level quantum system, such as $\log_2 d$ qubits. A d -dimensional *observable* is any self-adjoint $A \in \mathbb{C}^{d \times d}$; physically speaking, this is any real-valued property of the system. One can build an associated measuring device that takes in a quantum system in state ρ , and reads out a (stochastic) real number; we denote its expected value, the *expectation of A with respect to ρ* , by

$$\mathbf{E}_\rho[A] = \text{tr}(\rho A).$$

It is a basic fact of linear algebra that $\mathbf{E}_\rho[A] \geq 0$ whenever $A \geq 0$.

Note that if ρ and A are diagonal matrices then we reduce to the classical case, where the diagonal elements of ρ form a probability distribution on $[d]$ and the diagonal elements of A give a real-valued random variable.

We will use the term *quantum event* for an observable $A \in \mathbb{C}^{d \times d}$ with $0 \leq A \leq \mathbb{1}$; i.e., a self-adjoint operator with all its eigenvalues between 0 and 1. A state $\rho \in \mathbb{C}^{d \times d}$ assigns a probability $0 \leq \mathbf{E}_\rho[A] \leq 1$ to each event. We reserve the term *projector* for the special case when $A^2 = A$; i.e., when all of A 's eigenvalues are either 0 or 1. Note that we have not exactly paralleled the classical terminology, where an “event” is a random variable with all its values equal to 0 or 1, but: (i) it's convenient to have a brief term for observables A with $0 \leq A \leq \mathbb{1}$; (ii) the terminology “projector” is very standard. Of course, by the spectral theorem, every quantum event A may be written as

$$A = \sum_{i=1}^r \lambda_i \Pi_i, \text{ where each } 0 \leq \lambda_i \leq 1, \text{ and } \Pi_i \text{'s are pairwise orthogonal projectors.} \quad (1)$$

A *quantum measurement* \mathcal{M} , also known as a positive-operator valued measure (POVM), is a sequence $\mathcal{M} = (A_1, \dots, A_k)$ of quantum events with $A_1 + \dots + A_k = \mathbb{1}$. Since

$$\mathbf{E}_\rho[A_1] + \dots + \mathbf{E}_\rho[A_k] = \mathbf{E}_\rho[A_1 + \dots + A_k] = \mathbf{E}_\rho[\mathbb{1}] = 1,$$

a state ρ and a measurement \mathcal{M} determine a probability distribution p on $[k]$ defined by $p_i = \mathbf{E}_\rho[A_i]$ for $i = 1, \dots, k$. A common scenario is that of a *two-outcome measurement*, associated to any quantum event A ; this is the measurement $\mathcal{M} = (\bar{A}, A)$, where $\bar{A} = \mathbb{1} - A$.

For any quantum measurement \mathcal{M} , one can physically implement a measuring device that, given ρ , reports $\mathbf{i} \in [k]$ distributed according to p . Mathematically, an *implementation* of $\mathcal{M} = (A_1, \dots, A_k)$ is a sequence of d -column matrices M_1, \dots, M_k with $M_i^\dagger M_i = A_i$ for $i = 1, \dots, k$. Under this implementation, conditioned on the readout being $\mathbf{i} = i$, the state ρ

collapses to the new state $\rho|_{M_i}$, defined as follows:

$$\rho|_{M_i} = \frac{M_i \rho M_i^\dagger}{\mathbf{E}_\rho[M_i^\dagger M_i]} = \frac{M_i \rho M_i^\dagger}{\mathbf{E}_\rho[A_i]}.$$

Given \mathcal{M} , we will define the *canonical implementation* to be the one in which $M_i = \sqrt{A_i}$. In particular, if we have any quantum event A and we canonically implement the associated two-outcome measurement (\overline{A}, A) , then measuring ρ and conditioning on A occurring yields the new state

$$\rho|_{\sqrt{A}} = \frac{\sqrt{A} \rho \sqrt{A}}{\mathbf{E}_\rho[A]}.$$

More generally, we have the mathematical notion of a *quantum operation* S on d -dimensional states, defined by d -column matrices M_1, \dots, M_k such that

$$M_1^\dagger M_1 + \dots + M_k^\dagger M_k \leq \mathbb{1}.$$

The result of applying S to a state ρ is (the sub-normalized state)

$$S(\rho) = M_1 \rho M_1^\dagger + \dots + M_k \rho M_k^\dagger.$$

Finally, we will use the following special case of the well-known Naimark dilation theorem:

Theorem 2.1 (Naimark). *If $A \in \mathbb{C}^{d \times d}$ is a quantum event, then there exists a projector Π operating on the space \mathbb{C}^{2d} such that, for any $\rho \in \mathbb{C}^{d \times d}$,*

$$\mathbf{E}_{\rho \otimes |0\rangle\langle 0|}[\Pi] = \mathbf{E}_\rho[A].$$

2.3 Quantum state distances

Just as with classical probability distributions, there are a variety of distances/divergences between two quantum states $\rho, \sigma \in \mathbb{C}^{d \times d}$. In fact, for every classical “ f -divergence” there is a corresponding “measured quantum f -divergence”, which is the maximal classical divergence that can be achieved by performing the same measurement on ρ and σ . In this way, classical total variation distance precisely corresponds to quantum trace distance, the Bhattacharyya coefficient precisely corresponds to quantum fidelity, etc. See, e.g., [BOW19, Sec. 3.1.2] for further review; here we will simply directly define some quantum distances.

The *trace distance* $d_{\text{tr}}(\rho, \sigma)$ between states ρ and σ is defined by

$$d_{\text{tr}}(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\|_1 = \max_{0 \leq A \leq \mathbb{1}} |\mathbf{E}_\rho[A] - \mathbf{E}_\sigma[A]|.$$

Here the second equality is known as Helstrom’s theorem [Hel76], and the maximum is over all quantum events $A \in \mathbb{C}^{d \times d}$. (Indeed, Helstrom shows the maximum is achieved by a projector.) The *fidelity* $F(\rho, \sigma)$ between states ρ and σ is defined by

$$F(\rho, \sigma) = \|\sqrt{\rho} \sqrt{\sigma}\|_1 = \text{tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}}. \quad (2)$$

This can be used to define the *squared Bures distance* $d_{\text{Bures}}(\rho, \sigma)^2 = d_{\text{Bures}^2}(\rho, \sigma)$, viz.,

$$d_{\text{Bures}^2}(\rho, \sigma) = 2(1 - F(\rho, \sigma)).$$

It follows from the work of Fuchs and Caves [FC95] that $\frac{1}{2}d_{\text{Bures}^2}(\rho, \sigma) \leq d_{\text{tr}}(\rho, \sigma) \leq d_{\text{Bures}}(\rho, \sigma)$ for all states ρ and σ .

Below we give a simpler formula for fidelity in the case when σ is a conditioned version of ρ (such results are sometimes known under the name ‘‘gentle measurement’’; see [Wat18, Cor. 3.15]):

Proposition 2.2. *Let $\rho \in \mathbb{C}^{d \times d}$ and $M \in \mathbb{C}^{d \times d}$ an observable. Then $F(\rho, \rho|_M)^2 = \frac{\mathbf{E}_\rho[M]^2}{\mathbf{E}_\rho[M^2]}$. In particular, for a projector Π we get $F(\rho, \rho|_\Pi) = \sqrt{\mathbf{E}_\rho[\Pi]}$, and for conditioning on the occurrence of a quantum event A (under the canonical implementation), $F(\rho, \rho|_{\sqrt{A}}) = \frac{\mathbf{E}_\rho[\sqrt{A}]}{\sqrt{\mathbf{E}_\rho[A]}}$.*

Proof. Using the definition of $\rho|_M$ and the second formula for fidelity in Equation (2),

$$F(\rho, \rho|_M)^2 = \frac{\text{tr}(\sqrt{\sqrt{\rho}M\rho M^\dagger\sqrt{\rho}})^2}{\mathbf{E}_\rho[M^\dagger M]} = \frac{\text{tr}(\sqrt{\sqrt{\rho}M\sqrt{\rho}\sqrt{\rho}M\sqrt{\rho}})^2}{\mathbf{E}_\rho[M^2]} = \frac{\text{tr}(\sqrt{\rho}M\sqrt{\rho})^2}{\mathbf{E}_\rho[M^2]} = \frac{\mathbf{E}_\rho[M]^2}{\mathbf{E}_\rho[M^2]}.$$

Below we give a further formula for $F(\rho, \rho|_{\sqrt{A}})$ using the spectral decomposition of A . (We remark that it may be obtained as a special case of the theorem of Fuchs and Caves [FC95].)

Proposition 2.3. *Let $\rho \in \mathbb{C}^{d \times d}$ be a quantum state, and let $A \in \mathbb{C}^{d \times d}$ be a quantum event with spectral decomposition $A = \sum_{i=1}^r \lambda_i \Pi_i$ as in Equation (1). Let p be the probability distribution on $[r]$ determined by measurement $\mathcal{M} = (\Pi_1, \dots, \Pi_r)$ on ρ , and let q be the one determined by \mathcal{M} on $\rho|_{\sqrt{A}}$. Then $F(\rho, \rho|_{\sqrt{A}}) = \text{BC}(p, q)$.*

Proof. By definition,

$$\mathbf{E}_{\rho|_{\sqrt{A}}}[\Pi_i] \cdot \mathbf{E}_\rho[A] = \text{tr}(\sqrt{A}\rho\sqrt{A}\Pi_i) = \mathbf{E}_\rho[\sqrt{A}\Pi_i\sqrt{A}] = \mathbf{E}_\rho[\lambda_i\Pi_i] = \lambda_i p_i,$$

and hence $q_i = \lambda_i p_i / \mathbf{E}_\rho[A]$. It follows that

$$\text{BC}(p, q) = \frac{\sum_i \sqrt{\lambda_i} p_i}{\sqrt{\mathbf{E}_\rho[A]}} = \frac{\mathbf{E}_\rho[\sqrt{A}]}{\sqrt{\mathbf{E}_\rho[A]}},$$

and the proof is complete by Proposition 2.2. ■

2.4 Naive expectation estimation

Lemma 2.4. *Let $E \in \mathbb{C}^{d \times d}$ be a quantum event and let $0 < \epsilon, \delta < \frac{1}{2}$. Then there exists $n = O(\log(1/\delta)/\epsilon^2)$ (not depending on E) and a measurement $\mathcal{M} = (A_0, \dots, A_n)$ for such that, for any quantum state $\rho \in \mathbb{C}^{d \times d}$,*

$$\mathbf{P}\left[\left|\frac{\mathbf{k}}{n} - \text{tr}(\rho E)\right| > \epsilon\right] \leq \delta,$$

where $\mathbf{k} \in \{0, \dots, n\}$ is the random outcome of the measurement \mathcal{M} applied to the state $\rho^{\otimes n}$.
 Moreover, for parameters $0 \leq \tau, c \leq 1$, there exists a quantum event B such that

$$\begin{aligned} |\operatorname{tr}(\rho E) - \tau| > c + \epsilon &\implies \mathbf{E}_{\rho^{\otimes n}}[B] \geq 1 - \delta \text{ and} \\ |\operatorname{tr}(\rho E) - \tau| \leq c - \epsilon &\implies \mathbf{E}_{\rho^{\otimes n}}[B] \leq \delta. \end{aligned}$$

Additionally, if E is a projector, then so is B .

Proof. Let $E_1 = E$ and $E_0 = \mathbb{1} - E$. For all $x \in \{0, 1\}^n$, let $E_x \in (\mathbb{C}^{d \times d})^{\otimes n}$ be defined by $E_x = E_{x_1} \otimes E_{x_2} \otimes \dots \otimes E_{x_n}$. For $k = 0, \dots, n$, let $A_k \in (\mathbb{C}^{d \times d})^{\otimes n}$ be the quantum event defined by

$$A_k = \sum_{\substack{x \in \{0, 1\}^n \\ |x| = k}} E_x.$$

Let \mathcal{M} be the measurement defined by $\mathcal{M} = \{A_0, \dots, A_n\}$.

Thus, if $\mathbf{k} \in \{0, \dots, n\}$ is the random outcome of measuring $\rho^{\otimes n}$ against \mathcal{M} , then \mathbf{k} is distributed as Binomial($n, \operatorname{tr}(\rho E)$). Hence, if $n = O(\log(1/\delta)/\epsilon^2)$, then, by a Chernoff bound,

$$\mathbf{P} \left[\left| \frac{\mathbf{k}}{n} - \operatorname{tr}(\rho E) \right| \geq \epsilon \right] \leq 2 \exp(-2\epsilon^2/n) \leq \delta.$$

Let parameters $\tau, c \in [0, 1]$ be given and let the function $f : [0, 1] \rightarrow \{0, 1\}$ be defined by

$$f(t) = \begin{cases} 1, & |t - \tau| \geq c, \\ 0, & \text{otherwise.} \end{cases}$$

Finally, let the quantum event B be defined by

$$B = \sum_{k=0}^n f(k/n) A_k.$$

Thus, if $\mathbf{k} \sim \text{Binomial}(n, \operatorname{tr}(\rho E))$, then

$$\mathbf{E}_{\rho^{\otimes n}}[B] = \sum_{k=0}^n \mathbf{P}[\mathbf{k} = k] \cdot f(k/n) = \mathbf{E}[f(\mathbf{k}/n)] = \mathbf{P} \left[\left| \frac{\mathbf{k}}{n} - \tau \right| \geq c \right].$$

If $c + \epsilon \leq |\operatorname{tr}(\rho E) - \tau|$, then $|\operatorname{tr}(\rho E) - \mathbf{k}/n| < \epsilon$ implies $|\mathbf{k}/n - \tau| \geq c$. Hence,

$$\mathbf{E}_{\rho^{\otimes n}}[B] = \mathbf{P} \left[\left| \frac{\mathbf{k}}{n} - \tau \right| \geq c \right] \geq \mathbf{P} \left[\left| \frac{\mathbf{k}}{n} - \operatorname{tr}(\rho E) \right| < \epsilon \right] \geq 1 - \delta.$$

If $c - \epsilon \geq |\operatorname{tr}(\rho E) - \tau|$, then $|\operatorname{tr}(\rho E) - \mathbf{k}/n| < \epsilon$ implies $|\mathbf{k}/n - \tau| < c$. Hence,

$$\mathbf{E}_{\rho^{\otimes n}}[\overline{B}] = \mathbf{P} \left[\left| \frac{\mathbf{k}}{n} - \tau \right| < c \right] \geq \mathbf{P} \left[\left| \frac{\mathbf{k}}{n} - \operatorname{tr}(\rho E) \right| < \epsilon \right] \geq 1 - \delta.$$

If E is a projector, then A_k is a projector and $A_k A_\ell = A_\ell A_k = 0$ for all $k, \ell \in \{0, \dots, n\}$. Since B is a sum of orthogonal projectors A_k with $k \in \{0, \dots, n\}$, it follows that B is a projector. \blacksquare

2.5 Quantum union bound-style results

We will need a lemma (the ‘‘Damage Lemma’’) proved by Aaronson and Rothblum [AR19, Lemma 17]. In fact, we will require the following slight improvement, in which the summation ends at $n - 1$ (rather than n as in [AR19]):

Lemma 2.5. *Let S_1, \dots, S_n be arbitrary quantum operations on d -dimensional quantum states. Let ρ be a quantum state on \mathbb{C}^d with $p_i = \text{tr}(S_i(\rho)) > 0$ for all $i \in [n]$. For all $E \subseteq [n]$,*

$$|p_E - q_E| \leq \frac{2}{q_{[n-1] \setminus E}} \cdot \sum_{i=1}^{n-1} d_{\text{tr}} \left(\frac{S_i(\rho)}{\text{tr}(S_i(\rho))}, \rho \right),$$

where $q_1, \dots, q_n \in [0, 1]$ are defined by $q_1 = p_1$ and, for $2 \leq i \leq n$,

$$q_i = \frac{\text{tr}(S_i(\dots S_1(\rho)))}{\text{tr}(S_{i-1}(\dots S_1(\rho)))}.$$

Proof. For $i = 1, \dots, n$, let $\epsilon_i = d_{\text{tr}}(S_i(\rho)/p_i, \rho)$ and let S'_i be the quantum operation defined by

$$S'_i(\sigma) = \begin{cases} S_i(\sigma), & i \in E, \\ S_i(\sigma)/q_i, & i \notin E. \end{cases}$$

Following the proof of [AR19, Lemma 17], we obtain

$$\|S'_{n-1}(\dots S'_1(\rho)) - p_{E \cap [n-1]} \rho\|_1 \leq \frac{2}{q_{[n-1] \setminus E}} \cdot (\epsilon_1 + \dots + \epsilon_{n-1}).$$

If $n \notin E$, then $p_E = p_{E \cap [n-1]} = \text{tr}(p_{E \cap [n-1]} \rho)$ and $q_E = q_{E \cap [n-1]} = \text{tr}(S'_{n-1}(\dots S'_1(\rho)))$. Since $|\text{tr}(X)| \leq \|X\|_1$ for any self-adjoint matrix X , it follows that

$$\begin{aligned} |q_E - p_E| &= |\text{tr}(S'_{n-1}(\dots S'_1(\rho))) - p_{E \cap [n-1]} \text{tr}(\rho)| \\ &\leq \|S'_{n-1}(\dots S'_1(\rho)) - p_{E \cap [n-1]} \rho\|_1 \\ &\leq \frac{2}{q_{[n-1] \setminus E}} \cdot (\epsilon_1 + \dots + \epsilon_{n-1}). \end{aligned}$$

If $n \in E$, then $q_E = \text{tr}(S'_n(S'_{n-1}(\dots S'_1(\rho))))$ and $p_E = \text{tr}(p_{E \cap [n-1]} S'_n(\rho))$. For all self-adjoint matrices X , $|\text{tr}(X)| \leq \|X\|_1$ and $\|S(X)\|_1 \leq \|X\|_1$ for all quantum operations S . Hence,

$$\begin{aligned} |p_E - q_E| &= |\text{tr}(S'_n(S'_{n-1}(\dots S'_1(\rho)))) - p_{E \cap [n-1]} \text{tr}(S'_n(\rho))| \\ &\leq \|S'_n(S'_{n-1}(\dots S'_1(\rho)) - p_{E \cap [n-1]} \rho)\|_1 \\ &\leq \|S'_{n-1}(\dots S'_1(\rho)) - p_{E \cap [n-1]} \rho\|_1 \\ &\leq \frac{2}{q_{[n-1] \setminus E}} \cdot (\epsilon_1 + \dots + \epsilon_{n-1}). \quad \blacksquare \end{aligned}$$

Finally, for the ‘‘unique decoding’’ part of our Hypothesis Selection routine we will use a related result, Gao’s *quantum Union Bound* [Gao15]:

Lemma 2.6. For each of $i = 1, \dots, m$, let $\Pi_i^1 \in \mathbb{C}^{d \times d}$ be a projector and write $\Pi_i^0 = \mathbb{1} - \Pi_i^1$. Then for any quantum state $\rho \in \mathbb{C}^{d \times d}$,

$$\mathbf{E}_\rho[(\Pi_1^1 \cdots \Pi_m^1)(\Pi_1^1 \cdots \Pi_m^1)^\dagger] \geq 1 - 4 \sum_{i=1}^m \mathbf{E}_\rho[\Pi_i^0].$$

Corollary 2.7. In the preceding setting, suppose that $x \in \{0, 1\}^m$ is such that $\mathbf{E}_\rho[\Pi_i^{x_i}] \geq 1 - \epsilon$ for all $1 \leq i \leq m$. If an algorithm sequentially measures ρ with (Π_1^0, Π_1^1) , measures the resulting state with (Π_2^0, Π_2^1) , measures the resulting state with (Π_3^0, Π_3^1) , etc., then the probability that the measurement outcomes are precisely x_1, x_2, \dots, x_m is at least $1 - 4\epsilon m$.

3 χ^2 -stable Threshold Reporting

Our goal in this section is to prove [Theorem 1.2](#) and to show how this classical result applies to quantum states and measurements. We begin with some preparatory facts.

The following is well known [[Ber95](#)]:

Proposition 3.1. For \mathbf{S} a random variable and $f : \mathbb{R} \rightarrow \mathbb{R}$ Lipschitz, $\mathbf{Var}[f(\mathbf{S})] \leq \mathbf{Var}[\mathbf{S}]$.

Proof. Let \mathbf{S}' be an independent copy of \mathbf{S} . Since the function f is Lipschitz, we always have $\frac{1}{2}(f(\mathbf{S}) - f(\mathbf{S}'))^2 \leq \frac{1}{2}(\mathbf{S} - \mathbf{S}')^2$. The result follows by taking expectations of both sides. ■

We will also use the following simple numerical inequality:

Lemma 3.2. Fix $0 \leq p \leq 1$, $q = 1 - p$. Then for $C = (e - 1)^2 \leq 3$, we have

$$q + pe^{2\lambda} \leq (1 + Cpq\lambda^2) \cdot (q + pe^\lambda)^2 \quad \forall \lambda \in [0, 1].$$

Proof. Since $(q + pe^\lambda)^2 \geq (q + p)^2 = 1$ for $\lambda \geq 0$, it suffices to show

$$q + pe^{2\lambda} \leq (q + pe^\lambda)^2 + Cpq\lambda^2 \quad \forall \lambda \in [0, 1].$$

But $(q + p\Lambda^2) - (q + p\Lambda)^2 = pq(\Lambda - 1)^2$ when $p + q = 1$, so it is further equivalent to show

$$(e^\lambda - 1)^2 \leq C\lambda^2 \quad \forall \lambda \in [0, 1].$$

But this indeed holds with $C = (e - 1)^2$, as it is equivalent to $e^\lambda \leq 1 + (e - 1)\lambda$ on $[0, 1]$, which follows from convexity of $\lambda \mapsto e^\lambda$. ■

We now do a simple calculation showing how much a random variable changes (in χ^2 -divergence) when conditioning on an event. In using the below, the typical mindset is that B is an event that “rarely” occurs, so $\mathbf{P}[\overline{B}]$ is close to 1.

Proposition 3.3. Let \mathbf{S} be a discrete random variable, and let B be an event on the same probability space with $\mathbf{P}[B] < 1$. For each outcome s of \mathbf{S} , define $f(s) = \mathbf{P}[B \mid \mathbf{S} = s]$. Then

$$d_{\chi^2}((\mathbf{S} \mid \overline{B}), \mathbf{S}) = \mathbf{Var}[f(\mathbf{S})] / \mathbf{P}[\overline{B}]^2.$$

Proof. We have the likelihood ratio $\mathbf{P}[\mathbf{S} = s \mid \overline{B}] / \mathbf{P}[\mathbf{S} = s] = (1 - f(s)) / \mathbf{P}[\overline{B}]$, by Bayes' theorem. Hence,

$$d_{\chi^2}((\mathbf{S} \mid \overline{B}), \mathbf{S}) = \mathbf{E} \left[\left(1 - \frac{1 - f(\mathbf{S})}{\mathbf{P}[\overline{B}]} \right)^2 \right] = \frac{1}{\mathbf{P}[\overline{B}]^2} \mathbf{E}[(f(\mathbf{S}) - \mathbf{P}[B])^2] = \mathbf{Var}[f(\mathbf{S})] / \mathbf{P}[\overline{B}]^2,$$

where the last step uses $\mathbf{E}[f(\mathbf{S})] = \mathbf{P}[B]$. \blacksquare

We can now prove [Theorem 1.2](#), which we restate for convenience:

Theorem. *Let $\mathbf{S} \sim \text{Binomial}(n, p)$ and write $q = 1 - p$. Assume that \mathbf{X} is an independent Exponential random variable with mean at least $\text{stddev}[\mathbf{S}] = \sqrt{pqn}$ (and also at least 1). Let B be the event that $\mathbf{S} + \mathbf{X} > \theta n$, and assume that $\mathbf{P}[B] < \frac{1}{4}$. Then*

$$d_{\chi^2}((\mathbf{S} \mid \overline{B}), \mathbf{S}) \lesssim \left(\mathbf{P}[B] \cdot \frac{\text{stddev}[\mathbf{S}]}{\mathbf{E}[\mathbf{X}]} \right)^2.$$

Proof. Write $\lambda = 1/\mathbf{E}[\mathbf{X}]$, so $\mathbf{X} \sim \text{Exponential}(\lambda)$ and we have the assumptions $\lambda \leq \frac{1}{\sqrt{pqn}}$ and $\lambda \leq 1$. Using [Proposition 3.3](#) and $\mathbf{P}[\overline{B}] > \frac{3}{4}$, it suffices to show

$$\mathbf{Var}[f(\mathbf{S})] \lesssim \mathbf{P}[B]^2 \cdot pqn\lambda^2,$$

where

$$f(s) = \mathbf{P}[\mathbf{X} > \theta n - s] = \min(1, g(s)), \quad g(s) = \exp(-\lambda(\theta n - s)).$$

Since $y \mapsto \min(1, y)$ is Lipschitz, [Proposition 3.1](#) tells us that $\mathbf{Var}[f(\mathbf{S})] \leq \mathbf{Var}[g(\mathbf{S})]$. $\mathbf{Var}[g(\mathbf{S})]$ can be computed using the moment-generating function of $\mathbf{S} \sim \text{Binomial}(n, p)$, namely $\mathbf{E}[\exp(t\mathbf{S})] = (q + pe^t)^n$:

$$\begin{aligned} \mathbf{E}[g(\mathbf{S})] &= \mathbf{E}[\exp(-\lambda(\theta n - \mathbf{S}))] = \exp(-\lambda\theta n) \cdot (q + pe^\lambda)^n, \\ \mathbf{E}[g(\mathbf{S})^2] &= \mathbf{E}[\exp(-2\lambda(\theta n - \mathbf{S}))] = \exp(-2\lambda\theta n) \cdot (q + pe^{2\lambda})^n. \end{aligned}$$

Thus

$$\begin{aligned} \mathbf{Var}[g(\mathbf{S})] &= \mathbf{E}[g(\mathbf{S})]^2 \cdot \left(\frac{\mathbf{E}[g(\mathbf{S})^2]}{\mathbf{E}[g(\mathbf{S})]^2} - 1 \right) = \mathbf{E}[g(\mathbf{S})]^2 \cdot \left(\left(\frac{q + pe^{2\lambda}}{(q + pe^\lambda)^2} \right)^n - 1 \right) \\ &\leq \mathbf{E}[g(\mathbf{S})]^2 \cdot ((1 + 3pq\lambda^2)^n - 1) \quad (\text{Lemma 3.2}) \\ &\lesssim \mathbf{E}[g(\mathbf{S})]^2 \cdot pqn\lambda^2 \quad (\text{as } \lambda^2 \leq \frac{1}{pqn}) \end{aligned}$$

and it therefore remains to establish

$$\mathbf{E}[g(\mathbf{S})] = \exp(-\lambda\theta n) \cdot (q + pe^\lambda)^n \lesssim \mathbf{P}[B]. \quad (3)$$

Intuitively this holds because $g(s)$ should not be much different than $f(s)$, and $\mathbf{E}[f(\mathbf{S})] = \mathbf{P}[B]$ by definition. Formally, we consider two cases: $p \geq \frac{1}{n}$ (intuitively, the main case) and $p \leq \frac{1}{n}$.

Case 1: $p \geq \frac{1}{n}$. In this case we use that $\mathbf{P}[\mathbf{S} > pn] \geq \frac{1}{4}$ (see, e.g., [Doe18]), and hence: (i) it must be that $\theta \geq p$, since we are assuming $\mathbf{P}[B] = \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n] < \frac{1}{4}$; and, (ii) $\mathbf{P}[B] \geq \mathbf{P}[\mathbf{S} > pn] \cdot \mathbf{P}[\mathbf{X} \geq (\theta - p)n] \geq \frac{1}{4} \exp(-\lambda(\theta - p)n)$, where the first inequality used independence of \mathbf{S} and \mathbf{X} and the second inequality used $(\theta - p)n \geq 0$ (by (i)). Thus to establish **Inequality (3)**, it remains to show

$$\exp(-\lambda\theta n) \cdot (q + pe^\lambda)^n \lesssim \exp(-\lambda(\theta - p)n) \iff q + pe^\lambda \lesssim \exp(\lambda p).$$

But this last inequality indeed holds, as for $\lambda \leq 1$ we have $q + pe^\lambda \leq q + p(1 + 2\lambda) = 1 + 2p\lambda$, and $1 + 2p\lambda \leq 2 \exp(\lambda p)$ always.

Case 2: $p \leq \frac{1}{n}$. Here $q + pe^\lambda \leq 1 + 2p\lambda \leq 1 + \frac{2}{n}$, and so $(q + pe^\lambda)^n \lesssim 1$, meaning that **Inequality (3)** follows from $\mathbf{P}[B] \geq \mathbf{P}[\mathbf{X} > \theta n] = \exp(-\lambda\theta n)$. ■

3.1 The quantum version

Having established **Theorem 1.2**, we now show how this result applies to quantum states and measurements. Specifically, we prove that for any quantum event $A \in \mathbb{C}^{d \times d}$, there exists a corresponding event $B \in (\mathbb{C}^{d \times d})^{\otimes n}$ which, when measured against $\rho^{\otimes n}$, exhibits the same statistics as the classical event $\mathbf{S} + \mathbf{X} > \theta n$ from **Theorem 1.2** with $\mathbf{S} \sim \text{Binomial}(n, \text{tr}(\rho A))$. Moreover, we also relate the fidelity between the states $\rho^{\otimes n}$ and $\rho^{\otimes n}|_{\sqrt{\mathbb{1}-B}}$ (i.e. the state $\rho^{\otimes n}$ conditioned on the event $\mathbb{1} - B$) to the Bhattacharyya coefficient between \mathbf{S} and $(\mathbf{S} | \mathbf{S} + \mathbf{X} \leq \theta n)$ (i.e. \mathbf{S} conditioned on the event $\mathbf{S} + \mathbf{X} \leq \theta n$).

Lemma 3.4. *Let $\rho \in \mathbb{C}^{d \times d}$ represent an unknown quantum state and let $A \in \mathbb{C}^{d \times d}$ be a quantum event. Let $n \in \mathbb{N}$, let $\lambda > 0$, and let $\theta \in [0, 1]$ be an arbitrary threshold. Let \mathbf{S} and \mathbf{X} be classical random variables with distributions defined by $\mathbf{S} \sim \text{Binomial}(n, \mathbf{E}_\rho[A])$ and $\mathbf{X} \sim \text{Exponential}(\lambda)$. There exists a quantum event $B \in (\mathbb{C}^{d \times d})^{\otimes n}$ such that $\mathbf{E}_{\rho^{\otimes n}}[B] = \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n]$ and*

$$\mathbf{F}\left(\rho^{\otimes n}, \rho^{\otimes n}|_{\sqrt{\mathbb{1}-B}}\right) = \text{BC}((\mathbf{S} | \mathbf{S} + \mathbf{X} \leq \theta n), \mathbf{S}).$$

Proof. Let $\varrho = \rho^{\otimes n}$. By **Theorem 2.1**, we can assume, without loss of generality, that the quantum event A is a projector. Let $A_1 = A$ and $A_0 = \mathbb{1} - A$. For all $x \in \{0, 1\}^n$, let $A_x \in (\mathbb{C}^{d \times d})^{\otimes n}$ denote the event defined by $A_x = A_{x_1} \otimes A_{x_2} \otimes \cdots \otimes A_{x_n}$. For $k \in \{0, \dots, n\}$, let $E_k \in (\mathbb{C}^{d \times d})^{\otimes n}$ be the event defined by

$$E_k = \sum_{\substack{x \in \{0, 1\}^n \\ |x| = k}} A_x.$$

Since A is a projector, A_x is also a projector and $A_x A_y = A_y A_x = 0$ for all $x, y \in \{0, 1\}^n$ with $x \neq y$. Thus, each E_k is a sum of orthogonal projectors, so E_k is a projector as well and $E_k E_\ell = E_\ell E_k = 0$ for all $k, \ell \in \{0, \dots, n\}$ with $k \neq \ell$. Moreover,

$$\sum_{k=0}^n E_k = \sum_{x \in \{0, 1\}^n} A_x = \mathbb{1}.$$

Let $B \in (\mathbb{C}^{d \times d})^{\otimes n}$ denote the quantum event defined by

$$B = \sum_{k=0}^n \mathbf{P}[\mathbf{X} + k > \theta n] \cdot E_k.$$

The statistics of the measurement $\{E_k \mid k = 0, \dots, n\}$ applied to ϱ are distributed as $\text{Binomial}(n, \text{tr}(\rho A))$, so $\mathbf{E}_\varrho[E_k] = \mathbf{P}[\mathbf{S} = k]$. Hence,

$$\mathbf{E}_\varrho[B] = \sum_{k=0}^n \mathbf{P}[\mathbf{X} + k > \theta n] \cdot \mathbf{E}_\varrho[E_k] = \sum_{k=0}^n \mathbf{P}[\mathbf{X} + k > \theta n] \cdot \mathbf{P}[\mathbf{S} = k] = \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n].$$

For all $\ell \in \{0, \dots, n\}$,

$$\sqrt{\mathbb{1} - B} \cdot E_\ell = E_\ell \cdot \sqrt{\mathbb{1} - B} = \sqrt{\mathbf{P}[\mathbf{X} + \ell \leq \theta n]} \cdot E_\ell.$$

Hence,

$$\begin{aligned} \text{tr}(\varrho |_{\sqrt{\mathbb{1} - B}} \cdot E_\ell) &= \frac{1}{\mathbf{E}_\varrho[B]} \cdot \text{tr}(\sqrt{\mathbb{1} - B} \cdot \varrho \cdot \sqrt{\mathbb{1} - B} \cdot E_\ell) \\ &= \frac{1}{\mathbf{E}_\varrho[B]} \cdot \text{tr}(E_\ell \cdot \sqrt{\mathbb{1} - B} \cdot \varrho \cdot \sqrt{\mathbb{1} - B} \cdot E_\ell) \\ &= \frac{\mathbf{P}[\mathbf{X} + \ell \leq \theta n]}{\mathbf{E}_\varrho[B]} \cdot \text{tr}(E_\ell \cdot \varrho \cdot E_\ell) \\ &= \frac{\mathbf{P}[\mathbf{X} + \ell \leq \theta n]}{\mathbf{E}_\varrho[B]} \cdot \mathbf{E}_\varrho[E_\ell] \\ &= \frac{\mathbf{P}[\mathbf{X} + \ell \leq \theta n]}{\mathbf{P}[\mathbf{S} + \mathbf{X} \leq \theta n]} \cdot \mathbf{P}[\mathbf{S} = \ell]. \end{aligned}$$

Thus, the measurement $\{E_k \mid k = 0, \dots, n\}$ applied to $\varrho |_{\sqrt{\mathbb{1} - B}}$ yields statistics distributed as $(\mathbf{S} \mid \overline{B})$. Therefore, by [Proposition 2.3](#),

$$F(\varrho, \varrho |_{\sqrt{\mathbb{1} - B}}) = \sum_{k=0}^n \sqrt{\text{tr}(\varrho \cdot E_k)} \sqrt{\text{tr}(\varrho |_{\sqrt{\mathbb{1} - B}} \cdot E_k)} = \text{BC}((\mathbf{S} \mid \mathbf{S} + \mathbf{X} \leq \theta n), \mathbf{S}). \quad \blacksquare$$

Using [Lemma 3.4](#), we obtain the following “quantum version” of [Theorem 1.2](#):

Corollary 3.5. *Let $\rho \in \mathbb{C}^{d \times d}$ represent an unknown quantum state and let $A \in \mathbb{C}^{d \times d}$ be a quantum event. Let $n \in \mathbb{N}$, let $\lambda > 0$, and let $\theta \in [0, 1]$ be an arbitrary threshold. Fix $p = \mathbf{E}_\rho[A]$ and let \mathbf{S} and \mathbf{X} be defined as in [Theorem 1.2](#). If p, λ, n , and θ satisfy the conditions of [Theorem 1.2](#), then there exists a quantum event $B \in (\mathbb{C}^{d \times d})^{\otimes n}$ such that $\mathbf{E}_{\rho^{\otimes n}}[B] = \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n]$ and*

$$d_{\text{Bures}}\left(\rho^{\otimes n}, \rho^{\otimes n} |_{\sqrt{\mathbb{1} - B}}\right) \lesssim \mathbf{E}_{\rho^{\otimes n}}[B] \cdot \frac{\text{stddev}[\mathbf{S}]}{\mathbf{E}[\mathbf{X}]}.$$

Moreover,

$$\mathbf{E}_{\rho^{\otimes n}}[B] \leq \exp(-n\lambda(\theta - (e - 1)\mathbf{E}[A])).$$

Proof. Let $\varrho = \rho^{\otimes n}$. By [Lemma 3.4](#), there exists a quantum event $B \in (\mathbb{C}^{d \times d})^{\otimes n}$ such that $\mathbf{E}_\varrho[B] = \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n]$ and $F(\varrho, \varrho|_{\sqrt{1-B}}) = \text{BC}((\mathbf{S} | \mathbf{S} + \mathbf{X} \leq \theta n), \mathbf{S})$. Note that, for all distributions μ and ν , $1 - \text{BC}(\mu, \nu) \leq d_{\chi^2}(\mu, \nu)$. Hence, by [Lemma 3.4](#), it follows that

$$\begin{aligned} d_{\text{Bures}}(\rho^{\otimes n}, \rho^{\otimes n}|_{\sqrt{1-B}}) &= \sqrt{2(1 - F(\varrho, \varrho|_{\sqrt{1-B}}))} \\ &= \sqrt{2(1 - \text{BC}((\mathbf{S} | \mathbf{S} + \mathbf{X} \leq \theta n), \mathbf{S}))} \\ &= d_{\text{H}}((\mathbf{S} | \mathbf{S} + \mathbf{X} \leq \theta n), \mathbf{S}) \\ &\leq \sqrt{d_{\chi^2}((\mathbf{S} | \mathbf{S} + \mathbf{X} \leq \theta n), \mathbf{S})} \\ &\lesssim \mathbf{E}_{\rho^{\otimes n}}[B] \cdot \frac{\text{stddev}[\mathbf{S}]}{\mathbf{E}[\mathbf{X}]} \quad (\text{by } \text{Corollary 1.3}). \end{aligned}$$

Since $\mathbf{E}_\varrho[B] = \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n]$,

$$\begin{aligned} \mathbf{E}_\varrho[B] &= \mathbf{P}[\mathbf{S} + \mathbf{X} > \theta n] \\ &\leq \mathbf{E}[\exp(-\lambda(\theta n - \mathbf{S}))] && (\text{by } \mathbf{P}[\mathbf{X} > t] \leq \exp(-\lambda t)) \\ &= \exp(-\lambda \theta n) \mathbf{E}[\exp(\lambda \mathbf{S})] \\ &= \exp(-\lambda \theta n)(1 - \text{tr}(\rho A) + \text{tr}(\rho A)e^\lambda)^n && (\mathbf{E}[\exp(\lambda \mathbf{S})] \text{ is the m.g.f. of } \mathbf{S}) \\ &= \exp(-\lambda \theta n)(1 + \text{tr}(\rho A)(e^\lambda - 1))^n \\ &\leq \exp(-\lambda \theta n)(1 + \text{tr}(\rho A)(e - 1)\lambda)^n && (\text{by } e^x \leq 1 + (e - 1)x \text{ for } x \in [0, 1]) \\ &\leq \exp(-\lambda \theta n) \exp((e - 1)n\lambda \text{tr}(\rho A)) && (\text{by } 1 + x \leq e^x \text{ for } x \in \mathbb{R}) \\ &= \exp(-n\lambda(\theta - (e - 1)\text{tr}(\rho A))). \end{aligned}$$

4 Threshold Search

In this section, we prove [Theorem 1.1](#).

4.1 Preliminary reductions

We begin with several reductions that allow us to reduce to the case of projectors, and to the case when ϵ , δ , and the θ_i 's are all fixed constants.

Reduction to projectors. Let $\rho \in \mathbb{C}^{d \times d}$ denote the unknown quantum state and let $A_1, \dots, A_m \in \mathbb{C}^{d \times d}$ be the observables in the quantum Threshold Search problem (which we assume are given in an online fashion). Note that this problem depends only on the probabilities $\mathbf{E}_\rho[A_i]$ with $i = 1, \dots, m$. Thus if we extend the unknown state ρ to $\rho \otimes |0\rangle\langle 0|$, then by Naimark's [Theorem 2.1](#), there exists a projector $\Pi_i \in \mathbb{C}^{d \times d} \otimes \mathbb{C}^{2 \times 2}$ for each A_i such that $\mathbf{E}_{\rho \otimes |0\rangle\langle 0|}[\Pi_i] = \mathbf{E}_\rho[A_i]$ for all $i = 1, \dots, m$. Since the state $\rho \otimes |0\rangle\langle 0|$ can be prepared without knowing ρ and this extension increases the dimension of the quantum system only by a constant factor, by replacing ρ by $\rho \otimes |0\rangle\langle 0|$ and each A_i by the corresponding Π_i , it follows that we can assume, without loss of generality, that the observables A_1, \dots, A_m are projectors.

Reduction to 3/4 vs. 1/4. Let $0 < \epsilon < \frac{1}{2}$ be given, and recall that in the Threshold Search problem the algorithm is presented with a stream of projector/threshold pairs (A_i, θ_i) , with the goal of distinguishing the cases $\mathbf{E}_\rho[A_i] > \theta_i$ and $\mathbf{E}_\rho[A_i] \leq \theta_i - \epsilon$. We may have the algorithm use [Lemma 2.4](#) (the latter part, with $\tau = 0$, $c = \theta_i - \epsilon/2$, $\delta = 1/4$, and ϵ replaced by $\epsilon/2$), which establishes that for some $n_0 = O(1/\epsilon^2)$, each A_i may be replaced with a projector $B_i \in (\mathbb{C}^{d \times d})^{\otimes n_0}$ satisfying

- i. if $\mathbf{E}_\rho[A_i] > \theta_i$, then $\mathbf{E}_{\rho^{\otimes n_0}}[B_i] > 3/4$;
- ii. if $\mathbf{E}_\rho[A_i] \leq \theta_i - \epsilon$, then $\mathbf{E}_{\rho^{\otimes n_0}}[B_i] \leq 1/4$.

Thus we can reduce to the “3/4 vs. 1/4” version of Threshold Search at the expense of paying an extra factor of $n_0 = O(1/\epsilon^2)$ in the copy complexity. Note that the parameter d has increased to d^{n_0} , as well, but (crucially) our [Theorem 1.1](#) has no dependence on the dimension parameter.

Reduction to a promise-problem version, with fixed δ . So far we have reduced proving [Theorem 1.1](#) to proving the following:

Theorem 4.1. *There is an algorithm that, given $m \in \mathbb{N}$ and $0 < \delta < \frac{1}{2}$, first obtains $n^* = O(\log^2 m + \log(1/\delta)) \cdot \log(1/\delta)$ copies $\rho^{\otimes n^*}$ of an unknown state $\rho \in \mathbb{C}^{d \times d}$. Next, a sequence of projectors $A_1, \dots, A_m \in \mathbb{C}^{d \times d}$ is presented to the algorithm (possibly adaptively). After each A_t , the algorithm may either select t , meaning halt and output the claim “ $\mathbf{E}_\rho[A_t] > 1/4$ ”, or else pass to the next projector. If the algorithm passes on all m projectors, the algorithm must claim “ $\mathbf{E}_\rho[A_i] \leq 3/4$ for all i ”. Except with probability at most δ , the algorithm’s output is correct.*

The main work we will do is to show the following similar result:

Lemma 4.2. *There is an algorithm that, given $m \in \mathbb{N}$, first obtains $n = O(\log^2 m)$ copies $\rho^{\otimes n}$ of an unknown state $\rho \in \mathbb{C}^{d \times d}$. Next, a sequence of projectors $A_1, \dots, A_m \in \mathbb{C}^{d \times d}$, obeying the promise that $\mathbf{E}_\rho[A_j] > 3/4$ for at least one j , is presented to the algorithm. After each A_t , the algorithm may either halt and select t , or else pass to the next projector. With probability at least 0.1, the algorithm selects a t with $\mathbf{E}_\rho[A_t] \geq 1/3$.*

One needs a slight bit of care to reduce [Theorem 4.1](#) to [Lemma 4.2](#) while maintaining the online nature of the algorithm:

Proof of [Theorem 4.1](#), assuming [Lemma 4.2](#). We will use the algorithm in [Lemma 4.2](#) as a kind of “subroutine” for the main theorem. Our first step is to augment this subroutine in the following way:

- Given parameter δ for the main theorem, the subroutine will use a parameter $\delta' = \delta/(C \log(1/\delta))$, where C is a universal constant to be chosen later.
- n is increased from $O(\log^2 m)$ to $n' = O(\log^2 m) + O(\log(1/\delta'))$, where the first $O(\log^2 m)$ copies of ρ are used as usual, and the additional $O(\log(1/\delta'))$ copies are reserved as a “holdout”.

- If ever the subroutine is about to halt and select t , it first performs a “failsafe” check: It applies [Lemma 2.4](#) with $\tau = 0$, $c = .3$, $\epsilon = .03$, $\delta = \delta'$, and measures with the holdout copies. (Note that $c + \epsilon < 1/3$ and also $c - \epsilon > 1/4$.) If “ B ” occurs, the subroutine goes ahead and selects t ; otherwise, the algorithm not only passes, but it “aborts”, meaning that it automatically passes on all subsequent A_i ’s without considering them.

We make two observations about this augmented subroutine:

- When run under the promise that $\mathbf{E}_\rho[A_j] > 3/4$ for at least one j , it still selects a t satisfying $\mathbf{E}_\rho[A_t] \geq 1/3$ with probability at least 0.05. This is because the “failsafe” causes an erroneous change of mind with probability at most δ' , and we may assume $\delta' \leq 0.05$ (taking C large enough).
- When run *without* the promise that $\mathbf{E}_\rho[A_j] > 3/4$ for at least one j , the failsafe implies that the probability the algorithm ever selects a t with $\mathbf{E}_\rho[A_t] < 1/4$ is at most δ' .

With the augmented subroutine in hand, we can now give the algorithm that achieves [Theorem 4.1](#). The main algorithm will obtain $n^* = n' \cdot L$ copies of ρ , where $L = O(\log(1/\delta))$; these are thought of as L “batches”, each with of n' copies. As the projectors A_i are presented to the main algorithm, it will run the augmented subroutine “in parallel” on each batch. If any batch wants to halt accept a certain A_t , then the overall algorithm halts and outputs “ $\mathbf{E}_\rho[A_t] > 1/4$ ”. Otherwise, if all the batches pass on A_t , so too does the overall algorithm. Of course, if the overall algorithm passes on all A_i ’s, it outputs “ $\mathbf{E}_\rho[A_i] \leq 3/4$ for all i ”.

We now verify the correctness of this main algorithm. First, *if* there exists some A_j with $\mathbf{E}_\rho[A_j] > 3/4$, the probability of the algorithm wrongly outputting “ $\mathbf{E}_\rho[A_i] \leq 3/4$ for all i ” is at most $(1 - .05)^L$, which can be made smaller than δ by taking the hidden constant in $L = O(\log(1/\delta))$ suitably large. On the other hand, thanks to the “failsafe” and a union bound, the probability the algorithm ever wrongly outputs “ $\mathbf{E}_\rho[A_t] > 1/4$ ” is at most $L\delta' = L \cdot \delta / (C \log(1/\delta))$, which is again at most δ provided C is taken large enough. ■

4.2 The main algorithm (proof of [Lemma 4.2](#))

In this section, we will prove [Lemma 4.2](#). Let $n = n(m)$ and $\lambda = \lambda(m)$ be parameters to be fixed later and let $\theta = 2/3$. As stated in [Lemma 4.2](#), we may explicitly assume there exists $i \in [m]$ with $\mathbf{E}_\rho[A_i] \geq 3/4$. For each projector A_i , let B_i denote the event obtained from [Lemma 3.4](#). The algorithm proceeds as follows:

Let ϱ denote the current quantum state, with $\varrho = \rho^{\otimes n}$ initially. Given projector A_i , let B_i be the event obtained from [Lemma 3.4](#). Measure the current state ϱ with (\overline{B}_i, B_i) using the canonical implementation. If B_i occurs, halt and select i ; otherwise, pass.

Note that the n copies of ρ are only prepared once and reused, and that the current state ϱ collapses to a new state after each measurement.

The algorithm has the following modes of failure:

(FN) the algorithm passes on every observable because the event \overline{B}_i occurs for every $i \in [m]$;

(FP) the algorithm picks an observable A_j with $\mathbf{E}_\rho[A_j] < 1/3$.

We want to show that the algorithm does not make errors of type FP or FN with probability at least 0.1. To this end, we introduce the following notation.

Notation 4.3. For $i = 1, \dots, m$, let:

1. \mathbf{S}_i be a random variable distributed as $\text{Binomial}(n, \mathbf{E}_\rho[A_i])$;
2. $p_i = \mathbf{E}_{\rho^{\otimes n}}[B_i]$ be the probability that B_i would occur if $\rho^{\otimes n}$ were measured with (\overline{B}_i, B_i) ;
3. $\varrho_0 = \rho^{\otimes n}$ and let ϱ_i be the quantum state after the i th measurement, *conditioned* on the event \overline{B}_j occurring for all $1 \leq j \leq i$;
4. $r_i = \mathbf{E}_{\varrho_{i-1}}[\overline{B}_i]$ be the probability that the event \overline{B}_i occurs assuming all the events \overline{B}_j with $1 \leq j \leq i-1$ occurred;
5. $q_i = r_1 \cdots r_i$ be the probability that *all* of the events \overline{B}_j with $1 \leq j \leq i$ occur.

Note that the p_i 's refer to a ‘‘hypothetical,’’ whereas the r_i 's and q_i 's concern what actually happens over the course of the algorithm. In particular, q_m is the probability that the algorithm passes on every observable. The following claim shows that, as long as the noise expectation $\mathbf{E}[\mathbf{X}] = 1/\lambda$ used in [Lemma 3.4](#) is sufficiently large, the probability of a false negative (FN) is bounded above by 4/5:

Claim 4.4. *For $\mathbf{E}[\mathbf{X}] = \Omega(\sqrt{n})$, there exists $t \in [m]$ such that $q_t \leq 4/5$. Moreover, if $t > 1$, then $q_{t-1} \geq 3/4$ and $p_1 + \cdots + p_{t-1} \leq 1/4$.*

Proof. By [Lemma 3.4](#), $p_i = \mathbf{E}_{\rho^{\otimes n}}[B_i] = \mathbf{P}[\mathbf{S}_i + \mathbf{X} > \theta n]$. Let $k \in [m]$ be such that $\mathbf{E}_\rho[A_k] \geq 3/4$. Thus, \mathbf{S}_k is a binomial random variable with mean at least $3/4$. Since $\theta = 2/3 < 3/4$, for $n = \Omega(1)$ a sufficiently large constant,

$$p_k = \mathbf{P}[\mathbf{S}_k + \mathbf{X} > \theta n] \geq \mathbf{P}[\mathbf{S}_k > (2/3)n] \geq 1 - \exp(-1/4).$$

Therefore, there exists a minimal $t \in [m]$ such that $(1 - p_1) \cdots (1 - p_t) \leq \exp(-1/4)$. If $t = 1$, then $q_1 = 1 - p_1 \leq \exp(-1/4) \leq 4/5$. Otherwise, since t is minimal, it follows that $(1 - p_1) \cdots (1 - p_{t-1}) \geq \exp(-1/4)$. Hence,

$$\exp(-1/4) \leq (1 - p_1) \cdots (1 - p_{t-1}) \leq \exp(-(p_1 + \cdots + p_{t-1})),$$

whence $p_1 + \cdots + p_{t-1} \leq 1/4$. Thus, by [Lemma 2.5](#) and [Corollary 3.5](#),

$$\begin{aligned} |(1 - p_1) \cdots (1 - p_t) - q_t| &\leq 2 \sum_{i=1}^{t-1} d_{\text{tr}}\left(\rho^{\otimes n}, \rho^{\otimes n} \Big|_{\sqrt{1-B_i}}\right) \\ &\lesssim \sum_{i=1}^{t-1} \mathbf{E}_{\rho^{\otimes n}}[B_i] \cdot \frac{\text{stddev}[\mathbf{S}_i]}{\mathbf{E}[\mathbf{X}]} \leq \frac{\sqrt{n}}{\mathbf{E}[\mathbf{X}]} \cdot (p_1 + \cdots + p_{t-1}) \leq \frac{1}{4} \cdot \frac{\sqrt{n}}{\mathbf{E}[\mathbf{X}]} \end{aligned}$$

By a similar argument,

$$|(1 - p_1) \cdots (1 - p_{t-1}) - q_{t-1}| \lesssim \frac{\sqrt{n}}{\mathbf{E}[\mathbf{X}]} \cdot (p_1 + \cdots + p_{t-2}) \leq \frac{1}{4} \cdot \frac{\sqrt{n}}{\mathbf{E}[\mathbf{X}]}.$$

Therefore, for $\mathbf{E}[\mathbf{X}] = \Omega(\sqrt{n})$, $q_t \leq 4/5$ and $q_{t-1} \geq 3/4$. \blacksquare

Assuming $\mathbf{E}[\mathbf{X}] = \Omega(\sqrt{n})$, let $t \in [m]$ be as in [Claim 4.4](#). Since $q_m \leq q_t \leq 4/5$, it follows that the probability the algorithm makes an FN error is at most $4/5$. In fact, since $q_t \leq 4/5$, the algorithm will pick an index $i \leq t$ with probability at least $1/5$. Thus, to show that the algorithm succeeds w.p. at least 0.1 , it suffices to show that w.h.p. the algorithm does not pick an index $i \in \mathcal{B}$, where $\mathcal{B} \subseteq [m]$ is the subset defined by

$$\mathcal{B} = \{i \in [m] \mid 1 \leq i \leq t \text{ and } \mathbf{E}_\rho[A_i] < 1/3\}.$$

First, we show that an event B_i with $i \in \mathcal{B}$ is unlikely to occur when the initial state $\rho^{\otimes n}$ is measured against (\overline{B}_i, B_i) :

Claim 4.5. *If $n = O(\log^2(m/\eta))$, then $p_i \leq \eta/m$ for all $i \in \mathcal{B}$.*

Proof. By [Corollary 3.5](#), for all $i \in [m]$,

$$p_i = \mathbf{E}_{\rho^{\otimes n}}[B_i] \leq \exp(-n\lambda(\theta - (e-1)\mathbf{E}_\rho[A_i])).$$

Since $\theta = 2/3$ and $i \in \mathcal{B}$, then $\mathbf{E}_\rho[A_i] < 1/3$ and $\theta - (e-1)\mathbf{E}_\rho[A_i] \geq 0.09$. Since $n\lambda = O(\sqrt{n})$ and $n = O(\log^2(m/\eta))$, it follows that

$$p_i = \mathbf{E}_{\rho^{\otimes n}}[B_i] \leq \exp(-0.09\lambda n) \lesssim \frac{\eta}{m}, \quad \text{for all } i \in \mathcal{B}. \quad \blacksquare$$

The following claim shows that $r_{\mathcal{B}}$, i.e. the probability that the algorithm rejects all $i \in \mathcal{B}$, is sufficiently high:

Claim 4.6. *For $n = O(\log^2(m/\eta))$, $r_{\mathcal{B}} \geq \exp(-\eta) - 0.01$ if $t > 1$.*

Proof. Since $r_1 \cdots r_{t-1} = q_{t-1}$, it follows that $r_{[t-1]\setminus\mathcal{B}} \geq q_{t-1} \geq 3/4$. Hence, by [Lemma 2.5](#),

$$\left| \prod_{i \in \mathcal{B}} (1 - p_i) - r_{\mathcal{B}} \right| \leq \frac{2}{r_{[t-1]\setminus\mathcal{B}}} \sum_{i=1}^{t-1} d_{\text{tr}}(\rho^{\otimes n}, \rho^{\otimes n}|_{\sqrt{1-B_i}}) \lesssim \frac{\sqrt{n}}{\mathbf{E}[\mathbf{X}]}.$$

Since

$$\prod_{i \in \mathcal{B}} (1 - p_i) \geq \prod_{i \in \mathcal{B}} \left(1 - \frac{\eta}{m}\right) \geq \left(1 - \frac{\eta}{m}\right)^m,$$

it follows that $r_{\mathcal{B}} \geq \exp(-\eta) - 0.01$ for $m = \Omega(1)$ and $\mathbf{E}[\mathbf{X}] = \Omega(\sqrt{n})$ sufficiently large. \blacksquare

Thus, we can lower bound the probability that the algorithm succeeds in picking an index $i \in [m]$ with $\mathbf{E}_\rho[A_i] \geq 1/3$, proving [Lemma 4.2](#):

Proof of Lemma 4.2. Fix $\eta = 0.01$, so that indeed $n = O(\log^2 m)$ as promised. If $t = 1$, then $q_1 = 1 - p_1 \leq 4/5$. Hence, $p_1 \geq 1/5$. If $1 \in \mathcal{B}$, then, by Claim 4.5, $p_1 \leq \eta/m < 1/5$, which is a contradiction. Therefore, the algorithm picks the index $1 \in [m]$ with probability at least $p_1 \geq 1/5$ and $1 \notin \mathcal{B}$, i.e. $\mathbf{E}_\rho[A_1] \geq 1/3$.

Otherwise, suppose $t > 1$. By Claim 4.6, it follows that

$$0.98 \leq \exp(-\eta) - 0.01 \leq r_{\mathcal{B}} = \exp\left(-\sum_{i \in \mathcal{B}} 1 - r_i\right).$$

Recall that $1 - r_i$ is the probability that the algorithm selects index i . Hence, the algorithm picks an index $i \in \mathcal{B}$ with probability at most

$$\sum_{i \in \mathcal{B}} 1 - r_i \leq -\log(0.98) \leq 0.1.$$

Therefore, the algorithm picks an index $i \leq t$ with $\mathbf{E}_\rho[A_i] \geq 1/3$ with probability at least $1 - q_t - 0.1 \geq 0.1$. \blacksquare

5 Shadow Tomography and Hypothesis Selection

5.1 Shadow Tomography

We begin by describing how to deduce our online Shadow Tomography result, Theorem 1.4, from our online Threshold Search result, Theorem 1.1. As mentioned earlier, this deduction is known [Aar20] to follow almost immediately from a mistake-bounded learning algorithm for quantum states due to Aaronson, Chen, Hazan, Kale, and Nayak [ACH⁺19], described below. We will fill in a few details that are not spelled out in [Aar20].

Mistake-bounded learning scenario. Consider the following kind of interaction between a “student” and a “teacher”, given parameters $d \in \mathbb{N}$ and $0 < \epsilon < \frac{1}{2}$. There is a quantum state $\rho \in \mathbb{C}^{d \times d}$ that is unknown to the student (and possibly also unknown to the teacher). The teacher presents a sequence of quantum events A_1, A_2, A_3, \dots (possibly adaptively) to the student. Upon receiving A_t , the student must output a prediction $\hat{\mu}_t$ of $\mu_t = \mathbf{E}_\rho[A_t]$. After each prediction, the teacher must either “pass”, or else declare a “mistake” and supply a value μ'_t .

Theorem 5.1 ([ACH⁺19]). *Assume the following **Teacher Properties** hold for each t :*

- If $|\hat{\mu}_t - \mu_t| > \epsilon$, the teacher always declares “mistake”.
- If $|\hat{\mu}_t - \mu_t| \leq \frac{3}{4}\epsilon$, the teacher always passes.
- If the teacher ever declares “mistake”, the supplied value μ'_t always satisfies $|\mu'_t - \mu_t| \leq \frac{1}{4}\epsilon$.

Then there is an algorithm for the student that causes at most $C_0(\log d)/\epsilon^2$ “mistakes” (no matter how many events are presented), where C_0 is a universal constant.

The above theorem is similar to, but not quite the same, as “Theorem 1” in [ACH⁺19]. However it is easy to check that [ACH⁺19]’s Section 3.3 (“Proof of Theorem 1”) applies equally well to establish [Theorem 5.1](#) above.⁴

To use this theorem for the online Shadow Tomography problem, it only remains for the Shadow Tomography algorithm *to implement the teacher’s role itself*, given copies of ρ . This will be done using our Threshold Search algorithm; let us first slightly upgrade it so that (i) it is concerned with $\mathbf{E}_\rho[A_i] \approx \theta_i$ rather than $\mathbf{E}_\rho[A_i] < \theta_i$; (ii) if it finds j with $\mathbf{E}_\rho[A_j] \not\approx \theta_j$, then it also reports a very good estimate of $\mathbf{E}_\rho[A_j]$.

Lemma 5.2. *Consider the version of quantum Threshold Search where the inputs are the same, but the algorithm should correctly (except with probability at most δ) output:*

- “ $|\mathbf{E}_\rho[A_j] - \theta_j| > \frac{3}{4}\epsilon$, and in fact $|\mathbf{E}_\rho[A_j] - \mu'_j| \leq \frac{1}{4}\epsilon$ ”, for some particular j and value μ'_j ; or else,
- “ $|\mathbf{E}_\rho[A_i] - \theta_i| \leq \epsilon$ for all i ”.

Then as in [Theorem 1.1](#), the problem can be solved in an online fashion using

$$n'_{\text{TS}}(m, \epsilon, \delta) = \frac{\log^2 m + L}{\epsilon^2} \cdot O(L) \quad (L = \log(1/\delta))$$

copies of ρ .

Proof. Given m, ϵ, δ , we obtain $n = n_{\text{TS}}(2m, \frac{1}{4}\epsilon, \delta/2) + c \log(1/\delta)/\epsilon^2$ copies of ρ , where c is a universal constant to be specified later. This n indeed has the asymptotic form of n'_{TS} given above. We save the $c \log(1/\delta)/\epsilon^2$ copies as a “holdout”, and use the remaining copies to apply [Theorem 1.1](#) (with parameters $2m, \frac{1}{4}\epsilon, \delta/2$), converting our given observable/threshold pairs $(A_1, \theta_1), \dots, (A_m, \theta_m)$ to a “simulated input” of

$$(A_1, \theta_1 + \epsilon), (\mathbb{1} - A_1, 1 - \theta_1 + \epsilon), \dots, (A_m, \theta_m + \epsilon), (\mathbb{1} - A_m, 1 - \theta_m + \epsilon).$$

Except with probability at most $\delta/2$ we get a correct answer from the simulation, from which we can derive a correct final output as described below.

If the simulation passes on all $2m$ pairs, then [Theorem 1.1](#) tells us that we must have

$$\mathbf{E}_\rho[A_i] \leq \theta_i + \epsilon \quad \text{and} \quad \mathbf{E}_\rho[\mathbb{1} - A_i] \leq 1 - \theta_i + \epsilon$$

for all i , and therefore we may correctly output “ $|\mathbf{E}_\rho[A_i] - \theta_i| \leq \epsilon$ for all i ”.

On the other hand, suppose the simulation halts by outputting

$$\text{“}\mathbf{E}_\rho[A_j] > \theta_j + \epsilon - \frac{1}{4}\epsilon\text{”} \quad \text{or} \quad \text{“}\mathbf{E}_\rho[\mathbb{1} - A_j] > 1 - \theta_j + \epsilon - \frac{1}{4}\epsilon\text{”}$$

for some particular j . Then our algorithm can correctly output “ $|\mathbf{E}_\rho[A_j] - \theta_j| > \frac{3}{4}\epsilon$ ”. Furthermore, at this point the algorithm may use the holdout copies of ρ to obtain an estimate μ'_j of $\mathbf{E}_\rho[A_j]$ (in the naive way) that satisfies $|\mathbf{E}_\rho[A_j] - \mu'_j| \leq \frac{1}{4}\epsilon$ except with probability at most $\delta/2$, provided c is large enough. ■

⁴Briefly: the RTFL/MMW algorithm will only do an update in the “mistake” rounds. The loss is taken to be $|\hat{\mu}_t - \mu'_t|$. On any mistake, we have $|\hat{\mu}_t - \mu_t| > \frac{3}{4}\epsilon$ and $|\mu'_t - \mu_t| \leq \frac{1}{4}\epsilon$, hence the student incurs loss at least $\frac{1}{2}\epsilon$. On the other hand, answering according to the true μ_t would only incur loss at most $\frac{1}{4}\epsilon$. The regret calculation bounding the number of mistakes is now the same.

With [Lemma 5.2](#) in place, we can obtain our online Shadow Tomography algorithm:

Proof of [Theorem 1.4](#). Define

$$R = \lceil C_0(\log d)/\epsilon^2 \rceil + 1, \quad \delta_0 = \delta/R, \quad n_0 = n'_{\text{TS}}(m, \epsilon, \delta_0).$$

The number of copies of ρ used by our online Shadow Tomography algorithm will be $n = Rn_0$, which is indeed

$$n = \frac{(\log^2 m + L)(\log d)}{\epsilon^4} \cdot O(L)$$

for $L = \log(\frac{\log d}{\delta\epsilon})$, as claimed.

Upon receiving n copies of ρ , our Shadow Tomography algorithm partitions it into R “batches” of size n_0 each. The idea is that each batch will be devoted to (up to) one “mistake” of the “student”. We now describe the algorithm, and then give its analysis.

To begin, recall that our Shadow Tomography algorithm receives quantum events A_1, A_2, \dots in an online fashion. As it receives them it will, in parallel, run:

- the mistake-bounded learning algorithm of [Theorem 5.1](#) (implementing the student’s algorithm);
- the Threshold Search algorithm from [Lemma 5.2](#) (to implement the teacher), initially using only the first batch of $\rho^{\otimes n_0}$.

The estimates $\hat{\mu}_t$ output by the student will serve as the θ_t values used in [Lemma 5.2](#).

Whenever the Threshold Search algorithm “passes” on an $(A_t, \hat{\mu}_t)$ pair, the teacher also “passes”, and $\hat{\mu}_t$ serves as the Shadow Tomography algorithm’s final estimate for $\mathbf{E}_\rho[A_t]$. On the other hand, if the Threshold Search algorithm outputs “ $|\mathbf{E}_\rho[A_t] - \hat{\mu}_t| > \frac{3}{4}\epsilon$, and in fact $|\mathbf{E}_\rho[A_t] - \mu'_t| \leq \frac{1}{4}\epsilon$ ”, then the teacher will declare a “mistake” and supply the value μ'_t to the student. This μ'_t will also serve as the Shadow Tomography algorithm’s final estimate for $\mathbf{E}_\rho[A_t]$. Furthermore, at this point the teacher will abandon any remaining copies of ρ in the current batch, and will use a “fresh” batch $\rho^{\otimes n_0}$ for the subsequent application of [Lemma 5.2](#). We refer to this as moving on to the next “round”.

Let us now show that with high probability there are at most $R - 1$ mistakes and hence at most R rounds. (If the Shadow Tomography algorithm tries to proceed to an $(R + 1)$ th round, and thereby runs out of copies of ρ , we simply declare an overall failure.) By a union bound, in the first R applications of [Lemma 5.2](#), the probability of *any* incorrect answer is at most $R\delta_0 = \delta$. Below we will show that if there are no incorrect answers, then the “Teacher Properties” of [Theorem 5.1](#) hold, and therefore the total number of mistakes is indeed at most $\lceil C_0(\log d)/\epsilon^2 \rceil = R - 1$ with probability at least $1 - \delta$.

It remains to verify that — assuming correct answers from all uses of [Lemma 5.2](#) — our Shadow Tomography algorithm satisfies the Teacher Properties of [Theorem 5.1](#) and also that all m estimates for $\mathbf{E}_\rho[A_i]$ produced by the algorithm are correct to within an additive ϵ . Let us first observe that in each application of [Lemma 5.2](#), we will never supply more than m quantum events. The main thing to observe is that if our Threshold Search routine from [Lemma 5.2](#) ever passes on some $(A_t, \hat{\mu}_t)$ pair, it *must* be that $|\mathbf{E}_\rho[A_t] - \hat{\mu}_t| \leq \epsilon$; the reason is that passing implies the Threshold Search algorithm is prepared to output “ $|\mathbf{E}_\rho[A_i] - \theta_i| \leq \epsilon$ for all i ”. On the other hand, it’s immediate from [Lemma 5.2](#) that if the teacher declares

“mistake” on some $(A_t, \hat{\mu}_t)$ pair, then indeed we have $|\mathbf{E}_\rho[A_t] - \hat{\mu}_t| > \frac{3}{4}\epsilon$, and the supplied correction μ'_t satisfies $|\mathbf{E}_\rho[A_t] - \mu'_t| \leq \frac{1}{4}\epsilon$ (as is necessary for the Teacher Properties, and is more than sufficient for the Shadow Tomography guarantee). ■

5.2 Hypothesis Selection

In this section we establish our quantum Hypothesis Selection result, [Theorem 1.5](#). This theorem effectively has three different bounds, and we prove them via [Propositions 5.3, 5.5](#) and [5.6](#).

Recall that in the Hypothesis Selection problem there are given fixed hypothesis states $\sigma_1, \dots, \sigma_m \in \mathbb{C}^{d \times d}$, as well as access to copies of an unknown state $\rho \in \mathbb{C}^{d \times d}$. We write

$$\eta = \min_i \{d_{\text{tr}}(\rho, \sigma_i)\}, \quad i^* = \operatorname{argmin}_i \{d_{\text{tr}}(\rho, \sigma_i)\},$$

with the quantity η being unknown to the algorithm. Recall that Helstrom’s theorem implies that for each pair $i \neq j$, there is a quantum event A_{ij} such that

$$\mathbf{E}_{\sigma_i}[A_{ij}] - \mathbf{E}_{\sigma_j}[A_{ij}] = d_{\text{tr}}(\sigma_i, \sigma_j),$$

and furthermore we may take $A_{ji} = \overline{A_{ij}} = \mathbb{1} - A_{ij}$. These events *are* known to the algorithm.

One way to solve the quantum Hypothesis Selection problem is to simply use Shadow Tomography as a black box. Given parameters $0 < \epsilon, \delta < \frac{1}{2}$ for the former problem, we can run Shadow Tomography with parameters $\epsilon/2, \delta$, and the $\binom{m}{2}$ quantum events $(A_{ij} : i < j)$. Then except with probability at most δ , we obtain values $\hat{\mu}_{ij}$ with $|\mathbf{E}_\rho[A_{ij}] - \hat{\mu}_{ij}| \leq \epsilon/2$ for all i, j . Now we can essentially use any classical Hypothesis Selection algorithm; e.g., the “minimum distance estimate” method of Yatracos [[Yat85](#)]. We select as our hypothesis σ_k , where k is a minimizer of

$$\hat{\Delta}_k = \max_{i < j} \max_{\sigma_k} |\mathbf{E}[A_{ij}] - \hat{\mu}_{ij}|.$$

Recalling $\eta = d_{\text{tr}}(\rho, \sigma_{i^*})$, we have

$$\hat{\Delta}_k \leq \hat{\Delta}_{i^*} \leq \max_{i < j} \{|\mathbf{E}_{\sigma_{i^*}}[A_{ij}] - \mathbf{E}_\rho[A_{ij}]| + |\mathbf{E}_\rho[A_{ij}] - \hat{\mu}_{ij}|\} \leq \eta + \epsilon/2, \quad (4)$$

where the last inequality used the Helstrom theorem again, and the Shadow Tomography guarantee. We now obtain the following result (with the proof being an almost verbatim repeat of the one in [[DL01](#), Thm. 6.3]):

Proposition 5.3. *The above-described method selects σ_k with $d_{\text{tr}}(\rho, \sigma_k) \leq 3\eta + \epsilon$ (except with probability at most δ), using a number of copies of ρ that is the same as in Shadow Tomography (up to constant factors).*

Proof. By the triangle inequality for d_{tr} we have

$$\begin{aligned} d_{\text{tr}}(\sigma_k, \rho) &\leq \eta + d_{\text{tr}}(\sigma_k, \sigma_{i^*}) = \eta + |\mathbf{E}_{\sigma_k}[A_{ki^*}] - \mathbf{E}_{\sigma_{i^*}}[A_{ki^*}]| \\ &\leq \eta + |\mathbf{E}_{\sigma_k}[A_{ki^*}] - \hat{\mu}_{ki^*}| + |\mathbf{E}_{\sigma_{i^*}}[A_{ki^*}] - \hat{\mu}_{ki^*}| \leq \eta + \hat{\Delta}_k + \hat{\Delta}_{i^*}, \end{aligned}$$

and the result now follows from [Inequality \(4\)](#). ■

Now we give a different, incomparable method for Hypothesis Selection. It will use the following “decision version” of quantum Threshold Search, which we prove at the end of [Appendix A](#) (see [Corollary A.4](#)):

Corollary 5.4. *Consider the scenario of quantum Threshold Search (i.e., one is given parameters $0 < \epsilon_0, \delta_0 < \frac{1}{2}$, and m_0 event/threshold pairs (A_i, θ_i)). Suppose one is further given values η_1, \dots, η_m . Then using just $n_0 = O(\log(m_0/\delta_0)/\epsilon_0^2)$ copies of ρ , one can correctly output (except with probability at most δ):*

- “there exists j with $|\mathbf{E}_\rho[A_j] - \theta_j| > \eta_j$ ”; or else,
- “ $|\mathbf{E}_\rho[A_i] - \theta_i| \leq \eta_i + \epsilon$ for all i ”.

Indeed, the algorithm can be implemented by a projector applied to $\rho^{\otimes n_0}$.

Returning to Hypothesis Selection, let us define

$$\Delta_k = \max_{i < j} |\mathbf{E}_{\sigma_k}[A_{ij}] - \mathbf{E}_\rho[A_{ij}]|,$$

and note that $\Delta_{i^*} \leq \eta$, by Helstrom’s theorem. Let us also assume the algorithm has a candidate upper bound $\bar{\eta}$ on η . Now suppose our algorithm is able to find ℓ with $\Delta_\ell \leq \bar{\eta} + \epsilon$. Then the proof of [Proposition 5.3](#) similarly implies that σ_ℓ satisfies $d_{\text{tr}}(\sigma_k, \rho) \leq 2\eta + \bar{\eta} + \epsilon$.

Now let \mathcal{T}_k denote the following Threshold Decision instance (as in [Corollary 5.4](#)): $\epsilon_0 = \epsilon$, $\delta_0 = 1/3$, $m_0 = \binom{m}{2}$, the quantum events are all the A_{ij} ’s, the thresholds are $\theta_{ij} = \mathbf{E}_{\sigma_k}[A_{ij}]$, each “ η_{ij} ” is $\bar{\eta}$. Then [Corollary 5.4](#) gives us a projector B_k on $(\mathbb{C}^d)^{\otimes n_0}$, where $n_0 = O(\log(m)/\epsilon^2)$, with the following property: When it is used to measure $\varrho = \rho^{\otimes n_0}$,

$$\Delta_k \leq \bar{\eta} \implies \mathbf{E}_\varrho[\bar{B}_k] \geq 2/3, \quad \mathbf{E}_\varrho[\bar{B}_k] > 1/3 \implies \Delta_k \leq \bar{\eta} + \epsilon. \quad (5)$$

We can now apply our Threshold Search routine to the \bar{B}_k ’s (with all thresholds $\theta_k = 1/2$), using $n_{\text{TS}}(m, 1/6, \delta')$ copies of ϱ . Provided that indeed $\eta \leq \bar{\eta}$, we know there is at least one k (namely $k = i^*$) with $\Delta_k \leq \bar{\eta}$; thus except with probability at most δ' , the Threshold Search routine will find an ℓ with $\Delta_\ell \leq \bar{\eta} + \epsilon$.

If we wish to assume our Hypothesis Selection algorithm “knows” η , then we are done. Otherwise, we can search for the approximate value of η , as follows: We perform the above routine with $\bar{\eta} = 1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots$, stopping either when Threshold Search fails to find any ℓ or when $\bar{\eta} \leq \epsilon$. If we stop for the first reason, we know that our second-to-last $\bar{\eta}$ is at most 2η ; if we stop for the second reason, we know that $\eta \leq \epsilon$. Either way, assuming no failure on any of the Threshold Searches, we end with a guarantee of $d_{\text{tr}}(\sigma_k, \rho) \leq 4\eta + 2\epsilon$. To bound the overall failure probability we take $\delta' = \delta/\Theta(\log(1/\max\{\eta, \epsilon\}))$. It’s easy to check that the geometric decrease of $\bar{\eta}$ means we only use $O(n_{\text{TS}}(m, 1/6, \delta'))$ copies of ϱ , which is $O(n_{\text{TS}}(m, 1/6, \delta'))n_0$ copies of ρ . Finally, by tuning the constants we can make the final guarantee $d_{\text{tr}}(\sigma_k, \rho) \leq 3.01\eta + \epsilon$. We conclude:

Proposition 5.5. *The above-described method selects σ_ℓ with $d_{\text{tr}}(\rho, \sigma_\ell) \leq 3.01\eta + \epsilon$ (except with probability at most δ), using*

$$n = \frac{\log^3 m + L \cdot \log m}{\epsilon^2} \cdot O(L)$$

copies of ρ , where $L = \log(\log(1/\max\{\eta, \epsilon\})/\delta)$.

It remains to establish the last part of [Theorem 1.5](#), which operates under the assumption that $\eta < \frac{1}{2}(\alpha - \epsilon)$, where $\alpha = \min_{i \neq j} d_{\text{tr}}(\sigma_i, \sigma_j)$. Writing $\bar{\eta} = \frac{1}{2}(\alpha - \epsilon)$ (which is a quantity known to the algorithm), we have $\Delta_{i^*} = \eta \leq \bar{\eta}$, but $\Delta_k > \bar{\eta} + \epsilon$ for all $k \neq i^*$; the reason for this last claim is that

$$\Delta_k \geq |\mathbf{E}_{\sigma_k}[A_{i^*k}] - \mathbf{E}_{\rho}[A_{i^*k}]| \geq |\mathbf{E}_{\sigma_k}[A_{i^*k}] - \mathbf{E}_{\sigma_{i^*}}[A_{i^*k}]| - \eta = d_{\text{tr}}(\sigma_{i^*}, \sigma_k) - \eta \geq \alpha - \eta = 2\bar{\eta} + \epsilon - \eta > \bar{\eta} + \epsilon$$

where the second inequality above used Helstrom's theorem and $\eta = d_{\text{tr}}(\rho, \sigma_{i^*})$ and the last inequality used $\bar{\eta} > \eta$. Now if we perform Threshold Search to achieve [Inequality \(5\)](#) as before, except that we select $\delta_0 = \delta/(4m)$ rather than $1/3$, we'll get projectors B_1, \dots, B_m on $(\mathbb{C}^d)^n$ for $n = O(\log(m/\delta)/\epsilon^2)$ such that, for $\rho = \rho^{\otimes n}$,

$$\mathbf{E}_{\rho}[\bar{B}_{i^*}] \geq 1 - \delta/(4m), \quad \mathbf{E}_{\rho}[\bar{B}_k] \leq \delta/(4m) \quad \forall k \neq i^*.$$

It remains to apply the Quantum Union Bound (specifically, [Corollary 2.7](#)) to B_1, \dots, B_m and ρ to pick out i^* except with probability at most $4 \sum_i \delta/(4m) \leq \delta$. We conclude:

Proposition 5.6. *Using the assumption $\eta < \frac{1}{2}(\alpha - \epsilon)$, where $\alpha = \min_{i \neq j} d_{\text{tr}}(\sigma_i, \sigma_j)$, the above-described method selects σ_{i^*} (except with probability at most δ), using $n = O(\log(m/\delta)/\epsilon^2)$ copies of ρ .*

Acknowledgments

The authors are very grateful to John Wright, whose early contributions to this work provided invaluable understanding of the problem.

References

- [Aar16] Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016. Lectures from the Thérien Barbados Workshop on Computational Complexity. [2](#), [3](#)
- [Aar20] Scott Aaronson. Shadow tomography of quantum states. *SIAM Journal on Computing*, pages STOC18–368–STOC18–394, 6 2020. [2](#), [3](#), [5](#), [22](#), [30](#), [31](#)
- [ACH⁺19] Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. Online learning of quantum states. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124019, dec 2019. [2](#), [5](#), [22](#), [23](#)
- [AISW19] Jayadev Acharya, Ibrahim Issa, Nirmal Shende, and Aaron Wagner. Measuring quantum entropy. In *IEEE International Symposium on Information Theory (ISIT)*, pages 3012–3016. IEEE, 2019. [5](#)
- [AR19] Scott Aaronson and Guy Rothblum. Gentle measurement of quantum states and differential privacy, 2019. [2](#), [4](#), [12](#)

- [BCL20] Sebastien Bubeck, Sitan Chen, and Jerry Li. Entanglement is necessary for optimal quantum property testing, 2020. [5](#)
- [Ber95] Alain Berlinet. A note on variance reduction. *Statist. Probab. Lett.*, 25(4):357–360, 1995. [13](#)
- [BKS^W20] Mark Bun, Gautam Kamath, Thomas Steinke, and Zhiwei Steven Wu. Private hypothesis selection, 2020. [6](#)
- [BNS⁺15] Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis, 2015. [2](#)
- [BOW19] Costin Bădescu, Ryan O’Donnell, and John Wright. Quantum state certification. In *STOC’19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 503–514. ACM, New York, 2019. [5](#), [9](#)
- [BV08] Alain Berlinet and Igor Vajda. Divergence criteria for improved selection rules. *ÚTIA AV ČR (Praha 2008) Internal Publication*, 5, 2008. [5](#)
- [CM06] Matthias Christandl and Graeme Mitchison. The spectra of quantum states and the Kronecker coefficients of the symmetric group. *Comm. Math. Phys.*, 261(3):789–797, 2006. [5](#)
- [DFH⁺16] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Preserving statistical validity in adaptive data analysis, 2016. [2](#)
- [DL96] Luc Devroye and Gábor Lugosi. A universally acceptable smoothing factor for kernel density estimates. *Ann. Statist.*, 24(6):2499–2512, 1996. [5](#)
- [DL97] Luc Devroye and Gábor Lugosi. Nonasymptotic universal smoothing factors, kernel complexity and Yatracos classes. *Ann. Statist.*, 25(6):2626–2637, 1997. [5](#)
- [DL01] Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics. Springer-Verlag, New York, 2001. [5](#), [25](#)
- [Doe18] Benjamin Doerr. An elementary analysis of the probability that a binomial random variable exceeds its expectation. *Statist. Probab. Lett.*, 139:67–74, 2018. [15](#)
- [DR13] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–487, 2013. [2](#), [4](#)
- [FC95] Christopher Fuchs and Carlton Caves. Mathematical techniques for quantum communication theory. *Open Systems & Information Dynamics*, 3(3):345–356, 1995. [10](#)
- [Gao15] Jingliang Gao. Quantum union bounds for sequential projective measurements. *Physical Review A*, 92(5):052331, 2015. [12](#)

- [GKK⁺20] Sivakanth Gopi, Gautam Kamath, Janardhan Kulkarni, Aleksandar Nikolov, Zhiwei Steven Wu, and Huanyu Zhang. Locally private hypothesis selection, 2020. [6](#)
- [GS02] Alison Gibbs and Francis Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002. [4](#)
- [Hel76] Carl Helstrom. *Quantum Detection and Estimation Theory*. Academic Press, 1976. [9](#)
- [HHJ⁺17] Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. *IEEE Trans. Inform. Theory*, 63(9):5628–5641, 2017. [5](#), [7](#)
- [HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, Jun 2020. [2](#)
- [HLM17] Aram Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1598–1611. SIAM, Philadelphia, PA, 2017. [3](#), [30](#), [31](#)
- [HM02] Masahito Hayashi and Keiji Matsumoto. Quantum universal variable-length source coding. *Phys. Rev. A (3)*, 66(2):022311, 13, 2002. [5](#)
- [MdW16] Ashley Montanaro and Ronald de Wolf. A survey of quantum property testing. *Theory of Computing Library*, Graduate Surveys(7):1–81, 2016. [5](#)
- [MŠ07] Satyaki Mahalanabis and Daniel Štefankovič. Density estimation in linear time, 2007. [6](#)
- [NS11] Michael Nussbaum and Arleta Szkoła. An asymptotic error bound for testing multiple quantum hypotheses. *Ann. Statist.*, 39(6):3211–3233, 2011. [6](#)
- [OW15] Ryan O’Donnell and John Wright. Quantum spectrum testing. In *STOC’15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 529–538. ACM, New York, 2015. [5](#)
- [OW16] Ryan O’Donnell and John Wright. Efficient quantum tomography. In *STOC’16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 899–912. ACM, New York, 2016. [5](#), [7](#)
- [OW17] Ryan O’Donnell and John Wright. Efficient quantum tomography II. In *STOC’17—Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 962–974. ACM, New York, 2017. [5](#), [7](#)
- [QCR20] Yihui Quek, Clement Canonne, and Patrick Reberntrost. Robust quantum minimum finding with an application to hypothesis selection, 2020. [6](#)

- [Smi17] Adam Smith. Lecture notes for The Algorithmic Foundations of Adaptive Data Analysis. Lecture 7–10: Stability and adaptive analysis I., 2017. <https://adaptivedataanalysis.files.wordpress.com/2017/10/lect07-10-draft-v1.pdf>. 4
- [Wat18] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. 10
- [Yat85] Yannis Yatracos. Rates of convergence of minimum distance estimators and Kolmogorov’s entropy. *Ann. Statist.*, 13(2):768–774, 1985. 5, 25
- [Yu20a] Nengkun Yu. Quantum closeness testing: A streaming algorithm and applications, 2020. 5
- [Yu20b] Nengkun Yu. Sample efficient tomography via Pauli measurements, 2020. 5
- [Yu20c] Nengkun Yu. Sample optimal quantum identity testing via Pauli measurements, 2020. 5

A The quantum Threshold Decision problem

As mentioned at the end of [Section 1.1.1](#), Aaronson [[Aar20](#)] showed that the *decision* version of quantum Threshold Search can be done with $n = O(\log(m) \log(1/\delta)/\epsilon^2)$ copies, through the use of a theorem of Harrow, Lin, and Montanaro [[HLM17](#), Cor. 11]. In [Theorem A.2](#) below, we give a new version of the Harrow–Lin–Montanaro theorem, with a mild qualitative improvement. This improvement also lets us improve the quantum Threshold Decision copy complexity slightly, to $n = O(\log(m/\delta)/\epsilon^2)$ (see [Corollary A.4](#)).

First, a lemma:

Lemma A.1. *Let $X, Y \in \mathbb{C}^{d \times d}$, with $X \geq 0$. Then*

$$\mathbf{E}_{\rho}[XY] \leq \sqrt{\mathbf{E}_{\rho}[X]} \sqrt{\mathbf{E}_{\rho}[Y^{\dagger}XY]}.$$

Proof. This follows from the matrix form of Cauchy–Schwarz:

$$\begin{aligned} \operatorname{tr}(\rho XY) &= \operatorname{tr}(\sqrt{\rho} \sqrt{X} \cdot \sqrt{X} Y \sqrt{\rho}) \leq \sqrt{\operatorname{tr}(\sqrt{X} \sqrt{\rho} \sqrt{\rho} \sqrt{X})} \sqrt{\operatorname{tr}(\sqrt{\rho} Y^{\dagger} \sqrt{X} \sqrt{X} Y \sqrt{\rho})} \\ &= \sqrt{\operatorname{tr}(\rho X)} \sqrt{\operatorname{tr}(\rho Y^{\dagger} X Y)}. \quad \blacksquare \end{aligned}$$

Theorem A.2. *Let $0 \leq A_1, \dots, A_m \leq \mathbb{1}$ be d -dimensional observables and define $\#A = A_1 + \dots + A_m$. Let $\nu > 0$ and let B be the projector onto the span of eigenvectors of $\#A$ with eigenvalue at least ν . Then for any state $\rho \in \mathbb{C}^{d \times d}$, writing $p_{\max} = \max_i \{\mathbf{E}_{\rho}[A_i]\}$, we have*

$$p_{\max} - 2\sqrt{\nu} \leq \mathbf{E}_{\rho}[B] \leq \mathbf{E}_{\rho}[\#A]/\nu.$$

Remark A.3. One can read out a similar result in the work of Harrow, Lin, and Montanaro [HLM17, Cor. 11], except with a lower bound of $\mathbf{E}_\rho[B] \geq .632(p_{\max} - \nu)^2$. Note that unlike our bound, their lower bound is never close to 1, even when p_{\max} is very close to 1. It is this difference that leads to our slight improvement for the Threshold Decision problem. We speculate that the lower bound in our result can be sharpened further, to $(1 - O(\sqrt{\nu}))p_{\max}$.

Proof. The upper bound in the theorem is just “Markov’s inequality”; it follows immediately from $\#A \geq \nu B$ (and $\rho \geq 0$). As for the lower bound, suppose $p_{\max} = \mathbf{E}_\rho[A_j] = 1 - \delta$. Using the notation $\overline{B} = \mathbb{1} - B$, and defining $\beta = \mathbf{E}_\rho[\overline{B}A_j\overline{B}]$, we have

$$\beta \leq \mathbf{E}_\rho[\overline{B} \cdot \#A \cdot \overline{B}] < \nu,$$

since $A_j \leq \#A$ and $\overline{B} \cdot \#A \cdot \overline{B} < \nu \mathbb{1}$ by definition. On the other hand, write $p = \mathbf{E}_\rho[\overline{B}]$, so our goal is to show $p < \delta + 2\sqrt{\nu}$. Then

$$\begin{aligned} p = \mathbf{E}_\rho[\overline{B}] &= \mathbf{E}_\rho[A_j \cdot \overline{B}] + \mathbf{E}_\rho[\overline{A}_j \cdot \overline{B}] \leq \sqrt{\frac{\mathbf{E}_\rho[A_j]}{\rho}} \sqrt{\mathbf{E}_\rho[\overline{B}A_j\overline{B}]} + \sqrt{\frac{\mathbf{E}_\rho[\overline{A}_j]}{\rho}} \sqrt{\mathbf{E}_\rho[\overline{B}\overline{A}_j\overline{B}]} \\ &= \sqrt{1 - \delta} \sqrt{\beta} + \sqrt{\delta} \sqrt{p - \beta}, \end{aligned}$$

where the inequality is by [Lemma A.1](#), and the subsequent equality uses $p = \mathbf{E}_\rho[\overline{B}(A_j + \overline{A}_j)\overline{B}]$. The above deduction, together with $\beta < \nu$, yields an upper bound on p . Eschewing the tightest possible bound, we deduce from the above that

$$p \leq \sqrt{\beta} + \sqrt{\delta} \sqrt{p} < \sqrt{\nu} + \frac{\delta + p}{2} \implies p \leq 2\sqrt{\nu} + \delta. \quad \blacksquare$$

Given [Theorem A.2](#), it’s easy to obtain the following quantum Threshold Decision algorithm, similar to [[Aar20](#), Lem. 14]:

Corollary A.4. *In the scenario of quantum Threshold Search, suppose one only wishes to solve the decision problem, meaning the algorithm has only two possible outputs:*

- “there exists j with $\mathbf{E}_\rho[A_j] > \theta_j - \epsilon$ ”; or else,
- “ $\mathbf{E}_\rho[A_i] \leq \theta_i$ for all i ”.

This can be solved using just $n = O(\log(m/\delta)/\epsilon^2)$ copies of ρ , an indeed, the algorithm can be implemented by a projector applied to $\rho^{\otimes n}$.

Furthermore, the very similar-looking [Corollary 5.4](#) holds.

Proof. Writing $\varrho = \rho^{\otimes n}$, a standard Chernoff bound implies there are quantum events A'_1, \dots, A'_m satisfying

$$\mathbf{E}_\rho[A_i] > \theta \implies \mathbf{E}_\varrho[A'_i] \geq 1 - \delta/2, \quad \mathbf{E}_\rho[A_i] \leq \theta - \epsilon \implies \mathbf{E}_\varrho[A'_i] \leq \delta^3/(16m).$$

We apply [Theorem A.2](#) to A'_1, \dots, A'_m and ϱ , with $\nu = \delta^2/16$, obtaining the projector B with

$$\max_i \mathbf{E}_\varrho[A'_i] - \delta/2 \leq \mathbf{E}_\varrho[B] \leq (16/\delta^2) \mathbf{E}_\varrho[\#A'].$$

Now on one hand, if there exists j with $\mathbf{E}_\rho[A_j] > \theta$, we conclude $\mathbf{E}_\varrho[B] \geq 1 - \delta$. On the other hand, if $\mathbf{E}_\rho[A_i] \leq \theta - \epsilon$ for all i , then $\mathbf{E}_\varrho[\#A'] \leq m \cdot \delta^3/(16m)$ and hence $\mathbf{E}_\varrho[B] \leq \delta$. Thus the algorithm can simply measure B with respect to ϱ , reporting “there exists j with $\mathbf{E}_\rho[A_j] > \theta - \epsilon$ ” when B occurs, and “ $\mathbf{E}_\rho[A_j] \leq \theta$ for all i ” when \overline{B} occurs. This completes the main proof.

The “Furthermore” proof of [Corollary 5.4](#) is exactly the same, except we let A'_i be the quantum event that has

$$|\mathbf{E}_\rho[A_i] - \theta_i| > \eta_i + \epsilon \implies \mathbf{E}_\varrho[A'_i] \geq 1 - \delta/2, \quad |\mathbf{E}_\rho[A_i] - \theta_i| \leq \eta_i \implies \mathbf{E}_\varrho[A'_i] \leq \delta^3/(16m). \blacksquare$$