

Quantum Non-Identical Mean Estimation: Efficient Algorithms and Fundamental Limits

Jiachen Hu 


Peking University, Beijing, China

Tongyang Li 

Peking University, Beijing, China

Xinzhao Wang 


Peking University, Beijing, China

Yecheng Xue 

Peking University, Beijing, China

Chenyi Zhang 

Stanford University, CA, USA

Han Zhong 

Peking University, Beijing, China

Abstract

We systematically investigate quantum algorithms and lower bounds for mean estimation given query access to non-identically distributed samples. On the one hand, we give quantum mean estimators with quadratic quantum speed-up given samples from different bounded or sub-Gaussian random variables. On the other hand, we prove that, in general, it is impossible for any quantum algorithm to achieve quadratic speed-up over the number of classical samples needed to estimate the mean μ , where the samples come from different random variables with mean close to μ . Technically, our quantum algorithms reduce bounded and sub-Gaussian random variables to the Bernoulli case, and use an uncomputation trick to overcome the challenge that direct amplitude estimation does not work with non-identical query access. Our quantum query lower bounds are established by simulating non-identical oracles by parallel oracles, and also by an adversarial method with non-identical oracles. Both results pave the way for proving quantum query lower bounds with non-identical oracles in general, which may be of independent interest.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum query complexity

Keywords and phrases Quantum algorithms, Mean estimation, Non-identical samples, Query complexity

Digital Object Identifier 10.4230/LIPIcs.TQC.2024.9

Related Version *Full Version:* <https://arxiv.org/abs/2405.12838> [14]

Funding TL, XW, and YX were supported by the National Natural Science Foundation of China (Grant Numbers 62372006 and 92365117), and The Fundamental Research Funds for the Central Universities, Peking University.

1 Introduction

The problem of estimating the mean μ of a random variable X given its i.i.d. samples is a fundamental problem in statistics. For any random variable X with finite variance σ^2 , the median-of-means estimator can estimate μ to within additive error ϵ with failure probability $\leq \delta$ using $O(\frac{\sigma^2}{\epsilon^2} \log(\frac{1}{\delta}))$ samples. This sample complexity is known to be tight up to a constant multiplicative factor [7].



© Jiachen Hu, Tongyang Li, Xinzhao Wang, Yecheng Xue, Chenyi Zhang, and Han Zhong;
licensed under Creative Commons License CC-BY 4.0

19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024).

Editors: Frédéric Magniez and Alex Bredariol Grilo; Article No. 9; pp. 9:1–9:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

On the other hand, suppose that a quantum computer has access to a unitary U and its inverse such that $U|0\rangle$ encodes the random variable X coherently, and each application of U and U^\dagger as a black-box oracle can be regarded as a quantum analogue of getting a sample of the random variable X . Therefore, the application of U is sometimes called a *quantum experiment* [11]. Under this assumption, a quantum computer can estimate the mean of X with $O(\frac{\sigma}{\epsilon} \log(\frac{1}{\delta}))$ quantum experiments [17], which achieves quadratic speed-up compared to the classical counterpart. Such quantum mean estimators embrace various applications, including approximate counting [17, 6], data stream estimation [12], derivative pricing in finance [5], etc.

In some cases, we are interested in estimating the mean of “close” random samples, such as random samples with the same mean but different distributions. For example, it is ubiquitous that the measurements of random samples have small systematic errors. In such cases there may be small difference between the means of the actual distributions of the measured random samples, and our algorithms and lower bounds also take this into account. One specific example is to learn a linear system discussed below. In classical mean estimation, the same method for identical random variables also works for non-identical random variables. As long as the variance of all random variables is bounded by σ^2 , the median-of-means estimator can be directly adapted to these situations, yielding an algorithm with the same complexity. However, it is unclear whether similar results hold in the regime of quantum mean estimation. Therefore, it is a natural question whether we can achieve quantum speed-up for the mean estimation problem with non-identically distributed samples.

Below we provide a potential application for the quantum mean estimation with non-identically distributed samples.

Quantum Linear System

A classical linear dynamical system (LDS) is defined as

$$x_{t+1} = Ax_t + w_t, \quad x_t \in \mathbb{R}^n, \quad w_t \sim \mathcal{N}(\mathbf{0}, \sigma_w^2), \quad \|A\|_2 < 1, \quad x_1 = \mathbf{0} \tag{1}$$

where x_t is the state at time step t , and w_t is a random noise at step t . A well-known problem in LDS is to do the system identification: estimating the transition matrix A given a series of states starting from step 1. The standard approach to estimate transition matrix A in the classical linear system is ordinary least squares (OLS) [8, 20].

Consider the quantum counterpart of LDS (for example, when simulating a LDS on a quantum computer):

$$U_f|\psi_x\rangle|0\rangle = \int_{\mathbb{R}^n} \sqrt{f_w(w)}|\psi_x\rangle|\psi_{Ax+w}\rangle dw, \tag{2}$$

$$U_o|\psi_x\rangle|0\rangle = |\psi_x\rangle|x\rangle, \tag{3}$$

here $f_w(w)$ is the probability density function (pdf) of $\mathcal{N}(\mathbf{0}, \sigma_w^2)$, and $|\psi_x\rangle$ is an arbitrary embedding of the raw state x . It is natural to ask whether it is possible to estimate A by a quantum algorithm with desired speed-up in quantum linear systems. Actually, it is indeed possible with a procedure presented in Section 4.1.3. This estimation procedure uses multiple calls to U_f to construct a new oracle U_{t_0} for some step t_0 , which encodes a probability distribution over the matrix space with A as the mean value. However, the distribution encoded by U_{t_0} is different for different t_0 , though their means are all equal to A . Therefore, this problem presents another motivation of the quantum non-identical mean estimation problem.

In general, the quantum linear system problem described above is a special class of quantum estimation problem in which quantum probability oracles have a time-varying zero-mean noise. The distribution of noise at each step is different but all zero-mean. The number of samples at each step is limited.

1.1 Contributions

In this paper, we systematically analyze the sample complexity of the *quantum non-identical mean estimation problem* (see its formal definition in Task 8). Roughly speaking, the quantum algorithm is given T different random variables in turn and can get $m \in \mathbb{N}$ samples from each random variable. Suppose that the mean of every random variable is in $(\mu - c\epsilon, \mu + c\epsilon)$ for some constant $0 < c < 1$, the quantum non-identical mean estimation problem is to estimate μ up to additive error ϵ . If all random variables are bounded or sub-Gaussian (see definition in Definition 14), for accuracy ϵ and $m = \Omega(\log(\frac{1}{\epsilon}))$, we give quantum algorithms solving the quantum non-identical mean estimation problem with quadratic speed-up.

- **Theorem 1** (Informal versions of Theorem 12 and Theorem 16). *For the quantum non-identical mean estimation problem with sufficiently small accuracy ϵ ,*
- *if all random variables are bounded in $[L, H]$ and $m = \Omega(\log(\frac{H-L}{\epsilon}))$, there is a quantum algorithm that estimates μ to within additive error ϵ if $T = \Omega(\frac{H-L}{\epsilon})$. The algorithm uses $O(\frac{H-L}{\epsilon} \log(\frac{H-L}{\epsilon}))$ samples in total;*
 - *if all random variables are sub-Gaussian with parameter K and $m = \tilde{\Omega}(\log(\frac{K}{\epsilon}))$, there is a quantum algorithm that estimates μ to within additive error ϵ if $T = \tilde{\Omega}(\frac{K}{\epsilon})$. The algorithm uses $\tilde{O}(\frac{K}{\epsilon})$ samples in total.*

In the worst case, the variance of random variables bounded in $[L, H]$ can be $(H - L)^2/4$, so the optimal classical estimator needs $\Theta((H - L)^2/\epsilon^2)$ samples to estimate μ up to additive error ϵ . For normal random variables, their sub-Gaussian parameter K equals their standard deviation σ , so the optimal classical estimator needs $\Theta(K^2/\epsilon^2)$ samples to estimate μ up to additive error ϵ . Therefore, the quantum estimators in Theorem 1 achieve nearly quadratic speed-up compared to classical estimators.

On the other hand, for $m = 1$, we show that any algorithm with relatively small working register have no speed-up compared to classical estimators.

- **Theorem 2** (Informal version of Theorem 23). *Suppose all random variables in the quantum non-identical mean estimation problem with $m = 1$ have mean bounded by R and variance bounded by σ^2 . Let \mathcal{A} be a quantum query algorithm acting on query register Q , working register W such that the number of qubits in Q is larger than that in W by $\Omega(\log(\frac{R}{\epsilon}))$. It requires $T = \Omega(\frac{\sigma^2}{\epsilon^2})$ if there exists an algorithm \mathcal{A} solving this problem. The sample complexity of \mathcal{A} is $T = \Omega(\frac{\sigma^2}{\epsilon^2})$.*

For general $m \geq 1$, we give another sample complexity lower bound of estimating mean of Bernoulli random variables.

- **Theorem 3** (Informal version of Theorem 25). *Suppose all random variables in the quantum non-identical mean estimation problem with $m \geq 1$ are Bernoulli random variables with mean $\mu \in (0, 1)$, and the accuracy ϵ satisfies $\epsilon \leq \mu(1 - \mu)$ and $\epsilon = O(\frac{1}{m^2})$. It requires $T = \Omega(\frac{1}{\epsilon m^2})$ if there exists a quantum query algorithm solving this problem. The sample complexity is $mT = \Omega(\frac{1}{\epsilon m})$ in total.*

In Theorem 3, we take the Bernoulli random variables as a hard instance for the quantum non-identical mean estimation problem. Note that if $\epsilon = \Theta(\mu(1 - \mu))$, the classical optimal estimator needs $\Theta(\frac{\mu(1-\mu)}{\epsilon^2}) = \Theta(\frac{1}{\epsilon})$ samples to estimate the mean of the Bernoulli random

variable. Therefore, Theorem 3 shows that there is no quantum speed-up in this case if $m = O(1)$. However, it does not rule out the possibility of quantum speed-up for estimating the mean of Bernoulli random variables with $\epsilon = o(\mu(1 - \mu))$ or $m = \Omega(1)$. For example, if $\mu = \Theta(1)$, $\epsilon = o(1)$, and $m = \Omega(\log(\frac{1}{\epsilon}))$, the quantum estimator for bounded random variables in Theorem 1 can estimate μ up to error ϵ using $O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon}))$ samples while classical estimators need $\Omega(\frac{1}{\epsilon^2})$ samples.

In addition, Theorem 2 and Theorem 3 give two different lower bounds when $m = 1$. Compared with Theorem 3, the lower bound in Theorem 2 matches the classical upper bound for general distributions with variance σ^2 , but an additional requirement is that the register W has relatively small dimension.

Finally, we use Bernoulli random variable as an example to summary our systematical investigation on the quantum non-identical mean estimation problem.

- **Corollary 4.** For Bernoulli random variable with mean μ such that $\epsilon = \Theta(\mu(1 - \mu))$,
- if $m = \Omega(\log(1/\epsilon))$ and $T = \Omega(1/\epsilon)$, there exists an algorithm solving this problem using $O(\frac{1}{\epsilon} \log(1/\epsilon))$ quantum samples, achieving a near-quadratic speed-up;
- if $m = \Omega(\log(1/\epsilon))$ and $T = o(1/\epsilon m^2)$, there is no quantum algorithm solving this problem. There is an additional requirement that $\epsilon = O(1/m^2)$;
- if $m = O(1)$, there is no quantum speed-up for this problem.

Proof. This corollary comes directly from Theorem 1, Theorem 2, and Theorem 3. ◀

1.2 Techniques

1.2.1 Upper Bound

From a high-level perspective, our quantum algorithms for non-identical mean estimation encode the mean to an amplitude, use an uncomputation trick to be introduced below to align different oracles, and then use amplitude estimation to estimate the mean.

We start with the bounded case. Recall that this paper studies non-identically distributed samples and assumes that we have access to unitaries O_{X_1}, \dots, O_{X_T} , where

$$O_{X_i}|\mathbf{0}\rangle = \sum_{x \in E_i} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle. \tag{4}$$

The mean $\mu = \mu_i = \sum_{x \in E_i} p_i(x)x$ is equal for different $i \in [T]$ (In fact, these μ_i can be slightly different – see Remark 13 for more details), but each O_{X_i} has potentially different garbage states $|\psi_x^{(i)}\rangle$ and each can only be used for very limited times. Suppose that for any $i \in [T]$, the bounded random variable X_i satisfies $X_i \in [L, H]$. If we have sufficient access to any specific O_{X_i} , we can construct a unitary

$$U_i|\mathbf{0}\rangle|0\rangle = \sqrt{q}|\psi_1^{(i)}\rangle|1\rangle + \sqrt{1-q}|\psi_0^{(i)}\rangle|0\rangle \tag{5}$$

by one call to O_{X_i} and a series of controlled rotations [17], where $q = (\mu - L)/(H - L)$. Consequently, the mean is encoded to an amplitude and direct amplitude estimation provides mean estimation with quadratic quantum speedup. However, in the non-identical case, we do not have sufficient number of calls to any specific U_i to provide quadratic speedup. Furthermore, it is very difficult to use a mixture of different U_i in amplitude estimation [3]. This is due to the reason that amplitude estimation is based on Grover’s algorithm [9], which is essentially rotation in a two-dimensional plane spanned by two specific quantum states

related to U_i . In our case, different U_i may have different $|\phi_1^{(i)}\rangle$ and $|\phi_0^{(i)}\rangle$, which forms different rotation planes and thus their mixed use is invalid. However, we can use a small number of calls to U_i to construct a unitary such that

$$S_i|\mathbf{0}\rangle = \sqrt{1-\epsilon_i}|0\rangle\left(\sqrt{r}|1\rangle + \sqrt{1-r}|0\rangle\right) + \sqrt{\epsilon_i}|1\rangle|\text{garbage}_i\rangle \quad (6)$$

with r being a bijective function of q (the concrete value to be shown later) and ϵ_i being sufficiently small. Since the garbage state is small enough to be handled as an approximation error, S_i can be seen as an approximation of an unitary $S: |0\rangle \rightarrow \sqrt{r}|1\rangle + \sqrt{1-r}|0\rangle$. Therefore, We can then use these S_i instead of S to perform amplitude estimation, which provides estimation for r and thus q and μ .

The construction of S_i can be accomplished by an uncomputation trick [6] and fixed-point search [22]. Specifically, the uncomputation trick is to perform a unitary

$$V_i = (U_i^\dagger \otimes I)(I \otimes \text{CNOT})(U_i \otimes I) \quad (7)$$

instead of U_i , which enjoys a property that it extracts the value of q separated from a garbage state related to $|\phi_1^{(i)}\rangle$ and $|\phi_0^{(i)}\rangle$. The computing result of $\langle b|\langle 0|\langle \mathbf{0}|V_i|\mathbf{0}\rangle|0\rangle|0\rangle$ for $b \in \{0, 1\}$ tells that $V_i|\mathbf{0}\rangle|0\rangle|0\rangle$ only has components $|\mathbf{0}\rangle|0\rangle|0\rangle$, $|\mathbf{0}\rangle|0\rangle|1\rangle$, and a garbage state orthogonal to them. Besides, the amplitudes of the first two components are determined by q . In particular, it satisfies

$$V_i|\mathbf{0}\rangle|0\rangle|0\rangle = \sqrt{2q^2 - 2q + 1}|\mathbf{0}\rangle|0\rangle\left(\frac{q}{\sqrt{2q^2 - 2q + 1}}|1\rangle + \frac{1-q}{\sqrt{2q^2 - 2q + 1}}|0\rangle\right) + \sqrt{2q - 2q^2}|\text{garbage}_i\rangle, \quad (8)$$

where $|\text{garbage}_i\rangle$ is a unit garbage state and $(I \otimes \langle 0|\langle \mathbf{0}|)|\text{garbage}_i\rangle = 0$. Therefore, we can use fixed-point quantum search [22] to stably amplify the amplitude of the state $\frac{q}{\sqrt{2q^2 - 2q + 1}}|1\rangle + \frac{1-q}{\sqrt{2q^2 - 2q + 1}}|0\rangle$ and thus S_i is constructed with $r = \frac{q^2}{2q^2 - 2q + 1}$. See Theorem 12 for more details.

For a sub-Gaussian random variable with the absolute value of mean bounded by the sub-Gaussian parameter K , the probability of the random variable being more than a threshold related to K is sufficiently small and the mean of a truncated random variable can be a good enough approximation. Therefore, this case can be reduced to the case of bounded random variables. For general sub-Gaussian random variables X_1, \dots, X_T , a constant number of classical experiments provide an estimation $\hat{\mu}$ within K -additive error, thus $X_1 - \hat{\mu}, \dots, X_T - \hat{\mu}$ are sub-Gaussian random variables with the absolute value of mean bounded by K , which has been solved (see Theorem 16 for more details).

1.2.2 Lower Bound

We prove our two quantum query lower bounds using different techniques: the case $m = 1$ (Theorem 2) is established by simulating non-identical oracles by parallel oracles, and the case $m \geq 1$ (Theorem 3) is established by an adversarial method with non-identical oracles.

Simulating T Non-Identical Oracles by Constant T -Parallel Oracles

For the quantum non-identical mean estimation problem with $m = 1$, we give a sample complexity lower bound in Theorem 23 by constructing a quantum circuit with constant query depth simulating the original quantum circuit querying non-identical oracles. For any quantum query algorithm \mathcal{A} using the *state preparation oracle* U_x such that the state

$U_x|\mathbf{0}\rangle$ encodes the input, suppose that there is a sequence of unitary oracles that maps $|\mathbf{0}\rangle$ to the same state but have different effects acting on other states orthogonal to $|\mathbf{0}\rangle$. Suppose that the working register of \mathcal{A} is relatively small and \mathcal{A} queries T non-identical oracles. In Theorem 21, we prove that for any projection Π with small image space, there is a quantum algorithm \mathcal{A}' using two T -parallel queries such that

$$\|\Pi\mathcal{A}|\mathbf{0}\rangle\|^2 = \|(\Pi \otimes \langle \mathbf{0} |) \mathcal{A}'|\mathbf{0}\rangle|\mathbf{0}\rangle\|^2, \tag{9}$$

where a T -parallel query is to query T oracles simultaneously. This theorem builds a bridge between quantum algorithms with non-identical state preparation oracles and quantum algorithms with low query depth. If for any input x correct outputs of \mathcal{A} lie in a small space V_x , and let $\text{Im}(\Pi) = V_x$, then Theorem 21 shows that \mathcal{A} and \mathcal{A}' have the same probability to output a correct answer.

In Theorem 23, we prove that any quantum query algorithm \mathcal{A} starting from an efficiently preparable state $|\mathbf{0}\rangle$ can be modified to recover the query register to $|\mathbf{0}\rangle$ with a small overhead. This reduces the dimension of the subspace that the correct outputs of \mathcal{A} lie in, and then we use Theorem 21 to give a sample complexity lower bound of the quantum non-identical mean estimation problem with $m = 1$ based on the facts that parallelization only brings classical advantage to solving the quantum approximate counting problem [4], and the quantum approximate counting problem can be reduced to estimating the mean of Bernoulli random variables.

Adversarial Method with Non-Identical Oracles

Given a boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and access to a unitary oracle O_x which encodes the information of some $x \in \{0, 1\}^n$, the *generalized adversarial method* [13] gives a tight query complexity lower bound of computing $f(x)$. For any quantum query algorithm \mathcal{A} and $x \in \{0, 1\}^n$, let $|\psi_x^{(t)}\rangle$ be the quantum state after \mathcal{A} queries O_x for t times. Suppose \mathcal{A} can compute $f(x)$ with high probability for all $x \in \{0, 1\}^n$ using T queries, then we have $\langle \psi_x^{(T)} | \psi_y^{(T)} \rangle = 1 - \Omega(1)$ for all $x \in f^{-1}(0)$ and $y \in f^{-1}(1)$. Since $\langle \psi_x^{(0)} | \psi_y^{(0)} \rangle = 1$, to give a lower bound of T , it suffices to give an upper bound on the *progress* at time t , $\langle \psi_x^{(t-1)} | \psi_y^{(t-1)} \rangle - \langle \psi_x^{(t)} | \psi_y^{(t)} \rangle$, for all $x \in f^{-1}(0)$, $y \in f^{-1}(1)$, and $t \in [T]$. The generalized adversarial method assigns a weight Γ_{xy} to every pair of $x \in f^{-1}(0)$, $y \in f^{-1}(1)$, which proves an upper bound for the weighted progress at time t :

$$S_{t-1} - S_t = \sum_{x \in f^{-1}(0), y \in f^{-1}(1)} \Gamma_{xy} (\langle \psi_x^{(t-1)} | \psi_y^{(t-1)} \rangle - \langle \psi_x^{(t)} | \psi_y^{(t)} \rangle), \tag{10}$$

and hence gives a lower bound on T . However, they regard $|\psi_x^{(t-1)}\rangle, |\psi_y^{(t-1)}\rangle$ as free variables independent of previous states $|\psi_x^{(t')}\rangle, |\psi_y^{(t')}\rangle$ for $t' < t - 1$ while bounding the weighted progress at t , so their upper bound of $S_{t-1} - S_t$ is independent of t . Therefore, if the algorithm queries different oracles at different times, the adversarial method cannot give better lower bound than the case that all oracles are the same. In Lemma 24, we apply the adversarial method on the quantum approximate counting problem, but analyze the progress in another way which utilizes the connection between $|\psi_x^{(t)}\rangle$ and $|\psi_x^{(t')}\rangle$ for different t and t' . Specifically, we show that any quantum query algorithm solving the quantum approximate counting problem has progress upper bounded by $O(\frac{t}{n})$ at time t , where n is the number of items. The original adversarial method gives an $O(\frac{1}{\sqrt{n}})$ upper bound of the progress at any time t . Boyer et al. [2] gave a similar analysis of quantum search which utilizes the connection between states at different time t , and got a tight lower bound of quantum search

with a better constant factor compared to the hybrid argument. Since Reichardt [19] proved that the generalized adversarial method is asymptotically tight, we cannot expect more by exploring connections between states at different time with identical query oracles. However, if each oracle can only be queried a limited number of times, our bound in Lemma 24 is better than that obtained by the generalized adversarial method, since the progress bound $O(\frac{t}{n})$ is smaller in the early stages of the algorithm. We use this result to prove a query complexity lower bound of the quantum approximate counting problem with non-identical oracles. Since the quantum approximate counting problem can be reduced to estimating the mean of a Bernoulli random variable, we get a sample complexity lower bound of the quantum non-identical mean estimation problem in Theorem 3 for general m .

1.3 Organization

The rest of the paper is organized as follows. In Section 2 we formally define the input model and the quantum non-identical mean estimation problem, introduce the concept of parallel quantum query algorithms, and introduce quantum subroutines used in our algorithms. In Section 3 we give quantum algorithms for estimating the mean of non-identically distributed bounded or sub-Gaussian random variables with quadratic speed-up. In Section 4 we give two quantum query lower bounds of the quantum non-identical mean estimation problem based on reductions to low-depth quantum algorithms and the adversarial method with non-identical oracles, respectively.

2 Preliminaries

2.1 Notations

We denote $\{1, 2, \dots, n\}$ by $[n]$. We use $|\psi\rangle_{A,B}$ to indicate that the state $|\psi\rangle$ is in quantum registers A and B . For a quantum register A , we denote its number of qubits by n_A . For a boolean string $x \in \{0, 1\}^n$, we denote its Hamming weight $|\{i \in [n] \mid x_i = 1\}|$ by $|x|$. We abbreviate $|0^k\rangle$ as $|\mathbf{0}\rangle$ if k can be inferred from the context.

2.2 Input Model

We first recall the definition of random variables and the input model of the classical mean estimation problem.

► **Definition 5** (Random variable). *A finite random variable X is a function $X: \Omega \rightarrow E$ for some probability space (Ω, p) , where Ω is the finite sample space, p is a probability measure on Ω , and $E \subset \mathbb{R}$.*

Next, we assume that the random variable is the output of a quantum process O_X , and we can query O_X as an oracle to access X .

► **Definition 6** (Quantum random variable). *For any finite random variable X , a quantum random variable encoding X is a pair (\mathcal{H}, O_X) , where \mathcal{H} is a Hilbert space and O_X is a unitary operator on \mathcal{H} that performs the mapping*

$$O_X|\mathbf{0}\rangle = \sum_{x \in E} \sqrt{p(x)}|\psi_x\rangle|x\rangle \quad (11)$$

for some unknown garbage unit state $|\psi_x\rangle$.

Following the notation in [11], we call each application to U and U^\dagger a *quantum experiment*. We use the number of quantum experiments to measure the sample complexity of a quantum query algorithm.

► **Definition 7** (Quantum experiment). *Let (\mathcal{H}, O_X) be a quantum random variable. A quantum experiment is the process of applying O_X or its inverse O_X^\dagger or their controlled versions to a state in \mathcal{H} .*

Performing a quantum experiment of a quantum random variable (\mathcal{H}, O_X) can be regarded as a query to the unitary oracle O_X in the quantum query model, so the sample complexity is equivalent to the query complexity in this context, and we use the two terms interchangeably.

This input model is widely used in previous quantum mean estimation algorithms. The same oracle as defined in Definition 6 is used in [17]. Kothari and O’Donnell [16] used a similar input model except that they encode the probability distribution and the random variable mapping $\Omega \rightarrow \mathbb{R}$ in two oracles separately, and their algorithm also works well with the oracle in Definition 6. Hamoudi and Magniez [12, 11] used a more general input model called “q-random-variable”, where the value of the random variable is implicitly encoded in a register and can be compared with a constant or performed conditional Pauli rotations, and our oracle can be regarded as an instance of the “q-random-variable”. Since the oracle in Definition 6 already covers many common cases, we use it instead of the “q-random-variable” for simplicity and clarity. In fact, our quantum algorithm in Theorem 12 can also apply to the general “q-random-variable”.

The unitary O_X is a quantum generalization of the process generating a sample of X . Bennett [1] proved that any classical algorithm using time T and space S can be modified to be a reversible algorithm using time $O(T)$ and space $O(ST^\epsilon)$ for any $\epsilon > 0$, and hence can be simulated by a quantum circuit. Therefore, for any randomized algorithm \mathcal{A} , we can implement the oracle O_X in Definition 6 encoding the output distribution of \mathcal{A} with a small overhead.

Another natural way for a quantum algorithm to access a random variable is to assume that several copies of $|\psi_X\rangle = \sum_{x \in E} \sqrt{p(x)}|x\rangle$ encoding the information of X are given as the initial quantum state. This model is weaker than the one in Definition 6 since it does not provide access to a unitary preparing $|\psi_X\rangle$. Hamoudi [11] demonstrated that there is no quantum speed-up for the original mean estimation problem in this model. Therefore, it can be inferred that there is no quantum speed-up for the mean estimation problem of non-identically distributed random variables in this model, as it is a harder problem.

Based on the definition of quantum random variable, we define the mean estimation problem of non-identically distributed random variables formally as the following task.

► **Task 8** (Quantum non-identical mean estimation). *Let $(\mathcal{H}, O_{X_1}), \dots, (\mathcal{H}, O_{X_T})$ be a sequence of quantum random variables on the same Hilbert space \mathcal{H} . Assume there exists μ and $\delta \in (0, 1)$ such that each $\mu_i := \mathbb{E}[X_i]$ satisfies $|\mu_i - \mu| \leq \delta$ for all $i \in [T]$. Given the repetition parameter $m \in \mathbb{N}$ and accuracy ϵ such that $\delta < c\epsilon$ for some constant $c < 1$, the quantum non-identical mean estimation problem is to estimate μ to within additive error ϵ with probability at least $2/3$ using each O_{X_i} or $O_{X_i}^\dagger$ or their controlled versions at most m times.*

The non-identity of quantum random variables means more than the non-identity of classical random variables. Specifically, the difference between two quantum random variables $(\mathcal{H}, O_X), (\mathcal{H}, O_Y)$ lies in the following three aspects: the results of applying O_X and O_Y to states orthogonal to $|\mathbf{0}\rangle$, the garbage state $|\psi_x\rangle$, and the random variables they encode. In contrast, the difference between two classical random variables is solely determined by the third aspect. Consequently, the quantum mean estimation problem of non-identically distributed random variables is more challenging than its classical counterpart.

2.3 Parallel Quantum Query Algorithms

The classical parallel algorithm implies that the algorithm can perform multiple operations simultaneously, which has become increasingly important in recent years with the development of multi-core processors. In the quantum setting, there is an additional reason to consider parallel algorithms: quantum states are fragile and susceptible to disruption by environmental factors, specifically decoherence. By reducing the computation time, parallel quantum algorithms can reduce the probability of decoherence. One example is parallel quantum query algorithms which can make multiple queries simultaneously, where a p -parallel query is defined as making p parallel queries simultaneously. Zalka [23] gave an algorithm that makes $\sqrt{\frac{n}{p}}$ p -parallel queries to solve the unstructured search problem with 1 marked item among n items and showed that its query complexity is optimal. Subsequent works also analyzed the parallel quantum query complexity of quantum search [10], quantum walk [15], quantum counting [4], and Hamiltonian simulation [24].

2.4 Quantum Subroutines

► **Lemma 9** (Approximating unitary operators, Eq. (4.63) of [18]). *Let $\|\cdot\|$ be the operator 2-norm. For unitary operators $\{U_i\}_{i=1}^m, \{V_i\}_{i=1}^m$, it holds that*

$$\|U_m U_{m-1} \dots U_1 - V_m V_{m-1} \dots V_1\| \leq \sum_{j=1}^m \|U_j - V_j\|.$$

► **Lemma 10** (Amplitude estimation, Theorem 12 of [3]). *Given a unitary U satisfying*

$$U|0\rangle = \sqrt{p}|\phi_1\rangle|1\rangle + \sqrt{1-p}|\phi_0\rangle|0\rangle \quad (12)$$

for some $p \in [0, 1]$, there exists a quantum circuit C on a larger space such that the measurement outcome of $C|0\rangle|0\rangle, \tilde{p}$, satisfies

$$|\tilde{p} - p| \leq \frac{2\pi\sqrt{p(1-p)}}{M} + \frac{\pi^2}{M^2} \quad (13)$$

with probability $\frac{8}{\pi^2}$, where C has M calls to the controlled versions of $I - 2U|0\rangle\langle 0|U^\dagger$. Denote the algorithm by $\text{AmpEst}(U, M)$.

► **Lemma 11** (Fixed-point quantum search, [22]). *Let A be a unitary and Π be an orthogonal projector such that $\Pi A|0\rangle = \lambda|\phi\rangle$, where $\lambda \in \mathbb{R}$ and $|\phi\rangle$ is a normalized quantum state. There exists a quantum circuit $S_L = \text{FixSearch}(A, \Pi, \epsilon)$ such that $\| |\phi\rangle - S_L|0\rangle \| \leq \epsilon$, consisting of $O(\log(1/\epsilon)/\lambda)$ queries to A, A^\dagger , and $C_\Pi \text{NOT}$. Here $C_\Pi \text{NOT}$ is the Π -controlled NOT operator*

$$C_\Pi \text{NOT} = X \otimes \Pi + I \otimes (I - \Pi),$$

where X is the Pauli- X matrix.

3 Upper Bound

In this section, we first introduce an algorithm that solves Task 8 for bounded random variables, and then generalize it to sub-Gaussian variables.

9:10 Quantum Non-Identical Mean Estimation

■ **Algorithm 1** Mean Estimation of Bounded Random Variables.

- 1: **Input:** sequence of random variable oracle $\{O_{X_i}\}_{i=1}^T$, accuracy ϵ , mean difference δ , repetition parameter m , lower bound L , upper bound H
- 2: **Output:** mean estimation $\tilde{\mu}$
// Construct quantum circuit S_i
- 3: Construct unitary U_i

$$U_i : |\mathbf{0}\rangle|0\rangle \xrightarrow{O_{X_i} \otimes I} \sum_{x \in E_i} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle |0\rangle$$

$$\xrightarrow{\text{controlled rotation}} \sum_{x \in E_i} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle \left(\sqrt{\frac{x-L}{H-L}} |1\rangle + \sqrt{\frac{H-x}{H-L}} |0\rangle \right)$$

- 4: Let $V_i = (U_i^\dagger \otimes I)(I \otimes \text{CNOT})(U_i \otimes I)$
- 5: Let $S_i = \text{FixSearch}(V_i, |\mathbf{0}\rangle|0\rangle|0\rangle|\mathbf{0}\rangle \otimes I, \epsilon' = O(\epsilon^2/(H-L)^2))$
// Mean estimation using S_i
- 6: Let \tilde{p} be the output of $\text{AmpEst}(S, M = O(\frac{H-L}{\epsilon}))$, where S is arbitrarily replaced by S_1, \dots, S_T .
- 7: Output $\tilde{\mu} = \frac{\tilde{p} - \sqrt{\tilde{p}(1-\tilde{p})}}{2\tilde{p}-1}(H-L) + L$

3.1 Mean Estimation of Bounded Random Variables

In this subsection, we introduce an algorithm that solves Task 8 with quadratic speed-up given the condition that random variables X_1, \dots, X_T are bounded in $[L, H]$. According to the task, for each $i \in [T]$, oracle O_{X_i} can be used at most m times.

For clarity, we describe the algorithm with two phases. Let

$$|\phi_i\rangle = \frac{q_i}{\sqrt{2q_i^2 - 2q_i + 1}} |1\rangle + \frac{1 - q_i}{\sqrt{2q_i^2 - 2q_i + 1}} |0\rangle.$$

Here $q_i = \frac{\mu_i - L}{H - L} \in [0, 1]$. For each $i \in [T]$, We will construct a quantum circuit S_i that satisfies $S_i|\mathbf{0}\rangle \approx |\phi_i\rangle$ with m calls to O_{X_i} . Then we will prove that performing amplitude estimation with these S_i gives an ϵ -additive estimation of μ .

► **Theorem 12.** *Assume that all random variables X_1, \dots, X_T in Task 8 are bounded in $[L, H]$. Let m, ϵ, δ in Algorithm 1 satisfy $m = \Omega(\log(\frac{H-L}{\epsilon}))$, $\epsilon = O(\frac{(\mu-L)(H-\mu)}{H-L})$, and $\delta < \epsilon/2$. Algorithm 1 solves this task if $T = \Omega(\frac{H-L}{\epsilon})$, using $O(\frac{H-L}{\epsilon} \log(\frac{H-L}{\epsilon}))$ quantum experiments in total.*

Proof. We first prove that S_i in Line 5 satisfies $S_i|\mathbf{0}\rangle|0\rangle|0\rangle = \sqrt{1 - \epsilon_i}|\mathbf{0}\rangle|0\rangle|\phi_i\rangle + \sqrt{\epsilon_i}|\text{garbage}_i\rangle$. According to the construction of U_i in Line 3 of Algorithm 1, we have

$$U_i|\mathbf{0}\rangle|0\rangle = \sqrt{q_i}|\psi_1^{(i)}\rangle|1\rangle + \sqrt{1 - q_i}|\psi_0^{(i)}\rangle|0\rangle \quad (14)$$

for some unit states $|\psi_1^{(i)}\rangle$ and $|\psi_0^{(i)}\rangle$. Consider the V_i in Line 4 where we append a qubit to the register. For any $b \in \{0, 1\}$ we have

$$\begin{aligned}
 \langle b| \langle 0| \langle 0| V_i | \mathbf{0} \rangle | 0 \rangle | 0 \rangle &= ((U_i \otimes I) | \mathbf{0} \rangle | 0 \rangle | b \rangle)^\dagger (I \otimes \text{CNOT}) (U_i \otimes I) | \mathbf{0} \rangle | 0 \rangle | 0 \rangle \\
 &= \left(\sqrt{q_i} \langle b| \langle 1| \langle \psi_1^{(i)} | + \sqrt{1-q_i} \langle b| \langle 0| \langle \psi_0^{(i)} | \right) \left(\sqrt{q_i} | \psi_1^{(i)} \rangle | 1 \rangle | 1 \rangle + \sqrt{1-q_i} | \psi_0^{(i)} \rangle | 0 \rangle | 0 \rangle \right) \\
 &= \begin{cases} q_i & b = 1 \\ 1 - q_i & b = 0, \end{cases} \tag{15}
 \end{aligned}$$

which implies that

$$\begin{aligned}
 V_i | \mathbf{0} \rangle | 0 \rangle | 0 \rangle &= \sqrt{2q_i^2 - 2q_i + 1} | \mathbf{0} \rangle | 0 \rangle \left(\frac{q_i}{\sqrt{2q_i^2 - 2q_i + 1}} | 1 \rangle + \frac{1 - q_i}{\sqrt{2q_i^2 - 2q_i + 1}} | 0 \rangle \right) \\
 &\quad + \sqrt{2q_i - 2q_i^2} | \text{garbage}_i \rangle, \tag{16}
 \end{aligned}$$

where $|\text{garbage}_i\rangle$ is a unit garbage state and $(I \otimes \langle 0| \langle \mathbf{0}|) |\text{garbage}_i\rangle = 0$. Moreover, we define

$$|\phi_i\rangle = \frac{q_i}{\sqrt{2q_i^2 - 2q_i + 1}} | 1 \rangle + \frac{1 - q_i}{\sqrt{2q_i^2 - 2q_i + 1}} | 0 \rangle, \quad |s_i\rangle = V_i | \mathbf{0} \rangle | 0 \rangle | 0 \rangle. \tag{17}$$

Under these notations, we have

$$(| \mathbf{0} \rangle | 0 \rangle \langle 0| \langle \mathbf{0}| \otimes I) V_i | \mathbf{0} \rangle | 0 \rangle | 0 \rangle = \sqrt{2q_i^2 - 2q_i + 1} | \mathbf{0} \rangle | 0 \rangle | \phi_i \rangle. \tag{18}$$

Together with Lemma 11 and the fact that $\sqrt{2q_i^2 - 2q_i + 1} \geq \frac{1}{\sqrt{2}}$, we know that S_i in Line 5 satisfies

$$S_i | \mathbf{0} \rangle | 0 \rangle | 0 \rangle = \sqrt{1 - \epsilon_i} | \mathbf{0} \rangle | 0 \rangle | \phi_i \rangle + \sqrt{\epsilon_i} | \text{garbage}_i \rangle, \tag{19}$$

where $\epsilon_i \leq \epsilon'$ and S_i contains $O(\log \frac{1}{\epsilon'}) = O(\log(\frac{H-L}{\epsilon}))$ calls to V_i .

Let

$$q = \frac{\mu - L}{H - L} \in [0, 1], \quad |\phi\rangle = \frac{q}{\sqrt{2q^2 - 2q + 1}} | 1 \rangle + \frac{1 - q}{\sqrt{2q^2 - 2q + 1}} | 0 \rangle,$$

and S be a unitary such that

$$S | \mathbf{0} \rangle | 0 \rangle | 0 \rangle = | \mathbf{0} \rangle | 0 \rangle | \phi \rangle. \tag{20}$$

Performing an amplitude estimation using $\{S_i\}_{i=1}^T$ provides a result similar to an amplitude estimation using S , and thus provides a mean estimation with additive error $O(\epsilon)$. See the details in [14] appendix A.1.

Each V_i uses two quantum experiments, each S_i uses $O(\log(\frac{H-L}{\epsilon}))$ calls to V_i , and C' uses $M = O(\frac{H-L}{\epsilon})$ calls to controlled S_i . Therefore, the total number of quantum experiments is $O(\frac{H-L}{\epsilon} \log(\frac{H-L}{\epsilon}))$. ◀

► **Remark 13.** For every $i \in [T]$, S_i can be seen as an approximation of unitary S . The slight difference δ among different μ_i only causes a part of approximation error which is bounded by ϵ . Therefore, this difference is tolerable in our algorithm. See [14] equation (73) and (78) for more details.

3.2 Mean Estimation of Sub-Gaussian Random Variables

In this subsection, we consider the quantum non-identical mean estimation problem of sub-Gaussian random variables.

■ **Algorithm 2** Mean Estimation of Mean-Bounded sub-Gaussian Random Variable.

- 1: **Input:** sequence of random variable oracle $\{O_{X_i}\}_{i=1}^T$, accuracy ϵ , mean difference δ , repetition parameter m , upper bound for mean R , sub-Gaussian parameter K
- 2: **Output:** mean estimation $\tilde{\mu}$
- 3: Let $\Delta = K \max \left\{ \sqrt{4 \log\left(\frac{128K}{\epsilon}\right)}, \sqrt{2 \log\left(\frac{32R}{\epsilon}\right)} \right\}$, $L = -R - \Delta$, $H = R + \Delta$
- 4: Construct unitary $O_{\tilde{X}_i}$

$$\begin{aligned} O_{\tilde{X}_i} : |\mathbf{0}\rangle|\mathbf{0}\rangle &\xrightarrow{O_{X_i} \otimes I} \sum_{x \in E_i} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle |\mathbf{0}\rangle \\ &\xrightarrow{\text{CNOT}} \sum_{x \in [L, H]} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle |x\rangle + \sum_{x \in E_i \setminus [L, H]} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle |\mathbf{0}\rangle \end{aligned}$$

- 5: Output $\tilde{\mu} = \text{Algorithm 1}(\{O_{\tilde{X}_i}\}_{i=1}^T, \text{accuracy } \epsilon, \text{mean difference } \delta = \epsilon/2, m, L, H)$

► **Definition 14.** A random variable X is sub-Gaussian with parameter K if for all $t \geq 0$

$$\mathbb{P}[|X - \mathbb{E}[X]| \geq t] \leq 2 \exp\left(-\frac{t^2}{2K^2}\right). \quad (21)$$

We first give a quantum algorithm estimating the mean of non-identically distributed sub-Gaussian random variables with quadratic speed-up if the mean of the random variables are bounded by their sub-Gaussian parameter. This case can be reduced to the case of bounded random variables by truncation. Then, we show that this algorithm can be generalized to any sub-Gaussian random variable.

► **Lemma 15.** Suppose all random variables X_1, \dots, X_T in Task 8 are sub-Gaussian with parameter K and their mean satisfies $|\mu_i| \leq R$, $R \leq K$. Let $m, R, K, \epsilon, \delta$ in Algorithm 2 satisfies that $m = \Omega\left(\log\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon}\right)\right)$, $\epsilon = O(K)$, and $\delta < \epsilon/4$. Algorithm 2 solves Task 8 if $T = \Omega\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon}\right)$, using $O\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon} \log\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon}\right)\right)$ quantum experiments in total.

Quantum random variable \tilde{X}_i generated by oracle $O_{\tilde{X}_i}$ in Algorithm 2 is a truncated version of X_i . Calculation shows that the mean difference is within $\frac{\epsilon}{2}$, thus Algorithm 2 provides an estimation with $O(\epsilon)$ additive error.

Proof. See [14] appendix A.2. ◀

For general sub-Gaussian distributions, we first use $O(1)$ classical samples to estimate the mean of these sub-Gaussian random variables up to additive error $K/2$, and then shift the random variables by subtracting the approximate mean so that the shifted random variables have mean bounded by their sub-Gaussian parameter. After that, we can use Lemma 15 to estimate the mean of the shifted random variables.

► **Theorem 16.** Assume all random variables X_1, \dots, X_T in Task 8 are sub-Gaussian with parameter K . Let m, K, δ, ϵ in Algorithm 3 satisfy that $m = \Omega\left(\log\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon}\right)\right)$, $\epsilon = O(K)$, and $\delta < \epsilon/4$. Algorithm 3 solves Task 8 if $T = \Omega\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon}\right)$, using $O\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon} \log\left(\frac{K\sqrt{\log\left(\frac{K}{\epsilon}\right)}}{\epsilon}\right)\right)$ quantum experiments in total.

■ **Algorithm 3** Mean Estimation of sub-Gaussian Random Variable.

- 1: **Input:** sequence of random variable oracle $\{O_{X_i}\}_{i=1}^T$, accuracy ϵ , repetition parameter m , sub-Gaussian parameter K
- 2: **Output:** mean estimation $\tilde{\mu}$
- 3: Perform $N = \lceil 8 \log(20) \rceil$ times classical experiments on arbitrary X_i and let the average of the samples be $\hat{\mu}$
- 4: Construct unitary $O_{X'_i}$

$$\begin{aligned} O_{X'_i} : |\mathbf{0}\rangle|\mathbf{0}\rangle &\xrightarrow{O_{X_i} \otimes I} \sum_{x \in E_i} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x\rangle |\mathbf{0}\rangle \\ &\longrightarrow \sum_{x \in E_i} \sqrt{p_i(x)} |\psi_x^{(i)}\rangle |x - \hat{\mu}\rangle \end{aligned}$$

- 5: Output $\tilde{\mu} = \text{Algorithm 2}(\{O_{X'_i}\}_{i=1}^T, \text{accuracy } \epsilon, \text{mean difference } \delta = \epsilon/4, m, \text{upper bound for mean } R = K, \text{sub-Gaussian parameter } K)$

Proof. Classical experiment in Line 3 can be naturally implemented by quantum access to random variable. For any $i \in [T]$, by applying O_{X_i} to $|\mathbf{0}\rangle$ and measuring the second register in computational basis, we can get a classical sample of X_i . Since $\hat{\mu}$ is the average value of $N = \lceil 8 \log(20) \rceil$ samples, by the Hoeffding inequality for sub-Gaussian distributions [21], we have

$$\mathbb{P}[|\hat{\mu} - \mathbb{E}[\hat{\mu}]| \geq \frac{K}{2}] \leq 2 \exp\left(-\frac{N}{2K^2} \frac{K^2}{4}\right) \leq \frac{1}{10}. \quad (22)$$

In addition, since $|\mu_i - \mu| \leq \delta$ for all $i \in [T]$, we have

$$|\mathbb{E}[\hat{\mu}] - \mu| \leq \delta. \quad (23)$$

O'_{X_i} can be seen as quantum query to random variable $X'_i = X_i - \hat{\mu}$. With probability at least $\frac{9}{10}$, we have

$$|\mathbb{E}[X'_i]| = |\mathbb{E}[X_i] - \hat{\mu}| \leq |\mathbb{E}[X_i] - \mathbb{E}[\hat{\mu}]| + |\hat{\mu} - \mathbb{E}[\hat{\mu}]| \leq \delta + \frac{K}{2} \leq K. \quad (24)$$

Therefore, by Lemma 15 with $R = K$, $m = \Omega\left(\log\left(\frac{K\sqrt{\log(\frac{K}{\epsilon})}}{\epsilon}\right)\right)$ and $X'_i = X_i - \hat{\mu}$, we can estimate $\mu - \hat{\mu}$ with additive error $O(\epsilon)$ with probability at least $\frac{4}{5}$ using $O\left(\frac{K\sqrt{\log(\frac{K}{\epsilon})}}{\epsilon} \log\left(\frac{K\sqrt{\log(\frac{K}{\epsilon})}}{\epsilon}\right)\right)$ quantum experiments. Subtracting $\hat{\mu}$ from the estimate gives the final output of the algorithm which is an ϵ -additive estimate of μ with probability at least $\frac{4}{5} \cdot \frac{9}{10} \geq \frac{2}{3}$. ◀

4 Lower Bound

In this section, we prove sample complexity lower bounds for the quantum non-identical mean estimation problem in Task 8.

Let m be the repetition parameter defined Task 8. In Section 4.1, we give a sample complexity lower bound for $m = 1$, and show there is no quantum speed-up compared to classical algorithms. In Section 4.2, we give a sample complexity lower bound for $m \geq 1$.

4.1 Lower Bound for $m = 1$

Let X be a finite random variable with support E . Let (\mathcal{H}, O_X) be a quantum random variable in Definition 6, i.e.,

$$O_X|\mathbf{0}\rangle = \sum_{x \in E} \sqrt{p(x)}|\psi_x\rangle|x\rangle, \quad (25)$$

and we denote the output state by $|\psi_X\rangle$. A p -parallel query to O_X is to apply the unitary $O_X^{\otimes q}$ or $O_X^{\dagger \otimes q}$ for $q \leq p$.

Note that Eq. (25) only restricts the outcome of applying O_X on $|\mathbf{0}\rangle$, so the quantum random variable encoding the same X can be different. Throughout Section 4.1, we assume all quantum random variables encode the same finite random variable X . Given that $m = 1$, the algorithm can perform only one quantum experiment for each quantum random variable.

We use the quantum query model to analyze the sample complexity of the quantum non-identical mean estimation since every quantum experiment can be regarded as a query to the oracle O_X . A T -query quantum algorithm starts from an all-0 state $|\mathbf{0}\rangle_Q|\mathbf{0}\rangle_W$, and then interleaves fixed unitary operations U_0, U_1, \dots, U_T with queries. Suppose different oracles are queried at different time, and we denote the t -th oracle queried by the algorithm as $O_X^{(t)}$. Without loss of generality, we assume that all queries are applied to register $|\mathbf{0}\rangle_Q$ and U_0, U_1, \dots, U_T are applied to $|\mathbf{0}\rangle_Q|\mathbf{0}\rangle_W$. Whether to apply $O_X^{(t)}$ or $(O_X^{(t)})^\dagger$ needs to be determined in advance, and the choices can be represented by T boolean variables $a_1, \dots, a_T \in \{-1, 1\}$ such that

$$(O_X^{(t)})^{a_t} = \begin{cases} O_X^{(t)} & \text{if } a_t = 1, \\ (O_X^{(t)})^\dagger & \text{if } a_t = -1. \end{cases} \quad (26)$$

For any $1 \leq t \leq T$, let

$$|\psi^{(t)}\rangle := U_t(O_X^{(t)})^{a_t} \dots (O_X^{(1)})^{a_1} U_0|\mathbf{0}\rangle_Q|\mathbf{0}\rangle_W. \quad (27)$$

Hence the final state of the algorithm is $|\psi^{(T)}\rangle$.

At the end of the algorithm, we will measure $|\psi^{(T)}\rangle$ and let the projection onto the correct outputs be Π_c , and the success probability of the algorithm is hence

$$\|\Pi_c|\psi^{(T)}\rangle\|^2. \quad (28)$$

4.1.1 Reduction to Low-depth Quantum Algorithms

For a quantum circuit with oracles, the query depth is the maximum number of queries on any path from an input qubit to an output qubit. In this section, we prove that the behavior of a quantum algorithm querying T non-identical oracles can be simulated by a low query depth quantum algorithm with the same number of queries. Actually, we will show that the behavior of the algorithm can be simulated by a quantum circuit using two T -parallel queries.

For any $1 \leq t \leq T$, let

$$|\phi_{\text{beg}}^{(t)}\rangle := \begin{cases} |\mathbf{0}\rangle & \text{if } a_t = 1, \\ |\psi_X\rangle & \text{if } a_t = -1, \end{cases} \quad (29)$$

$$|\phi_{\text{end}}^{(t)}\rangle := \begin{cases} |\psi_X\rangle & \text{if } a_t = 1, \\ |\mathbf{0}\rangle & \text{if } a_t = -1, \end{cases} \quad (30)$$

so that

$$(O_X^{(t)})^{a_t} |\phi_{\text{beg}}^{(t)}\rangle = |\phi_{\text{end}}^{(t)}\rangle. \quad (31)$$

This is the only subspace that $(O_X^{(t)})^{a_t}$'s behavior is fixed and defined by Eq. (25).

For any $1 \leq t \leq T$, let

$$\Pi_{\text{beg}}^{(t)} := |\phi_{\text{beg}}^{(t)}\rangle\langle\phi_{\text{beg}}^{(t)}| \otimes I, \quad (32)$$

and

$$|\psi_{\text{eff}}^{(t)}\rangle_{Q,W} := (O_X^{(t)})^{a_t} \Pi_{\text{beg}}^{(t)} U_{t-1} (O_X^{(t-1)})^{a_{t-1}} \Pi_{\text{beg}}^{(t-1)} \cdots U_1 (O_X^{(1)})^{a_1} \Pi_{\text{beg}}^{(1)} U_0 |\mathbf{0}\rangle_Q |\mathbf{0}\rangle_W \quad (33)$$

$$= (|\phi_{\text{end}}^{(t)}\rangle\langle\phi_{\text{beg}}^{(t)}| \otimes I) U_{t-1} \cdots U_1 (|\phi_{\text{end}}^{(1)}\rangle\langle\phi_{\text{beg}}^{(1)}| \otimes I) U_0 |\mathbf{0}\rangle_Q |\mathbf{0}\rangle_W. \quad (34)$$

These states are fixed no matter what the queries $O_X^{(t)}$ are, since all queries in Eq. (33) are applied to the subspace that its behavior is defined by Eq. (25).

We show in the following lemma that $|\psi_{\text{eff}}^{(t)}\rangle$ can be prepared by a quantum algorithm using two t -parallel queries after post-selection.

► **Lemma 17.** *Given a T -query quantum algorithm acting on registers Q and W , for any $0 \leq t \leq T$, $|\psi_{\text{eff}}^{(t)}\rangle$ defined in Eq. (33) can be prepared by another quantum circuit V_t^{low} using two t -parallel queries to any unitary oracle O_X satisfying Eq. (25) after post-selection, namely,*

$$(I_{W,Q_t} \otimes \langle\mathbf{0}|_{Q_0,\dots,Q_{t-1}}) V_t^{\text{low}} |\mathbf{0}\rangle_{W,Q_0,\dots,Q_t}, \quad (35)$$

where Q_0, \dots, Q_t are $t+1$ registers with n_Q qubits.

Proof. For all $1 \leq t < T$, from the definition of $|\psi_{\text{eff}}^{(t)}\rangle$, it can be written as

$$|\psi_{\text{eff}}^{(t)}\rangle = |\phi_{\text{end}}^{(t)}\rangle |\phi_{\text{W}}^{(t)}\rangle \quad (36)$$

for some unnormalized state $|\phi_{\text{W}}^{(t)}\rangle$, then we have

$$|\phi_{\text{end}}^{(t+1)}\rangle |\phi_{\text{W}}^{(t+1)}\rangle = |\psi_{\text{eff}}^{(t+1)}\rangle = (|\phi_{\text{end}}^{(t+1)}\rangle\langle\phi_{\text{beg}}^{(t+1)}| \otimes I) U_t |\phi_{\text{end}}^{(t)}\rangle |\phi_{\text{W}}^{(t)}\rangle. \quad (37)$$

Apply $\langle\phi_{\text{end}}^{(t+1)}| \otimes I$ to both sides we have

$$|\phi_{\text{W}}^{(t+1)}\rangle = (\langle\phi_{\text{beg}}^{(t+1)}| \otimes I) U_t |\phi_{\text{end}}^{(t)}\rangle |\phi_{\text{W}}^{(t)}\rangle. \quad (38)$$

Define

$$|\psi_{\text{eff}}^{(0)}\rangle = |\mathbf{0}\rangle |\mathbf{0}\rangle, \quad |\phi_{\text{end}}^{(0)}\rangle = |\mathbf{0}\rangle, \quad |\phi_{\text{W}}^{(0)}\rangle = |\mathbf{0}\rangle, \quad (39)$$

so that Eq. (36) and Eq. (38) also hold for $t = 0$.

To construct the required circuit, We prove the following stronger statement.

► **Statement 18.** *Let O_X be any unitary satisfying Eq. (25), and $U_0^{\text{low}}, \dots, U_T^{\text{low}}$ be a sequence of quantum circuits satisfying $U_0^{\text{low}} = I$ and*

$$U_{t+1}^{\text{low}} = \begin{cases} ((U_t)_{Q_t,W} \otimes I) \cdot (U_t^{\text{low}} \otimes (O_X)_{Q_{t+1}}) & \text{if } a_{t+1} = 1, \\ ((U_t)_{Q_t,W} \otimes I) \cdot (U_t^{\text{low}} \otimes I_{Q_{t+1}}) & \text{if } a_{t+1} = -1, \end{cases} \quad (40)$$

9:16 Quantum Non-Identical Mean Estimation

for all $0 \leq t < T$. The quantum circuit U_t^{low} can prepare $|\psi_{\text{eff}}^{(t)}\rangle$ after post-selection, namely,

$$|\psi_{\text{eff}}^{(t)}\rangle = \left(I_{W, Q_t} \bigotimes_{i=1}^t \langle \phi_{\text{beg}}^{(i)} |_{Q_{i-1}} \right) U_t^{\text{low}} |\mathbf{0}\rangle_{W, Q_0, \dots, Q_t}, \quad (41)$$

for any $0 \leq t \leq T$.

Proof. See [14] appendix B.1. ◀

The number of queries in U_t^{low} is $|\{a_i = 1 \mid i \in [t]\}|$. Let

$$V_t^{\text{low}} = \bigotimes_{1 \leq i \leq t, a_i = -1} (O_X^\dagger)_{Q_i} U_t^{\text{low}}, \quad (42)$$

then from Eq. (41) we have

$$(I_{W, Q_t} \otimes \langle \mathbf{0} |_{Q_0, \dots, Q_{t-1}}) V_t^{\text{low}} |\mathbf{0}\rangle_{W, Q_0, \dots, Q_t}, \quad (43)$$

for all $0 \leq t \leq T$.

The number of queries in V_t^{low} is

$$|\{a_i = 1 \mid i \in [t]\}| + |\{a_i = -1 \mid i \in [t]\}| = t. \quad (44)$$

Conditioning on the state in registers Q_0, \dots, Q_{t-1} to be $|\mathbf{0}\rangle$, V_t^{low} prepares $|\psi_{\text{eff}}^{(t)}\rangle_{Q_t, W}$ and uses two t -parallel queries. ◀

Next, we demonstrate that $U_T |\psi_{\text{eff}}^{(T)}\rangle$ is the only useful component in the final state $|\psi^{(T)}\rangle$, since other parts can be controlled by $O_X^{(t)}$ to make the result worse. Before that, we prove the following useful lemma.

► **Lemma 19.** *For any T -query quantum algorithm acting on registers Q, W , and any finite random variable X on (Ω, p) , if $\dim \mathcal{H}_Q > 2 \dim \mathcal{H}_W$, there exists a sequence of quantum random variables $(\mathcal{H}_Q, O_X^{(1)}), \dots, (\mathcal{H}_Q, O_X^{(T-1)})$ such that for any $0 \leq t < T$*

$$|\psi^{(t)}\rangle = |\phi_{\text{beg}}^{(t+1)}\rangle |\phi_W^{(t+1)}\rangle + |\psi_{\perp}^{(t)}\rangle, \quad (45)$$

for some unnormalized state $|\psi_{\perp}^{(t)}\rangle$ orthogonal to $|\phi_{\text{beg}}^{(t+1)}\rangle \otimes \mathcal{H}_W$.

Proof. By induction. See the details in [14] appendix B.2. ◀

Now we prove that $U_T |\psi_{\text{eff}}^{(T)}\rangle$ is the only useful component in the final state $|\psi^{(T)}\rangle$.

► **Lemma 20.** *Suppose that X is a finite random variable. For any T -query quantum algorithm acting on registers Q, W , and any projection Π_c , if $\dim \mathcal{H}_Q > 2 \dim \mathcal{H}_W$ and $\dim \mathcal{H}_Q \geq 2 \dim \text{Im}(\Pi_c)$, then there exists a sequence of quantum random variables $(\mathcal{H}_Q, O_X^{(1)}), \dots, (\mathcal{H}_Q, O_X^{(T)})$ such that*

$$\|\Pi_c |\psi^{(T)}\rangle\|^2 = \|\Pi_c U_T |\psi_{\text{eff}}^{(T)}\rangle\|^2. \quad (46)$$

Proof. Note that

$$|\psi^{(T)}\rangle = U_T (O_X^{(T)})^{a_T} |\psi^{(T-1)}\rangle \quad (47)$$

$$= U_T (O_X^{(T)})^{a_T} (|\phi_{\text{beg}}^{(T)}\rangle |\phi_W^{(T)}\rangle + |\psi_{\perp}^{(T-1)}\rangle) \quad (48)$$

$$= U_T |\phi_{\text{end}}^{(T)}\rangle |\phi_W^{(T)}\rangle + U_T (O_X^{(T)})^{a_T} |\psi_{\perp}^{(T-1)}\rangle \quad (49)$$

$$= U_T |\psi_{\text{eff}}^{(T)}\rangle + U_T (O_X^{(T)})^{a_T} |\psi_{\perp}^{(T-1)}\rangle. \quad (50)$$

To satisfy Eq. (46), we need to find a unitary operator $O_X^{(T)}$ such that

$$\Pi_c U_T (O_X^{(T)})^{a_T} |\psi_{\perp}^{(T-1)}\rangle = 0, \quad (51)$$

which means

$$(O_X^{(T)})^{a_T} |\psi_{\perp}^{(T-1)}\rangle \in (U_T^{\dagger} \text{Im}(\Pi_c))^{\perp}. \quad (52)$$

By the same method in the proof of Lemma 19, we can construct oracle $O_X^{(T)}$. By similar argument to Lemma 19, we can prove that if

$$\dim \mathcal{H}_Q > \dim \mathcal{H}_W + \dim \text{Im}(\Pi_c), \quad (53)$$

there exists $O_X^{(T)}$ such that Eq. (46) holds. By assumptions that $\dim \mathcal{H}_Q > 2 \dim \mathcal{H}_W$ and $\dim \mathcal{H}_Q \geq 2 \dim \text{Im}(\Pi_c)$, we can conclude that Eq. (46) holds. \blacktriangleleft

In conclusion, there exists a sequence of quantum random variables such that the output of a T -query quantum algorithm can be simulated by a quantum algorithm using two T -parallel queries.

► **Theorem 21.** *For any T -query quantum algorithm \mathcal{A} acting on registers Q , W , and any projection Π_c , suppose that $\dim \mathcal{H}_Q > 2 \dim \mathcal{H}_W$ and $\dim \mathcal{H}_Q \geq 2 \dim \text{Im}(\Pi_c)$. Let $|\psi^{(T)}\rangle$ be the final state of the algorithm. There exists another quantum circuit U^{low} using two T -parallel queries such that for any finite random variable X , there is a sequence of quantum random variables $(\mathcal{H}_Q, O_X^{(1)}), \dots, (\mathcal{H}_Q, O_X^{(T)})$ satisfying*

$$\|\Pi_c |\psi^{(T)}\rangle\|^2 = \|(\Pi_c \otimes \langle \mathbf{0} |_{Q_0, \dots, Q_{T-1}}) U^{\text{low}} | \mathbf{0} \rangle_{W, Q_0, \dots, Q_T}\|^2, \quad (54)$$

where Q_0, \dots, Q_T are $T + 1$ registers with n_Q qubits.

Proof. Let V_T^{low} be the low-depth quantum circuit defined in Lemma 17, and U_T be the unitary in algorithm \mathcal{A} at time step T . By Lemma 17, the unitary $U^{\text{low}} = ((U_T)_{Q_T, W} \otimes I) V_T^{\text{low}}$ satisfies

$$(I \otimes \langle \mathbf{0} |_{Q_0, \dots, Q_{T-1}}) U^{\text{low}} | \mathbf{0} \rangle_{W, Q_0, \dots, Q_T} = U_T |\psi_{\text{eff}}^{(T)}\rangle_{Q_T, W}. \quad (55)$$

By Lemma 20, there exists a sequence of quantum random variables $(\mathcal{H}_Q, O_X^{(1)}), \dots, (\mathcal{H}_Q, O_X^{(T)})$ such that

$$\|\Pi_c |\psi^{(T)}\rangle\|^2 = \|\Pi_c U_T |\psi_{\text{eff}}^{(T)}\rangle\|^2 = \|(\Pi_c \otimes \langle \mathbf{0} |_{Q_0, \dots, Q_{T-1}}) U^{\text{low}} | \mathbf{0} \rangle_{W, Q_0, \dots, Q_T}\|^2. \quad (56)$$

\blacktriangleleft

4.1.2 Lower Bounds for Low-depth Quantum Mean Estimation Algorithms

Given an input $x = x_0 \dots x_{n-1} \in \{0, 1\}^n$, the quantum query to it is a unitary O_x such that

$$O_x |i\rangle |b\rangle = |i\rangle |b \oplus x_i\rangle \quad (57)$$

for all $i \in [n]$ and $b \in \{0, 1\}$.

9:18 Quantum Non-Identical Mean Estimation

The *approximate counting* problem is that given O_x , output an estimate of $|x|$ up to error ϵ with high probability. From another perspective, we can think of $[n]$ as a sample space Ω with uniform distribution P , and $X: \Omega \rightarrow \{0, 1\}$ is a Bernoulli random variable such that $X(i) = x_i$, and the mean of X is

$$p = \frac{|x|}{n}. \quad (58)$$

Note that

$$|0\rangle|0\rangle \xrightarrow{\text{Hardmard gates}} \sum_{i=1}^n \frac{1}{\sqrt{n}} |i\rangle|0\rangle \xrightarrow{I \otimes O_x} \sum_{i=1}^n \frac{1}{\sqrt{n}} |i\rangle|X(i)\rangle, \quad (59)$$

which means we can implement the oracle to X with one query to O_x . Hence, the approximate counting problem can be reduced to the mean estimation problem.

A k -parallel query call to x is

$$O_x^{\otimes k} |i_1, \dots, i_k, b_1, \dots, b_k\rangle = |i_1, \dots, i_k, b_1 \oplus x_{i_1}, \dots, b_k \oplus x_{i_k}\rangle \quad (60)$$

[4] proved a k -parallel query lower bound of the approximate counting problem.

► **Theorem 22** ([4]). *For any quantum query algorithm and boolean string $x \in \{0, 1\}^n$,*

$$\Omega\left(\frac{\binom{n-|x|}{\epsilon n} \binom{|x|+\epsilon n}{|x|}}{k \binom{n-|x|-1}{\epsilon n-1} \binom{|x|+\epsilon n-1}{|x|}}\right) = \Omega\left(\frac{p(1-p)}{\epsilon^2 k}\right) \quad (61)$$

k -parallel queries to O_x is necessary to estimate $p = \frac{|x|}{n}$ to within additive error ϵ .

By Theorem 22, if we want to use constant k -parallel queries to estimate p up to additive error ϵ , k needs to satisfy

$$\frac{p(1-p)}{\epsilon^2 k} = O(1), \quad (62)$$

which means

$$k = \Omega\left(\frac{p(1-p)}{\epsilon^2}\right). \quad (63)$$

Now we give a sample complexity lower bound of algorithms solving Task 8 with $m = 1$ using Theorem 21. The difficulty of directly applying Theorem 21 is that it requires $\dim \text{Im}(\Pi_c)$ to be small. To resolve it, we prove that any quantum mean estimator can be modified to recover the state in query register Q to $|0\rangle$ with a small overhead so that correct answers lie in a much smaller subspace.

► **Theorem 23.** *Suppose all random variables in Task 8 have variance bounded by σ^2 , and $|\mu| \leq R$. Let \mathcal{A} be a quantum query algorithm acting on registers Q, W solving the quantum non-identical mean estimation problem defined in Task 8 with repetition parameter $m = 1$ and accuracy $\epsilon/2$. Suppose that $\frac{1}{2}n_Q > n_W + 2 \log\left(\frac{2R}{\epsilon}\right) + 1$, then it requires $T = \Omega\left(\frac{\sigma^2}{\epsilon^2}\right)$ for the existence of such an algorithm \mathcal{A} , and \mathcal{A} needs $T = \Omega\left(\frac{\sigma^2}{\epsilon^2}\right)$ quantum experiments.*

Proof. Use the uncomputation trick to combine Theorem 21 and Theorem 22. See [14] appendix B.3. ◀

4.1.3 Implication for Quantum Linear Systems

As mentioned in the introduction, we can possibly estimate A by the following procedure.

For fixed integers $t_0, \gamma = \Theta(\log(\sqrt{n}/\delta))$ and any $0 \leq t < n$, suppose we have a register storing $|\psi_{t_0+2\gamma t}\rangle$. We measure $|\psi_{t_0+2\gamma t}\rangle$ to obtain a classical state $x_{t_0+2\gamma t}$, and get $|\psi_{t_0+2\gamma t+1}\rangle$ as the second register of $U_f|\psi_{t_0+2\gamma t}\rangle|0\rangle$ (note that $|\psi_{t_0+2\gamma t}\rangle$ has collapsed after the measurement), which encodes the randomness of $x_{t_0+2\gamma t+1}$ given $x_{t_0+2\gamma t}$. Similarly, we can also obtain $|\psi_{t_0+2\gamma t+1}^{-1}\rangle$ by querying U_f^{-1} . After that, we compute $U_f|\psi_{t_0+2\gamma t+1}\rangle|0\rangle$ and collect the second register as $|\psi_{t_0+2\gamma t+2}\rangle$, and do this computation for all $t_0 + 2\gamma t + 1$ to $t_0 + 2(\gamma + 1)t - 1$. Then we let $t = t + 1$ and repeat this process.

After such process, we have n classical samples at even steps $X_{t_0} := [x_{t_0}, x_{t_0+2\gamma}, \dots, x_{t_0+2n\gamma-2}] \in \mathbb{R}^{n \times n}$, and n quantum samples at odd steps. It holds that X_{t_0} is full rank with probability 1 given that

$$AX_{t_0} = [x_{t_0+1}, \dots, x_{t_0+2n\gamma-1}] + W_{t_0} + Z_{t_0} \quad (64)$$

where W_{t_0} is a zero-mean noise matrix and $\|Z_{t_0}\|_F \leq O(\delta)$. The matrix Z_{t_0} denotes the difference between $\mathbb{E}[x_{t_0+2\gamma t+1} \mid x_{t_0+2\gamma t}]$ and $\mathbb{E}[x_{t_0+2\gamma t+1} \mid x_{t_0+2\gamma t}, x_{t_0+2\gamma(t+1)}]$, which are close since $\|A^n\|_2 = O(-\exp(n))$. We define the quantum unitary U_{t_0} as

$$U_{t_0}|0\rangle := \int_W \sqrt{f_{t_0}(W)} |\psi_{t_0+1}, \dots, \psi_{t_0+2n\gamma-1}\rangle X_{t_0}^{-1} dW \quad (65)$$

where $f_{t_0}(W)$ is the pdf of random matrix $W_{t_0} X_{t_0}^{-1}$. Then we can use the quantum samples collected at steps $t_0 + 1, \dots, t_0 + 2n\gamma - 1$ as the return of query to U_{t_0} (or $U_{t_0}^{-1}$). Note that the mean of the random variable encoded by U_{t_0} is $O(\delta)$ -close to A in Frobenius norm according to (64). However, the distribution encoded in U_{t_0} are different for different t_0 since X_{t_0} are different. The lower bound presented in the previous section shows that this methods cannot achieve a desired quantum speed-up since the oracle U_{t_0} can only be queried once for each t_0 .

4.2 Lower Bounds for $m \geq 1$

Given a boolean string $|x| \in \{0, 1\}^n$ and $k \in [n]$, the task of distinguishing $|x| = k$ and $|x| = k + 1$ or $|x| = k - 1$ can be reduced to estimating $\frac{|x|}{n}$ to within $\frac{1}{n}$ additive error, which can be regarded as a mean estimation problem. Therefore, the query complexity lower bound for the first problem is also a lower bound for the second problem. As a result, we first prove the query complexity lower bound of the first problem given non-identical oracles.

We use the same quantum query algorithm model in Section 4.1, where the algorithm pre-determines U_0, \dots, U_T and needs to distinguish the cases between $|x| = k$ and $|x| = k + 1$ or $k - 1$ for any $1 \leq k < n$.

► **Lemma 24.** *Given a sequence of oracles O_{x_1}, \dots, O_{x_T} encoding boolean strings x_1, \dots, x_T in $\{0, 1\}^n$, suppose all strings have the same Hamming weight w and the algorithm can query each oracle at most m times in turn. For any $1 \leq k < n$ and $m = O(\sqrt{n})$, any quantum algorithm needs $\Omega(\frac{n}{m})$ queries in total to distinguish between $w = k$ and $w = k - 1$ or $k + 1$ with high probability.*

Proof. See [14] appendix B.4. ◀

Now we give a sample complexity lower bound of the quantum non-identical mean estimation problem with repetition parameter m .

► **Theorem 25.** *Suppose all random variables in Task 8 are Bernoulli random variables with mean $\mu \in (0, 1)$ such that $\epsilon \leq \mu(1 - \mu)$ and $\epsilon = O(\frac{1}{m^2})$. It requires $T = \Omega(\frac{1}{\epsilon m^2})$ if there exists a quantum algorithm which queries each random variable at most m times in turn solves this problem. Any such quantum query algorithm needs $mT = \Omega(\frac{1}{\epsilon m})$ quantum experiments in total.*

Proof. Let $n = \frac{1}{\epsilon}$ and $k = \mu n$. Since $\epsilon \leq \mu(1 - \mu)$, we have $1 \leq k \leq n - 1$. Given a boolean string $|x| \in \{0, 1\}^n$, the task of distinguishing $|x| = k$ and $|x| = k + 1$ or $|x| = k - 1$ can be reduced to estimating $\frac{|x|}{n}$ to within $\frac{1}{n}$ additive error. The latter problem can be regarded as estimating the mean of a Bernoulli random variable X to within additive error $\epsilon = \frac{1}{n}$. Since one query to O_X can be implemented by one query to O_x , the query complexity lower bound for the first problem is also a lower bound for the second problem. From $\epsilon = O(\frac{1}{m^2})$, we have $m = O(\frac{1}{\sqrt{\epsilon}}) = O(\sqrt{n})$. Therefore, by Lemma 24, any quantum algorithm solving the quantum non-identical mean estimation problem with repetition parameter m needs $\Omega(\frac{1}{\epsilon m})$ quantum experiments in total. ◀

References

- 1 Charles H. Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989. doi:10.1137/0218053.
- 2 Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *Fortschritte der Physik: Progress of Physics*, 46(4-5):493–505, 1998.
- 3 Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.
- 4 Paul Burchard. Lower bounds for parallel quantum counting. *CoRR*, abs/1910.04555, 2019. doi:10.48550/arXiv.1910.04555.
- 5 Shouvanik Chakrabarti, Rajiv Krishnakumar, Guglielmo Mazzola, Nikitas Stamatopoulos, Stefan Woerner, and William J. Zeng. A threshold for quantum advantage in derivative pricing. *Quantum*, 5:463, 2021. doi:10.22331/q-2021-06-01-463.
- 6 Arjan Cornelissen and Yassine Hamoudi. A sublinear-time quantum algorithm for approximating partition functions. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1245–1264. SIAM, 2023. doi:10.1137/1.9781611977554.ch46.
- 7 Paul Dagum, Richard Karp, Michael Luby, and Sheldon Ross. An optimal algorithm for Monte Carlo estimation. *SIAM Journal on Computing*, 29(5):1484–1496, 2000. doi:10.1137/S0097539797315306.
- 8 Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, 20(4):633–679, 2020. doi:10.1007/s10208-019-09426-y.
- 9 Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996. doi:10.1145/237814.237866.
- 10 Lov K. Grover and Jaikumar Radhakrishnan. Quantum search for multiple items using parallel queries. *arXiv preprint*, 2004. arXiv:quant-ph/0407217.
- 11 Yassine Hamoudi. Quantum sub-gaussian mean estimator. *CoRR*, abs/2108.12172, 2021. doi:10.48550/arXiv.2108.12172.
- 12 Yassine Hamoudi and Frédéric Magniez. Quantum chebyshev’s inequality and applications. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, 2019.
- 13 Peter Hoyer, Troy Lee, and Robert Spalek. Negative weights make adversaries stronger. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 526–535, 2007.

- 14 Jiachen Hu, Tongyang Li, Xinzhao Wang, Yecheng Xue, Chenyi Zhang, and Han Zhong. Quantum non-identical mean estimation: Efficient algorithms and fundamental limits, 2024. [arXiv:2405.12838](https://arxiv.org/abs/2405.12838).
- 15 Stacey Jeffery, Frederic Magniez, and Ronald De Wolf. Optimal parallel quantum query algorithms. *Algorithmica*, 79:509–529, 2017. doi:10.1007/s00453-016-0206-z.
- 16 Robin Kothari and Ryan O’Donnell. Mean estimation when you have the source code; or, quantum Monte Carlo methods. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1186–1215. SIAM, 2023. doi:10.1137/1.9781611977554.ch44.
- 17 Ashley Montanaro. Quantum speedup of Monte Carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2181):20150301, 2015.
- 18 Michael A. Nielsen and Isaac L. Chuang. Quantum computation and quantum information. *Phys. Today*, 54(2):60, 2001.
- 19 Ben W Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 544–551. IEEE, 2009. doi:10.1109/FOCS.2009.55.
- 20 Max Simchowitz and Dylan Foster. Naive exploration is optimal for online LQR. In *International Conference on Machine Learning*, pages 8937–8948. PMLR, 2020. URL: <http://proceedings.mlr.press/v119/simchowitz20a.html>.
- 21 Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- 22 Theodore J. Yoder, Guang Hao Low, and Isaac L. Chuang. Fixed-point quantum search with an optimal number of queries. *Physical review letters*, 113(21):210501, 2014.
- 23 Christof Zalka. Grover’s quantum searching algorithm is optimal. *Physical Review A*, 60(4):2746, 1999.
- 24 Zhicheng Zhang, Qisheng Wang, and Mingsheng Ying. Parallel quantum algorithm for Hamiltonian simulation. *Quantum*, 8:1228, 2024. doi:10.22331/q-2024-01-15-1228.