On the complexity of unique quantum witnesses and quantum approximate counting

Anurag Anshu[1], Jonas Haferkamp[2], Yeongwoo Hwang[1], and Quynh T. Nguyen[1]

[1]Harvard University
[2]Saarland University

September 19, 2025

Abstract

We study the long-standing open question on the power of unique witnesses in quantum protocols, which asks if UniqueQMA, a variant of QMA whose accepting witness space is 1-dimensional, contains QMA under quantum reductions.

This work rules out any black-box reduction from QMA to UniqueQMA by showing a quantum oracle separation between $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ and QMA. This provides a contrast to the classical case, where the Valiant-Vazirani theorem shows a black-box randomized reduction from UniqueNP to NP, and suggests the need for studying the structure of the ground space of local Hamiltonians in distilling a potential unique witness. Via similar techniques, we show, relative to a quantum oracle, that $\mathsf{QMA}^{\mathsf{QMA}}$ cannot decide quantum approximate counting, ruling out a quantum analogue of Stockmeyer's algorithm in the black-box setting. Our results employ a subspace reflection oracle, previously considered in [AK07; AKKT20; SY23], but we introduce new tools which allow us to exploit the unique witness constraint. We also show a strong "polarization" behavior of QMA circuits, which could be of independent interest in studying quantum polynomial hierarchies.

We then ask a natural question; what structural properties of the local Hamiltonian problem can we exploit? We introduce a physically motivated candidate by showing that the ground energy of local Hamiltonians that satisfy a computational variant of the eigenstate thermalization hypothesis (ETH) can be estimated through a UniqueQMA protocol. Our protocol can be viewed as a quantum expander test in a low energy subspace of the Hamiltonian and verifies a unique entangled state across two copies of the subspace. This allows us to conclude that if UniqueQMA is not equivalent to QMA, then QMA-hard Hamiltonians must violate ETH under adversarial perturbations (more accurately, further assuming the quantum PCP conjecture if ETH only applies to extensive energy subspaces). Under the same assumption, this also serves as evidence that chaotic local Hamiltonians, such as the SYK model may be computationally simpler than general local Hamiltonians.

Contents

1 Introduction

One of the central aims of quantum complexity theory is to explore the structure of ground, and more generally low-energy, subspaces in quantum many-body systems. A key question in this field is whether these low-energy subspaces of local Hamiltonians are computationally simpler than generic quantum subspaces.[1] A long-standing problem which formalizes this inquiry is the relationship between UniqueQMA and QMA [ABOBS22]. The complexity class UniqueQMA captures verification tasks in which positive instances have a 1-dimensional accepting subspace (unique quantum witness) while QMA places no restriction on the dimension of the accepting subspace. The open question is to determine the relationship between these classes: are they equivalent under quantum reductions, or does the uniqueness promise strictly reduce computational power? On one hand, an equivalence would imply that a polynomial-time quantum algorithm can *uniquely* verify some fixed state in a low-energy subspace of local Hamiltonians, a task that intuitively seems unattainable for generic quantum subspaces. On the other, a separation would formalize the physics intuition that "natural" Hamiltonians with a non-degenerate ground space and inverse polynomial spectral gap are computationally simpler than QMA-hard Hamiltonians and thus amenable to analysis.

This question is also motivated by the insights produced by studying the class UniqueNP. Classically, the celebrated Valiant-Vazirani isolation lemma [VV86] was used to show that UniqueNP contains NP under *randomized* reductions, i.e., $\mathsf{NP} \subseteq \mathsf{RP}^{\mathsf{UniqueNP}}$ (where RP is a one-sided error version of BPP), suggesting that the uniqueness promise does not significantly reduce the computational power of NP, at least under randomized reductions. The isolation lemma has underlaid various developments in theoretical computer science, including approximate counting [Sto83], Toda's theorem [Tod91], parallel computation [MVV87], one-way functions [GS88], and average-case complexity [BDCG89]. One could hope that studying the possibility of a "quantum isolation lemma" would sharpen our understanding in a similarly wide range of quantum problems.

Limitations of unique quantum witnesses. There has been a series of works attempting to resolve this question. Aharonov et al. [ABOBS22] (first appearing on arxiv in 2008) initiated the study of UniqueQMA and possible quantum analogues of the isolation lemma. There, the authors extended the hashing-based arguments in the Valiant-Vazirani isolation lemma to the classes MA and QCMA. However, they noted the failure of a similar strategy for QMA, citing the fact that two random quantum states have exponentially small overlap. A later work [JKKS+11] sought a completely different approach to obtain a *deterministic* isolation procedure in a special case, but also fell short of achieving a unique witness protocol for general QMA problems as their strategy worked only for accepting subspaces of polynomial dimension. This led [JKKS+11] to conjecture that a *quantum* reduction is required to obtain a unique verifier, suggesting the question of whether the inclusion $\mathsf{QMA} \subseteq \mathsf{BQP}^{\mathsf{UniqueQMA}}$ holds.[2] This discussion brings us to our first result.

Theorem 1.1 (Corollary 3.7 restated)**.** *There exists a* quantum *oracle relative to which* $\mathsf{QMA} \not\subseteq \mathsf{BQP}^{\mathsf{UniqueQMA}}$.

This oracle separation implies there is no black-box quantum isolation procedure for quantum witnesses, unlike its classical counterpart [VV86], and explains the failures of all approaches in [ABOBS22; JKKS+11]. It may also be viewed as some evidence for the inequality between UniqueQMA and QMA. Furthermore, it suggests the need to exploit the physical structure of quantum low-energy spaces to construct a UniqueQMA verifier for local Hamiltonians (see our result below).

Our Theorem 1.1 is based on a quantum oracle reflecting about an unknown subspace, and the task is to determine whether the subspace is non-empty (of some large dimensionality). Although this result in a sense is a continuation of a series of works on quantum oracle separations and unitary property testing [AK07; INNRY21; SY23; ABD24], the techniques used in our work differ in that they make careful use of the structure of the quantum witness. Previous separations used techniques which were not sensitive to the presence of a

[1] From an information-theoretic point of view, the ground space of a local Hamiltonian is known to be simpler as it has a polynomial-sized tensor network approximation [SWVC07]. However this tensor network is not known to be computationally tractable.

[2] In this work, is implicit that the oracularized complexity classes refers to the promise version (i.e. UniqueQMA actually refers to PromiseUniqueQMA). The precise definition of these oracularized classes requires some care, see Section 2.

witness (e.g. replacing the witness with the maximally mixed state), which would undermine the uniqueness promise in our case. We hope that the techniques presented in this work will help develop similar query lower bounds on other witness-restricted quantum classes, such as $\mathsf{QMA}(2)$ where constructing oracle separations have been challenging.

Along the way to Theorem 1.1 we also obtain a *classical* oracle separation between $\mathsf{UniqueQMA}$ and QMA (Corollary 3.3). In fact, we can even separate NP from $\mathsf{P}^{\mathsf{UniqueQMA}}$ under the same classical oracle. However, this separation is perhaps not as interesting as Theorem 1.1, for one would a priori expect randomized or quantum reductions were needed to reduce QMA to $\mathsf{UniqueQMA}$, in analogy to the classical case. Regardless, we note that *derandomizing* the Valiant-Vazirani lemma is a long-standing open question by itself [BBF98; KVM99], and this intermediate result rules out such a derandomization in the black-box setting, see Remark 3.4. Thus, we believe the relationship between $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ and QMA is the proper question through which to study the power of unique quantum witnesses.

On classical vs. quantum oracles. Both of our separations Theorem 1.1 and Theorem 1.2 are by quantum oracles. In our approach in Theorem 1.1, the quantumness is necessary as otherwise the classical version of our oracle problem is easily shown to be contained in $\mathsf{NP} \subseteq \mathsf{BQP}^{\mathsf{UniqueQMA}}$. A similar observation can be made for Theorem 1.2. Historically, when the class or resource is classical, a classical separation can be viewed as the "gold-standard," as evidenced by the long series of works aiming to separate QCMA and QMA with a classical oracle (see the exposition in [Zha24]). However, this decision should be made on a case-by-case basis. In our context, the subspace reflection oracle captures a way in which an algorithm might probe the ground space of a Hamiltonian. Naively, the best algorithmic way to 'access' this subspace is to use quantum phase estimation to, say, implement reflections about it. There is a priori no known reason to suspect that this subspace will be well structured for worst-case Hamiltonians and thus we instantiate the ground space with a subspace in a randomly chosen basis. Another consideration when using and interpreting quantum oracle separations is that they only rule out quantum relativizing proof strategies. Indeed, [Aar08; AK25] already point out that the proof technique of explicitly representing quantum amplitudes is classically relativizing but quantumly not (and, to our knowledge, this is the only known proof technique with such property). However, this technique has only been used used to show containments in very powerful classical classes like EXP and PSPACE. Thus, quantum oracle separations are still useful in ruling out black-box reductions at lower complexity regimes like $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ and QMA.

Improved quantum approximate counting lowerbounds. Classically, the question of uniqueness is closely related to approximate counting. In [VV86], a random hash is used to isolate a single solution to an NP circuit. The celebrated Stockmeyer algorithm [Sto83] uses similar tools to give a $\mathsf{BPP}^{\mathsf{NP}}$ algorithm for estimating the number of accepting witnesses to multiplicative precision. A quantum analogue of this is to multiplicatively estimate the dimensionality of the accepting subspace of a QMA circuit [BFS11; SZ09; BCGHZ24]. The complexity class capturing this quantum approximate counting problem is QXC. As usual, one should be careful when defining quantum analogues of classical classes. However, this problem is well motivated and shown in [BCGW22] to be connected to a number of physical relevant computational problems, such as computing quantum partition functions and estimating local observables of Gibbs states. QXC was further studied in [BCGHZ24] and shown to contain the problem of estimating the Kronecker coefficients of the symmetric group. Stockmeyer's approximate counting algorithm has led to many important applications, including the sampling-counting equivalence [JVV86] and even the theoretical foundations for quantum advantage experiments [HE23]. Thus, it is natural to ask whether there exists a quantum analogue: Could quantum approximate counting be solved in (quantum) polynomial time with access to an oracle solving QMA decision problems? In essence, do we have $\mathsf{QXC} \subseteq \mathsf{BQP}^{\mathsf{QMA}}$?

The phrase "quantum approximate counting" has appeared elsewhere in the literature. In [Kup09; AIK22], the authors study SBQP, the quantum analogue of SBP (where SBP is the multiplicatively precise analogue of BPP), which captures the task of estimating the acceptance probability of a BQP machine to multiplicative precision. This is not known to be equivalent to the definition of QXC (unlike SBP which is equivalent to classical approximate counting). We do not consider SBQP in this work. Additionally, in [AKKT20],

"quantum approximate counting" is used to refer to the task of designing quantum (BQP or QMA) algorithms for classical approximate counting, given quantum queries to a Boolean oracle. [SY23] consider a quantum generalization of the problem in [AKKT20], called the quantum approximate *dimension* problem; the task here is to decide whether an unknown subspace has dimension $\leq w$ or $\geq 2w$, given access to the corresponding reflection oracle. Using the quantum approximate dimension problem, we show

Theorem 1.2 (Corollary 4.9, restated)**.** *There exists a* quantum *oracle relative to which* $\mathsf{QXC} \not\subseteq \mathsf{BQP}^{\mathsf{QMA}}$*, and in fact,* $\mathsf{QXC} \not\subseteq \mathsf{QMA}^{\mathsf{QMA}}$*.*

Our Theorem 1.2 rules out the existence of a quantum analogue of Stockmeyer's theorem in the black-box setting. This is a strengthening of a result of [SY23], where the authors show a QMA query lower bound against this task via a reduction to the classical problem in [AKKT20]. Again, the quantum oracle is necessary in our approach, because the classical version of our oracle is the classical approximate counting problem, which is contained in $\mathsf{BPP}^{\mathsf{NP}} \subseteq \mathsf{BQP}^{\mathsf{QMA}}$. The polynomial methods used in [AKKT20; SY23] do not transfer to the relativized setting in the presence of QMA-oracle calls. Moreover, prior works [AIK22; ABD24] dealing with quantum versions of the polynomial hierarchy are all based on variants of the switching lemma, which are not known to be compatible with quantum oracles. Therefore, we need to develop new techniques to handle quantum oracle calls in "stacked" complexity classes like $\mathsf{QMA}^{\mathsf{QMA}}$.

Along the way we also prove a *classical* oracle separation $\mathsf{QXC} \not\subseteq \mathsf{QMA}$, in fact, $\mathsf{SBP} \not\subseteq \mathsf{P}^{\mathsf{QMA}}$ (Corollary 4.4), thus giving a much simpler and elementary proof of the QMA lower-bound against classical approximate counting of [AKKT20] (who showed $\mathsf{SBP} \not\subseteq \mathsf{QMA}$). Our separation is perhaps stronger as $\mathsf{P}^{\mathsf{QMA}}$ likely strictly contains QMA as it contains coQMA.

Unique witness from a well-connected ground space. Theorem 1.1 underscores the importance of gaining a deeper understanding of the structure of quantum ground spaces to extract a unique quantum witness. For this, it is crucial to identify classes of Hamiltonians for which we know how to verify a unique low-energy witness. However, the only such classes of Hamiltonian known so far are classical Hamiltonians [VV86], local Hamiltonians with ground states that can be prepared by a polynomial-sized quantum circuit [ABOBS22], and local Hamiltonians with a polynomial sized low-energy subspace [JKKS+11]. Our final result identifies a new class of Hamiltonians for which we can verify a unique low energy state. Our starting point is a simple observation already hinted at in the previous paragraphs: An algorithm that only uses subroutines treating the *entire* Hamiltonian H as a single operator cannot distinguish between states of the same energy as they are indistinguishable with respect to H. For example, quantum phase estimation, which only uses time evolutions of H, falls into this category. Thus, any quantum algorithm that aims to verify a unique low-energy state must use operators which are more fine-grained than H itself, namely the *individual* local terms in H.

The action of local operators on Hamiltonian eigenstates is in itself a challenging problem, with few general results known [AKL16]. Despite this, there are families of local Hamiltonians for which useful statements can be made. A particularly interesting such family is the set of local Hamiltonians that satisfy the Eigenstate Thermalization Hypothesis (ETH) from quantum statistical mechanics [Sre99]. This hypothesis roughly says that within an energy window, eigenstates are well-connected in the sense that local operators induce transitions from one eigenstate to many others. We adopt the following informal definition from [CB23].

> *ETH (informal).* There is a collection of polynomially many local observables $A_1, \ldots A_m$ such that for every eigenstate $|\psi\rangle$ with energy roughly E, the vector $A_i |\psi\rangle$ roughly stays within a small energy window around E, and has a fairly good overlap with all the eigenstates in this window.

See Section 5.1 for the formal definition of ETH (Hypothesis 5.3) and a discussion on its physical origins. From the computer science perspective, a well-connected low-energy eigenspace should have features of an expanding graph and hence allow one to single out a unique 'fixed point'. We solidify this intuition by generalizing the quantum expander test for verifying the maximally entangled state [AHLN+14] and show the following theorem:

Theorem 1.3 (Informal restatement of Theorem 5.5)**.** *There is a* UniqueQMA *protocol for estimating the ground energy of a local Hamiltonian satisfying ETH.*

The protocol verifies a unique quantum state in two copies of the subspace with energy around E, if ETH holds in that subspace. Assuming E is the ground energy itself, we can extract a unique ground state witness. If E is larger but sub-extensive (that is, n^c with $c < 1$, as a function of the number of qubits n), we can still obtain a unique witness for the purposes of UniqueQMA protocols due to known reductions concerning the local-Hamiltonian problem.[3]

A likely implication of Theorem 1.3 is that the complexity of local Hamiltonians satisfying ETH, such as chaotic quantum Hamiltonians including the SYK model [SV17], is significantly lower than that of general local Hamiltonians. It is unclear if the theorem would help establish that UniqueQMA and QMA are equivalent under quantum reductions, since we do not know if QMA-hard local Hamiltonians would satisfy the ETH assumption.

1.1 Techniques

The oracles. Our oracles will be of the form $\mathcal{O}_{\mathcal{S}} = \mathbb{I} - 2\Pi_{\mathcal{S}}$, corresponding to a reflection about an unknown subspace $\mathcal{S}$. Variations of this oracle have been used in other works. The case when $\mathcal{S}$ corresponds to a 1-dimensional subspace is exactly the oracle used by [AK07] to separate QCMA from QMA. When increasing the dimension and restricting the subspace $\mathcal{S}$ to be classical, this becomes the classical oracle used to separate QMA and classical approximate counting [AKKT20]; when the subspace is quantum, this separates QMA and quantum approximate counting [SY23]. Having said that, our results and proof techniques do not follow from these works in any straightforward way. For example, the [AK07] oracle would not work in our setting as this problem is easily seen to be in UniqueQMA. A better candidate is the oracle $\mathcal{O}_S$ corresponding to a large subspace S, as there does not seem to be a natural "unique" witness that the prover could provide. However, when constructing our query lower bounds (which are, as usual, the heart of nearly all oracle separations), new techniques are necessary to account for the uniqueness restriction on the witness. Indeed, prior oracle separations involving QMA avoid directly dealing with the witness either by conditioning on a particular classical witness [AK07], replacing the witness with the maximally mixed state [AKKT20; SY23], or reducing to a classical circuit lower bound [AIK22; ABD24]. Instead, our work carefully uses the requirement that the witness be unique.

Separating QMA from $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ We consider circuits with access to the unknown subspace reflection oracle $\mathcal{O}_S$, as well as an oracle $\mathcal{A}$ for UniqueQCircuitSat (more precisely, their "oracularized" versions), the natural complete problems for UniqueQMA. The oracle distinguishing task (simplified) is the following.

Task 1.4 (Non-empty vs empty)**.** *Given an oracle $\mathcal{O}_S$ encoding a reflection about an unknown subspace S, distinguish between the following cases:*

- *(YES case) S is a non-trivial subspace so that* $\dim(S) \geq 1$*, and*
- *(NO case) S is an empty subspace.*

The overall idea for both separations is to use a hybrid argument and separately consider steps of the $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ verifier where a query is made to the circuit sat oracle $\mathcal{A}$ and steps where the verifier directly queries the oracle $\mathcal{O}_S$. Therefore, our separation requires two ingredients.

1. First, we argue that the verifier's calls to $\mathcal{A}$ are either on circuits $\mathtt{C}$ which are themselves good verifiers for the oracle distinguishing task *or* are nearly useless. We show this by establishing a *polarization* property of oracular quantum circuits.

[3] If ETH only holds for extensive energies (that is, for energies $\geq cn$ for some small constant c,) then Theorem 1.3 does not apply; however it gives some implications assuming quantum PCP conjecture, see Section 5.3.

2. Second, we establish our oracle separation between UniqueQMA and QMA. This allows us to rule out the first case in the item above, as well as to argue that the verifier's direct calls to $\mathcal{O}_S$ are not too helpful.

We next give some details on the above items.

In a sense Item 2 is where the uniqueness property is crucially used. Imagine we have two classical subspaces S and T, where T is a subspace *extending* S, so that $T = S \oplus \Delta$. We will pick the dimension of S and T to be on the order superpoly(n); we will show that as $\dim(S), \dim(T) \ll 2^n$, a UniqueQMA verifier using poly(n) oracle queries cannot distinguish S and *random* extensions T. When S is taken from as a YES instance of an oracle distinguishing task with a unique witness ϕ_S, this observation essentially says that the "same" witness works for T. Flipping the order of quantification, we can equivalently say that for a random T, *nearly all* of its "roots" S have the same witness. But since $\dim(T) \gg \dim(S)$, the S's can be picked so that they are orthogonal. Finally, applying a variation of the hybrid method shows that the same witness must overlap too many orthogonal subspaces, yielding a contradiction. As mentioned previously, the oracle here can even be restricted to be classical (see Corollary 3.3).

However, when moving to Item 1, the use of quantum oracle is seemingly necessary. When considering the $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ verifier's direct calls to $\mathcal{O}_S$, the correlation between the verifier's state and the oracle is naturally captured by the overlap of some intermediate state and the projector onto the oracle subspace; we can get a lot of mileage by studying quantities of the form "$\mathrm{tr}[\Pi_S \psi]$" and averaging over subspaces S. The calls to the circuit sat oracles $\mathcal{A}$ require more careful handling. The challenge is, for each oracular circuit $\mathsf{C}^{\mathcal{O}_S}$ on which the outer verifier is called, we would like to show that the answer $\mathcal{A}(\mathsf{C}^{\mathcal{O}_S})$ cannot help tell apart the yes and no case oracles (with high probability). We establish this via a concentration argument, enabled by our use of quantum oracles. In our case, we use a version of Levy's lemma for the Grassmannian (the space of k-dimensional subspaces of $\mathbb{C}^N$), due to [GS23].

Lemma 1.5 (Levy's lemma on the Grassmannian (informal))**.** *Let Δ be uniform over $G_{N,k}$ and Π_Δ be corresponding projector. Given any quantum state $|\psi\rangle \in \mathbb{C}^{N'}$, with $N' \geq N$, we have that*

$$\Pr_{\Delta \sim G_{N,k}} [\mathrm{tr}[\Pi_\Delta |\psi\rangle\langle\psi|] \geq 2N^{-1/3}] \leq \exp\left(-\mathsf{poly}(N)\right) .$$

Via Lemma 1.5 we show that the acceptance probabilities of oracular quantum polynomial time circuits must *concentrate* when averaged over the random quantum subspace oracles. A similar observation is made in [INNRY21] to rule out the possibility of a search-to-decision reduction for QMA. We use this observation somewhat differently. Whereas [INNRY21] used concentration to argue that on across *yes* instances the circuit behaves similarly, we instead want to argue the same thing across yes *and no* instances. This requires a different and more intricate argument considering the average acceptance probabilities in the yes and no cases individually. In the specific form used in this work, we call this property "polarization" (see Lemma 3.12).

Lemma 1.6 (Circuit polarization (informal))**.** *Let $\mathsf{C}^{\mathcal{O}}$ be any $\mathsf{poly}(n)$-query* UniqueQMA *verifier with access to a subspace oracle $\mathcal{O}$. Let $\gamma \in \exp(-\mathsf{poly}(2^n))$. Then, either*

$$\Pr_{\text{random oracles}} [\mathcal{A}(\mathsf{C}^{\mathcal{O}_{\text{Yes}}}) \neq \mathcal{A}(\mathsf{C}^{\mathcal{O}_{\text{No}}})] \geq 1 - \gamma \qquad \text{or} \qquad \Pr_{\text{random oracles}} [\mathcal{A}(\mathsf{C}^{\mathcal{O}_{\text{Yes}}}) \neq \mathcal{A}(\mathsf{C}^{\mathcal{O}_{\text{No}}})] \leq \gamma .$$

The final obstacle when designing oracle separations for relativized (or "stacked") classes is the fact that the natural complete problem for UniqueQMA, quantum circuit sat, is defined as a *promise problem*. This subtlety makes defining the precise behavior of an oracular verifier non-trivial, but also gives us a surprising amount of flexibility. Indeed, this freedom in defining a quantum circuit sat oracle $\mathcal{A}$'s behavior on instances outside the promise range is what enabled [INNRY21] to demonstrate the impossibility of a generic search-to-decision reduction for QMA. In this work, we use the definition of an oracular verifier proposed by [AIK22], where we are free to "extend" the oracle to behave arbitrarily on instances outside the promise gap, see Section 2.

Separating QXC from $\mathsf{QMA}^{\mathsf{QMA}}$ The overall argument here is similar to above, except we consider a slightly different task.

Task 1.7 (Small vs large)**.** *Fix constants $0 < k_1 < k_2 < 1$. Given an oracle $\mathcal{O}^S$ encoding a reflection about an unknown subspace S, distinguish between the following cases:*

- *(YES case)* $\dim(S) \leq 2^{k_1 n}$*, and*
- *(NO case)* $\dim(S) \geq 2^{k_2 n}$*.*

Intuitively, the different parameters can be attributed to the fact that in UniqueQMA, the unique witness condition affords structure even within the YES case instances (thus, both subspaces of dimension $\approx 2^{k_1 n}$ and $\approx 2^{k_2 n}$ can be treated as YES cases). For QMA, we need to treat the dimension $\geq 2^{k_2 n}$ as a NO case in order to carry through similar arguments. As in the case of UniqueQMA, we will use an oracle separation between QMA and QXC as well a version of the polarization lemma for QMA verifiers. In both separations, the polarization property puts strong limitations on the way the outer verifiers can interact with the circuit sat oracles. As such, we hope these techniques will generalize beyond our setting to yield lower bounds against verifiers querying quantum oracles, such as those capturing constant stacks of QMA (the 'QMA hierarchy').

Finally, we make a remark on our (classical) oracle separation of QMA and QXC. One might be tempted to leverage an argument similar to [AK07] to remove the outer witness. This does not work, as the witness in our case is quantum and would require conditioning on a event with probability $\leq \exp\left(-2^{\mathsf{poly}(n)}\right)$. Instead, we use a hybridization argument similar to the UniqueQMA proof. When extending a (yes case) small subspace S to a (no cases) larger subspace $S \oplus \Delta$, the only way the QMA algorithm changes its behavior is that the witness overlaps with Δ. Thus, choosing many orthogonal extensions Δ leads to a contradiction. This recovers the QMA lower bound for classical approximate counting (with quantum queries) of [AKKT20], albeit by a much more straightforward proof without involving the machinery of Laurent polynomials used in [AKKT20]. As an easy consequence of our proof, we improve the latter result to show the same lower bound holds for $\mathsf{P}^{\mathsf{QMA}}$.

Protocol for unique verification. Finally, we switch gears and show that under an appropriate structural assumption, unique verification *is* possible for the local Hamiltonian problem. As mentioned earlier, we prove Theorem 1.3 by designing a quantum expander-like test [AHLN+14] for a given low-energy subspace that satisfies the eigenstate thermalization hypothesis. Whereas the protocol of [AHLN+14] yields a test for states which are maximally entangled over an explicitly known subspace, we use the ETH to extend their test to a subspace implicitly described by a local Hamiltonian H. The unique witness that we verify is a highly entangled state on two copies of the local Hamiltonian system, localized in a low-energy window where the ETH is assumed to hold. To accomplish "mixing," the verification protocol uses the set of local operators $A_1, \ldots, A_m$ from ETH, which will behave as approximate unitaries in the subspace.

More concretely, given a purported low-energy state, we first use quantum phase estimation to perform a projector $\{\Pi, I - \Pi\}$ on the eigensubspace corresponding the energy window satisfying ETH. We then would like to apply a set of "expanding unitaries" within this subspace. An idea is to simply use the observables A_i to generate unitaries of the form $e^{\mathrm{i}A_i}$ on each copy of the system. However, doing so will uncontrollably bring the system outside of the relevant subspace. We take ideas from the quantum Zeno effect and instead apply unitaries of the form $e^{\mathrm{i}\varepsilon A_i} \otimes e^{-\mathrm{i}\varepsilon A_i}$ for sufficiently small $\varepsilon > 0$. Finally, we apply the projector $\{\Pi, I - \Pi\}$ again. To analyze this protocol, we formulate a computational interpretation of ETH in Hypothesis 5.3 that is inspired by the pseudorandomness literature, which could be of independent interests. We then use the computational ETH to show that the resulting approximate unitaries in the energy window form an approximate quantum expander. One difference from standard quantum expanders is that the unique state we can verify is not necessarily the maximally entangled state; all we know is that it is a gapped stationary state guaranteed by the Perron-Frobenius theorem. We point out here that we also need a mild assumption on the Hamiltonian density of states (see Property 1). Additionally, we highlight that [CB23] also used quantum expander behaviors and other assumptions in a quantum Gibbs sampler, but our protocol is simpler since its goal is low-energy state verification instead of Gibbs state preparation.

It may appear surprising that a (variant of) the EPR testing protocol appears in the above computational protocol, while it was originally introduced in the communication setting [AHLN+14]. This is an example of the communication vs computation correspondence: creating the state to be tested is computationally hard while testing it can be easy, in analogy with the fact that creating a maximally entangled state requires large communication but testing it does not.

1.2 Physical relevance

On the surface, our first result can be viewed as evidence that UniqueQMA is not equal to QMA. An alternative interpretation is via a characterization due to [ABOBS22]; in that paper, the authors show that the class of local Hamiltonians with an inverse polynomial spectral gap is UniqueQMA complete. Thus, our oracle separations give evidence that the local Hamiltonian problem is easier for inverse polynomially gapped instances[4]. This is interesting from a physics perspective as many classes of physically relevant Hamiltonians are known to have an inverse polynomial spectral gap (e.g., ferromagnetic Heisenberg [KN97], 1D AKLT [AKLT04], Movassagh-Shor [MS14]). In particular, the presence of a spectral gap in these models immediately makes studying various properties more tractable [Orú14; LVV15]. Thus our work provides some evidence for the physical intuition that inverse polynomially gapped Hamiltonians are "easy" by separating such Hamiltonians from QMA, under an oracle.

Our third result has implications for the Eigenstate Thermalization Hypothesis. If UniqueQMA and QMA are not equivalent (under quantum reductions), then hard instances of the local Hamiltonian problem cannot satisfy ETH in its low energy window (near-extensive). This violation of ETH is robust against adversarial perturbations of near-extensive strength. Under the quantum PCP conjecture [AAV13], which states that there are local Hamiltonians such that ground energy estimation is QMA-hard even up to extensively scaling accuracy, the latter statement can be extended to extensive energy scales. In other words, if UniqueQMA and QMA are not equivalent, then quantum PCP Hamiltonians need to robustly violate ETH. Note that the same assumption implies QMA-hard Hamiltonians are not many-body localized [NH15] either (see Section 5.3). Thus, our results suggest that these Hamiltonians may correspond to an interesting phase of matter that is robust against adversarial perturbations.

We expect our unique verification protocol to work for families of random quantum Hamiltonians such as random spin systems [ES14] or the SYK model [SY93; Kit15a; Kit15b; HO22]. If random quantum Hamiltonians satisfy a sufficiently strong variant of ETH, our result implies that the average-case is considerably easier than the worst-case, under the assumption UniqueQMA and QMA are not equivalent. Indeed, randomized versions of NP-hard problems such as optimizing the Sherrington-Kirkpatrick model [SK75] have been proven to be average-case easy [Mon21].

1.3 Open problems

There are multiple avenues to continue this work. Here, we list a couple:

- Mulmuley, Vazirani, and Vazirani [MVV87] introduced a variant of the isolation lemma in [VV86], by introducing random weights to the edges of a graph to isolate a unique maximally weighted clique. The quantum analogue of this approach is to consider the quantum variant of the Clique problem, which is QMA hard [BS08]. It seems to suffer with issues analogous to those raised in [ABOBS22]: introducing poly(n) amount of randomness does not suffice to single out a unique quantum witness. Can the ideas in [MVV87] still be useful in some quantum problems?

- Theorem 1.2 demonstrates a quantum oracle to which $\mathsf{QXC}^{\mathcal{O}} \not\subseteq \mathsf{QMA}^{\mathsf{QMA}^{\mathcal{O}}}$. We expect our tools can be extended to showing separations for constants stacks of QMA (i.e. for the QMA hierarchy [AIK22])? Extending our argument to, e.g., $\mathsf{QMA}^{\mathsf{QMA}^{\mathsf{QMA}}}$ would require 1) showing polarization for $\mathsf{QMA}^{\mathsf{QMA}}$ circuits, and 2) a two-sided query lower bound, along the lines of [SY23], but for relativized classes.

[4] A more recent work [DGF22] also studied the role of the spectral gap in small promise gap regimes.

- Theorem 1.3 constructs a quantum algorithm that verifies some low energy state under ETH. Can we verify a more natural low energy state assuming ETH, such as the purification of Gibbs state (called the thermofield double state)? Note that, due to the Feynman-Kitaev circuit-to-Hamiltonian mapping [fSV02], uniquely verifying a state is equivalent to the state being a near ground state of a local Hamiltonian with inverse polynomial spectral gap. Such local Hamiltonians for the thermofield double state have been constructed under physics-motivated assumptions in [CFHL19].

- Our results hint at a separation between worst-case and average-case local Hamiltonians. Random spin Hamiltonians [ES14] or the SYK model [SY93; Kit15a; Kit15b] are expected to satisfy the ETH and are therefore solvable by a UniqueQMA machine. It is interesting to find more evidence for the easiness of average-case Hamiltonians. Recent work studies the complexity of approximating the maximal energy of the SYK model. Ref. [HO22] provides a quantum algorithm to prepare a state with energy $c\sqrt{n}$ – a constant fraction of the expected optimum. Ref. [ACKK24] finds further evidence that estimating the energy of the SYK might even be quantumly easy. More concretely, they show that the annealed and quenched free energies agree at inverse polynomial temperatures. This is in contrast to classical random spin systems where these quantities disagree when the system is in a "glassy" phase that is algorithmically hard.

- It is interesting to find more connections between quantum many-body physics assumptions such as ETH and quantum computer science. The works [CB23; SM23] show that ETH implies fast preparation of Gibbs quantum states, if there is no bottleneck throughout the eigenvalue distribution and Gibbs distribution. However, the latter assumption cannot be justified on general grounds. On the other hand, our results work perfectly fine without this assumption.

1.4 Organization of paper

In Section 2, we define basic notions such as complexity classes, local Hamiltonians, quantum oracles, norms, and introduce and prove some useful concentration lemmas. In Section 3, we give background on the UniqueQMA vs QMA questions and its physical relevance before proving the oracle separation in Theorem 1.1. We prove classical and quantum oracle separations regarding quantum approximate counting in Section 4. In Section 5, we overview the physical motivations and introduce the eigenstate thermalization hypothesis. Then we formulate a computational variant of ETH suitable for applications to complexity theory and describe a unique verification algorithm based on this variant. The Appendix sections include random matrix theory calculations relevant for the verification algorithm as well as a proof of norm properties.

1.5 Acknowledgments

In an earlier version of this paper, the oracle separation between UniqueQMA and QMA was given via a quantum oracle. We thank Chinmay Nirkhe and Anand Natarajan for pointing out that our proof did not in fact need quantumness in this case. This version has been updated to provide a classical oracle separation.

We thank Soonwon Choi, Sam Garratt, Yingfei Gu, Rahul Jain and Shivaji Sondhi for insightful discussions, and especially thank Soonwon Choi for sharing an example which shows that the computational ETH assumption is incorrect if the Gaussian prescription is applied at high energies (we only applying it at low energies). This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing, supported by DOE QSA grant number FP00010905. AA and QTN acknowledge support through the NSF Award No. 2238836. AA acknowledges support through the NSF award QCIS-FF: Quantum Computing & Information Science Faculty Fellow at Harvard University (NSF 2013303). QTN acknowledges support through the Harvard Quantum Initiative PhD fellowship. JH acknowledges support through the Harvard Quantum Initiative postdoctoral fellowship. YH is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2140743. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors(s) and do not necessarily reflect the views of the National Science Foundation.

2 Preliminaries

Notation A register R is a named finite-dimensional complex Hilbert space. If not otherwise specified, $N \triangleq 2^n$, with n defined as appropriate in context. In this work we work with both *classical* subspaces of the form $S \subseteq \{0,1\}^n$ as well as *quantum* subspaces $\mathcal{S} \subseteq \mathbb{C}^N$. To make this distinction more clear, we use standard letters for classical subspaces S, and $\mathcal{CALLIGRAPHIC}$ letters for quantum subspace $\mathcal{S}$. Additionally, the quantity $\exp(f(n))$ is understood to be the class of functions $2^{\Omega(f(n))}$ and $\mathsf{poly}(n) \triangleq \bigcup_{k\in\mathbb{N}} \mathcal{O}(n^k)$.

2.1 Complexity classes

Definition 2.1 (QMA)**.** A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in the complexity class *Quantum Merlin-Arthur* (or QMA) if there exists a polynomial-time classical algorithm yielding a description of a quantum circuit V which implements a quantum operator $U_x : \mathcal{H} \to \mathcal{H}$ with $\mathcal{H} = \mathcal{H}_{\mathsf{R}_{\text{in}}} \otimes \mathcal{H}_{\mathsf{R}_{\text{aux}}}$, such that

1. (Completeness) For all $x \in L_{\text{yes}}$, there exists a *witness* $|\phi\rangle \in \mathcal{H}_{\mathsf{R}_{\text{in}}}$ such that

$$\|\Pi_{\text{acc}} U_x |\phi\rangle_{\mathsf{R}_{\text{in}}} |0\rangle_{\mathsf{R}_{\text{aux}}} \|_2 \geq \frac{2}{3}$$

2. (Soundness) For all $x \in L_{\text{no}}$ and $|\phi\rangle \in \mathcal{H}_{\mathsf{R}_{\text{in}}}$,

$$\|\Pi_{\text{acc}} U_x |\phi\rangle_{\mathsf{R}_{\text{in}}} |0\rangle_{\mathsf{R}_{\text{aux}}} \|_2 \leq \frac{1}{3}$$

where Π_{acc} can be taken to be the projector into the $|0\rangle$ on the first qubit.

Definition 2.2 (Unique QMA)**.** A promise problem $L = (L_{\text{yes}}, L_{\text{no}})$ is in the complexity class *Unique Quantum Merlin-Arthur* (or $\mathsf{UniqueQMA}$) if in addition to the above properties of QMA, we have the further restriction that in the *completeness* case, there is some accepting state $|\phi\rangle$ and for all states orthogonal to $|\phi\rangle$, the verifier accepts with probability equal to the soundness case. In particular,

1. (Completeness) For all $x \in L_{\text{yes}}$, there exists a *unique witness* $|\phi\rangle \in \mathcal{H}_{\text{witness}}$ such that

$$\|\Pi_{\text{acc}} U_x |\phi\rangle_{\mathsf{R}_{\text{in}}} |0\rangle_{\mathsf{R}_{\text{aux}}} \|_2 \geq \frac{2}{3}$$

   and for all states $|\psi\rangle \perp |\phi\rangle$,

$$\|\Pi_{\text{acc}} U_x |\psi\rangle_{\mathsf{R}_{\text{in}}} |0\rangle_{\mathsf{R}_{\text{aux}}} \|_2 \leq \frac{1}{3}$$

Going forward, we rarely explicitly refer to Π_{acc} and U_x. Instead, for both the above classes, we use the notation V_x to refer to both the verification circuit, as well as the POVM measurement corresponding to the accept outcome. Thus, the completeness condition could be rewritten as the existence of $|\phi\rangle$ such that $\|V_x |\phi\rangle |0\rangle \|_2 \geq \frac{2}{3}$. Note we omit the register subscripts if clear from context, and write $|0\rangle$ to refer to the $\log\dim(\mathcal{H}_{\mathsf{R}_{\text{aux}}})$-qubit all-zeros state. Additionally, define the function $\text{acc}(V_x, \psi) \triangleq \|V_x |\psi\rangle |0\rangle \|_2$ and $\text{acc}(V_x) \triangleq \max_{\text{witnesses } |\psi\rangle} \text{acc}(V_x, \psi)$. If V_x takes in no witness, then $\text{acc}(V_x)$ is simply the acceptance probability of V_x on input x.

A natural complete problem for QMA is satisfiability of a polynomial-size quantum circuit.[5] Notationally, we denote circuits in `teletype`. We overload notation and for a circuit taking in an input state $|\psi\rangle$, we write $\text{acc}(\mathtt{C}, \psi)$ to denote the probability that $\mathtt{C}$ accepts on $|\psi\rangle$ and $\text{acc}(\mathtt{C})$ as the maximum over all inputs ψ. The natural complete problem for $\mathsf{UniqueQMA}$ is the unique circuit SAT problem.

Definition 2.3. Let $\mathtt{C}$ be a $\mathsf{poly}(n)$-size quantum circuit, taking in a n-qubit witness state. UniqueQCircuitSat(c, s) is the decision problem of deciding whether,

[5] Circuit vs. verifier, as a circuit $\mathtt{C}$ may not correspond to a language.

- (YES case) There exists a witness $|\psi\rangle$ such that $\mathrm{acc}(\mathtt{C}, \psi) \geq c$ and for all orthogonal witnesses $|\phi\rangle$, $\mathrm{acc}(\mathtt{C}, \psi) \leq s$, or
- (NO case) $\mathrm{acc}(\mathtt{C}) \leq s$.

In [JKKS+11], it was shown that this problem is indeed complete for UniqueQMA.

Theorem 2.4. UniqueQCircuitSat(c, s) *is* UniqueQMA*-complete with* $c - s \geq n^{-k}$ *for some constant* k*.*

Many of the arguments used in this paper involve showing concentration about the operator norm of a POVM. In the case of UniqueQMA, its acceptance probability depends not only on its highest eigenvalue, but it's second highest as well. Thus, we will consider the Ky Fan 2-norm of a matrix M.

Definition 2.5 (Ky Fan 2-norm, $\|M\|_{\uparrow 2}$)**.** For a matrix $M \in \mathbb{C}^{N \times N}$, the Ky Fan 2-norm is defined as

$$\|M\|_{\uparrow 2} = \sigma_1(M) + \sigma_2(M)\,,$$

where $\sigma_i(M)$ is the i-th singular value of M. Alternatively, we have the following variational definition:

$$\|M\|_{\uparrow 2} = \max_{\substack{U,V \in \mathbb{C}^{2 \times N} \\ UU^\dagger = VV^\dagger = \mathbb{I}_{2\times 2}}} \mathrm{Tr}[UMV^\dagger]\,.$$

The equivalence of these definitions can be seen by taking the singular value decomposition of M. The variational definition of the Ky Fan 2-norm is for showing the following "Cauchy-Schwarz"-like property.

Lemma 2.6 (Properties of the Ky Fan 2-norm)**.** *The Ky Fan 2-norm is a norm and therefore satisfies the triangle inequality and for any scalar* t*,* $\|t \cdot M\|_{\uparrow 2} = |t| \cdot \|M\|_{\uparrow 2}$*. Moreover, the Ky Fan 2-norm is* sub-multiplicative *and satisfies*

$$\|MN\|_{\uparrow 2} \leq \|M\|_{\uparrow 2} \cdot \|N\|_{\uparrow 2}\,.$$

Finally, QXC, introduced in [BCGW22], is the complexity class which naturally contains quantum approximate counting.

Definition 2.7 (Quantum approximate counting, QXC)**.** Let V be a m-sized QMA verification circuit which takes a n-qubit state $|\psi\rangle$ as input, with $m \in \mathsf{poly}(n)$. Then, given thresholds $0 < b < a \leq 1$ with $a - b \geq \frac{1}{\mathsf{poly}(n)}$, and an error parameter $\varepsilon \geq \frac{1}{\mathsf{poly}(n)}$, the approximate counting problem QXC is to compute an estimate D such that $(1-\varepsilon)D_a \leq D \leq (1+\varepsilon)D_b$, where D_a is defined as the dimension of the witness subspace for which V accepts with probability at least a, respectively for D_b.

2.2 Quantum oracles

Our first result is a oracle separation between QMA and UniqueQMA. There have been various types of oracles considered in previous works. Although the "gold standard" for oracle separations is arguably a classical oracle (which can be queried in superposition), quantum separations have required more complex oracles, such as the distributional oracles of [NN24]. In this work, we work with quantum oracles, as studied in [AK07]. Thus, in our setting, an oracle $\mathcal{O}$ refers to some arbitrary unitary operation. We will also allow controlled queries to $\mathcal{O}$, so that a call to $\mathcal{O}$ is in fact a call to the unitary,

$$(\mathbb{I} - \Pi) \otimes \mathbb{I} + \Pi \otimes \mathcal{O}\,,$$

where Π is a projector representing a control subspace. Given a family of oracles, we can define a "oracle promise problem,"

Definition 2.8 (oracle promise Problem)**.** A oracle promise problem is defined by two disjoint sets of oracles (unitary operators), Ω_{YES} and Ω_{NO}. We use $\Omega = (\Omega_{\mathrm{YES}}, \Omega_{\mathrm{NO}})$ to refer to the decision problem itself.

Then, any L-query algorithm deciding Ω should be able to receive an oracle $\mathcal{O}$ either in the YES or NO case and use L calls to $\mathcal{O}$ to decide which is the case. Formally, a oracular verifier for $\mathsf{UniqueQMA}^{\mathcal{O}}$ consists of alternating unitaries (consisting of local quantum gates) and calls to $\mathcal{O}$. The task is to (uniquely) determine whether the unknown oracle corresponds to a YES case or NO case instance.

Definition 2.9 ($\mathsf{UniqueQMA}$ Oracle Verifier)**.** Given a n-qubit oracle $\mathcal{O}$ and $m \in \mathsf{poly}(n)$, a m-qubit, L-query $\mathsf{UniqueQMA}$ oracular verifier $V^{\mathcal{O}}$ is defined over a n-qubit oracle register $\mathsf{R}_{\mathcal{O}}$, in a n_{aux} qubit auxiliary register R_{aux}. Additionally, there is a n_{in}-qubit input register R_{in}. Thus, $m = n_{\text{aux}} + n_{\text{in}}$. Given a n_{in}-qubit witness, the verifier then performs a sequence of unitaries on the full m-qubit subspace, interwoven with L queries to an oracle $\mathcal{O}$ on $\mathsf{R}_{\mathcal{O}}$, controlled on the remaining system. Finally, the verifier performs the measurement $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$ on the first qubit in the auxiliary register (which we label as the register R_{out}), rejecting if it obtains $|0\rangle\langle 0|$ the and accepting if obtains $|1\rangle\langle 1|$. See Figure 1. Thus, the measurement corresponding to the "accept" case can be written as

$$\begin{aligned} V &\triangleq (\mathbb{I} \otimes |1\rangle\langle 1|_{\mathsf{R}_{\text{out}}}) U_L \mathcal{O} U_{L-1} \dots U_1 \mathcal{O} U_0 (\mathbb{I} \otimes |0^{n_{\text{aux}}}\rangle\langle 0^{n_{\text{aux}}}|_{\mathsf{R}_{\text{aux}}}), \\ \Pi_{\text{accept}} &\triangleq V^\dagger V. \end{aligned} \tag{1}$$

We say that V *uniquely solves* the oracle promise problem $\Omega = (\Omega_{\text{YES}}, \Omega_{\text{NO}})$ if

- If $\mathcal{O} \in \Omega_{\text{YES}}$ then there exists a witness $|\psi\rangle_{\mathsf{R}_{\text{in}}}$ so that

$$\text{acc}(V, \psi) = \| V |\psi\rangle |0\rangle_{\mathsf{R}_{\text{aux}}} \|_2^2 = \text{tr}\left[\Pi_{\text{accept}} |\psi\rangle\langle\psi|_{\mathsf{R}_{\text{in}}} \otimes |0\rangle\langle 0|_{\mathsf{R}_{\text{aux}}}\right] \geq 1 - \varepsilon\,,$$

  and for any witness $|\psi'\rangle$ where $|\psi'\rangle$ is orthogonal to $|\psi\rangle$, $\text{acc}(V, \psi') \leq \varepsilon$.

- If $\mathcal{O} \in \Omega_{\text{NO}}$ then $\text{acc}(V) \leq \varepsilon$; for all witnesses the verifier accepts with probability at most ε.

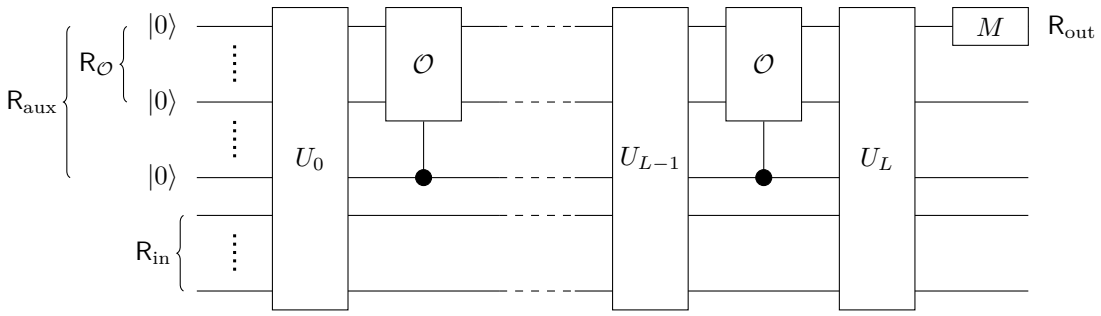

Figure 1: Circuit diagram for the verifier V. $\mathsf{R}_{\mathcal{O}}$ is the oracle register, contained in the auxiliary register R_{aux}, initialized to $|0\rangle_{\mathsf{R}_{\text{aux}}}$. R_{in} is the input register. The final measurement M is performed on R_{out}, which is the first qubit of the auxiliary register.

Remark 2.10 (Error Reduction)*.* In [JKKS+11], it was observed that the in-place amplification scheme of [MW05] works for $\mathsf{UniqueQMA}$. In particular, we may assume that $\varepsilon \in \mathcal{O}(1/\exp(n))$, at the expense of polynomial blowup in L, the number of queries to the oracle.

Remark 2.11 (Controlled Oracles)*.* Without loss of generality, we may assume that each oracle $\mathcal{O}$ is controlled on $|0\rangle\langle 0|$, on the last qubit of the auxiliary register. This is because we can write a general projector as $\Pi = U\left(\sum_{x \in [\text{rank}(\Pi)]} |x\rangle\langle x|\right) U^\dagger$. The U and $U^\dagger$ can be absorbed into the unitaries preceding and following the oracle call. The computational basis projector can be turned into one-qubit control by: compute the indicator function $\mathbf{1}_{x \in [\text{rank}(\Pi)]}$ and write output on an ancilla qubit, apply $\mathcal{O}$ controlled on the output qubit, and uncompute. As such, we can choose the control subspace as $\Pi = |0\rangle\langle 0| \otimes \mathbb{I}$.

The standard oracle that we consider in this work are variations of an oracle considered in prior works: the subspace reflection oracle $\mathcal{O}^{\mathcal{S}} = \mathbb{I} - 2\Pi_{\mathcal{S}}$, where $\Pi_{\mathcal{S}}$ corresponds to a projection onto a k-dimensional subspace $\mathcal{S}$. In [AK07], $\mathcal{O}^{\mathcal{S}}$ with $k = 1$ was used to separate QCMA from QMA; in [AKKT20], $\mathcal{S}$ was restricted to be a classical subspace to separate QMA from classical approximate counting; and finally [SY23] quantized the prior work to separate QMA from quantum approximate counting. We borrow notation from these last two works to define the oracle promise problems we study.

Definition 2.12 (Quantum approximate dimension)**.** In the oracle promise problem $\mathsf{QApxDim}(k_1, k_2, N)$, the problem is to decide whether an n-qubit unknown oracle corresponds to a reflection about a subspace of dimension $\geq k_2$ (the YES case), or a subspace of dimension $\leq k_1$ (the NO case), promised that one of these two is the case. For technical reasons, we restrict $k_2 \leq 2^{2n/3}$.

If we are further promised that the subspace is in the computational basis, then we have the classical problem $\mathsf{ApxDim}(k_1, k_2, N)$.

Remark 2.13 (Classical oracles). We use the term "classical oracle" to refer to phase oracles $\mathcal{O}^S = \mathbb{I} - 2\Pi_S$, where S is restricted to be in the computational basis. The truly "classical" analogue would be the oracle mapping $|x\rangle |b\rangle \to |x\rangle |b \oplus \mathbf{1}_{x \in S}\rangle$. However, as our separations holds under controlled oracles, these notions are equivalent.

In this work, we also consider oracle separations between relativized classes, e.g. $\mathsf{BQP}^{\mathsf{UniqueQMA}}$. We instantiate verifiers belong to a relativized class ${\mathcal{C}_1}^{\mathcal{C}_2}$ as a verifier for $\mathcal{C}_1$, with access to an black-box oracle $\mathcal{O}$ for a $\mathcal{C}_2$-complete problem. In our case, $\mathcal{C}_2$ is either UniqueQMA or QMA and thus $\mathcal{A}$ is a classical oracle and the corresponding complete problem will be UniqueQCircuitSat or QCircuitSat respectively. However, one difficulty in defining it in this way is that these languages correspond *promise* problems, but there is not guarantee that an algorithm querying $\mathcal{A}$ will do so only on valid instances. To deal with this issue, we take an approach used in prior works [AIK22; INNRY21] and imagine that the outer verifier is given access to an oracle $\mathcal{A}$ for a decision "extension" of the promise problem.

Definition 2.14 (Promise oracle extension)**.** For any oracle given by the partial function $\mathcal{O} : \{0,1\}^n \to \{0,1,\perp\}$, an oracle extension is a total function $\mathcal{A} : \{0,1\}^n \to \{0,1\}$ such that $\mathcal{O}(x) = \mathcal{A}(x)$ for all $x \in \mathrm{Dom}(\mathcal{O})$.

Thus, we take the following definition of a verifier accessing a promise oracle $\mathcal{O}$, as suggested in [AIK22]:

$$V^{\mathcal{O}}(x) \triangleq \begin{cases} 1 & \text{if } V^{\mathcal{A}}(x) = 1 \text{ for every } \mathcal{A} \text{ extending } \mathcal{O}, \\ 0 & \text{if } V^{\mathcal{A}}(x) = 0 \text{ for every } \mathcal{A} \text{ extending } \mathcal{O}, \\ \perp & \text{otherwise.} \end{cases}$$

As stated in [AIK22], it is not obvious that this is the "right" way to define oracle access to a promise problem. Nonetheless, we believe this to be a fairly reasonable definition; when querying an oracle for, e.g., QuantumCircuitSat on an instance falling outside of the promise gap, why should the verifier expect a "reasonable" response from oracle?

Lastly, we remark that whenever we have a verifier of the form ${\mathcal{C}_1}^{\mathcal{C}_2}$, we assume that the oracle for $\mathcal{C}_2$ acts on instances of size polynomial in the outer problem instance. This is required technically as many of our arguments go through union bounding over possible inputs to the oracle for $\mathcal{C}_2$. Moreover, as noted in [AKKT20], there is a distinguisher between (classical) subspaces of dimension k and $2k$ with a witness of size roughly $\mathcal{O}(k)$. Thus, without a bound on the witness register, this distinguisher would preclude a $\mathsf{P}^{\mathsf{QMA}}$ query lower bound for approximate counting.

2.3 Local Hamiltonians

A local Hamiltonian is a Hermitian matrix of the form $H = \sum_i h_i$, with bounded operator norm $\|h_i\| \leq 1$ and each h_i acts on a constant number of qubits. We will also assume that each qubit is acted upon a constant number of local terms. This definition allows the Hamiltonian to act on a high dimensional geometry in a local way. We say that a Hamiltonian is *geometrically-local* if it is defined on a low-dimensional lattice.

Definition 2.15 ((c,s)-Local Hamiltonian Problem)**.** Consider a family of Hamiltonians $\mathcal{H}$ and parameters a, b with $b - a \geq \frac{1}{\mathsf{poly}(n)}$. For any $H \in \mathcal{H}$, let $\lambda_0(H) \leq \lambda_1(H) \leq \cdots \leq \lambda_N(H)$ denote the eigenvalues of H. Then, given $H \in \mathcal{H}$, the (c,s)-local Hamiltonian problem is the problem of deciding between the two cases, promised that one of the cases holds,

1. (YES) $\lambda_0(H) \leq c$, or
2. (NO) $\lambda_0(H) \leq s$.

The local Hamiltonian problem was shown to be complete for the complexity class QMA by Kitaev [fSV02]. Much like the Local Hamiltonian problem LH is complete for QMA, we have the following natural complete problem for UniqueQMA.

Definition 2.16 ((c,s)-Unique Local Hamiltonian Problem)**.** Consider a family of local Hamiltonian $\mathcal{H}$, as above. Then, the (c,s)-*Unique Local Hamiltonian Problem* (or UniqueLH(c,s)) is the problem of deciding between the two cases, promised that one of the cases holds.

1. (YES) $\lambda_0(H) \leq c$ and $\lambda_1(H) \geq s$, or
2. (NO) $\lambda_0(H) \geq s$.

Then by a straightforward adaptation of Kitaev's proof, we have that

Theorem 2.17 (UniqueLH is Complete for UniqueQMA)**.** UniqueLH(c,s) *is complete for* UniqueQMA *with* $s - c \geq n^{-k}$ *for some constant* k*.*

2.4 Probability and concentration

Many of the arguments in this paper involve drawing random k-dimensional subspaces and showing that their intersection with a fixed state is small on average. To argue this formally, we use a version of Levy's Lemma which holds for the *Grassmannian* $G_{\mathbb{C}^N,k}$, which is the set of dimension-k subspaces of $\mathbb{C}^N$. In the literature, $G_{\mathbb{C}^N,k}$ is often used to refer to the set of projectors onto the corresponding subspaces. In particular,

$$G_{\mathbb{C}^N,k} \triangleq \{\Pi \in \mathbb{C}^{N\times N} \mid \Pi^2 = \Pi, \mathrm{rank}(\Pi) = k\}.$$

In this work, we will work with subspaces and the corresponding projectors interchangeably, using a letter like $\mathcal{S}$ to refer to the subspace, and $\Pi_{\mathcal{S}}$ to refer to the subspace projector. We will abbreviate $G_{\mathbb{C}^N,k}$ to $G_{N,k}$. The unique unitarily-invariant Haar measure over $G_{N,k}$ is denoted $\mathrm{Haar}_{N,k}$. [GS23] gives the following analogue of Levy's Lemma over the Grassmannian:

Lemma 2.18 (Levy's lemma on the Grassmannian [GS23])**.** *Let* $F : G_{N,k} \to \mathbb{R}$ *be a* κ*-Lipschitz function defined on the Grassmannian. Then,*

$$\Pr_{\mathcal{S}\sim\mathrm{Haar}_{N,k}}[|F(\mathcal{S}) - \mathop{\mathbb{E}}_{\mathcal{S}'} F(\mathcal{S}')| \geq \eta] \leq \exp\left(-\frac{(N-1)\eta^2}{16\kappa^2}\right).$$

F is κ-Lipschitz if for all $\mathcal{S}, \mathcal{S}' \in G_{N,k}$, $|F(\mathcal{S}) - F(\mathcal{S}')| \leq \kappa\|\Pi_{\mathcal{S}} - \Pi_{\mathcal{S}'}\|_2$, where $\|\cdot\|_2$ is the Frobenius norm.

Additionally, we introduce the notion of subspace extensions. Given a k_1 dimension subspace $\mathcal{S} \subseteq \mathbb{C}^N$, we define a dimension $k_2 > k_1$ extension $\mathcal{T}$ of $\mathcal{S}$ as $\mathcal{T} = \mathcal{S} \oplus \Delta$, with $\dim(\Delta) = k_2 - k_1$. Moreover, we define the Haar-random distribution over dimension k_2 extensions of $\mathcal{S}$ as $\mathcal{Q}_{\uparrow\mathcal{S},k_2}$ with,

$$\Pr_{T\sim\mathcal{Q}_{\uparrow\mathcal{S},k_2}}[\mathcal{T} = \mathcal{T}^\star] \triangleq \Pr_{\Delta\sim\mathrm{Haar}\left(G_{\mathbb{C}^N\setminus\mathcal{S},k_2-k_1}\right)}[\mathcal{S}\oplus\Delta = \mathcal{T}^\star].$$

Extensions T for classical subspaces S are defined similarly, but where $\Delta \subseteq \{0,1\}^n$. Notationally, we write this as $T \sim \mathcal{C}_{\uparrow S,k_2}$ to differentiate it from the quantum distribution.

Most often, we will use the above lemma to bound the probability that a quantum state $|\psi\rangle$ overlaps a Haar-random subspace $\mathcal{S}$ drawn from a N-dimensional ambient space.

Lemma 2.19 (Concentration of overlap)**.** *Let $\mathcal{S} \sim \mathcal{G}_{N,k}$ be a Haar-random k-dimensional subspace of $\mathbb{C}^N$ with $k \leq N^{2/3}$ and $\Pi_{\mathcal{S}}$ be the projector onto S. Given any quantum state $|\psi\rangle \in \mathbb{C}^{N'}$, with $N' \geq N$, we have that*

$$\Pr_{\mathcal{S}}[\mathrm{tr}[\Pi_{\mathcal{S}} |\psi\rangle\langle\psi|] \geq 2N^{-1/3}] \leq \exp\left(-N^{1/3}\right).$$

Proof. Define $f(\mathcal{S}) = \mathrm{tr}[\Pi_{\mathcal{S}} |\psi\rangle\langle\psi|]$. Then, $\mathbb{E}_{\mathcal{S}} f(\mathcal{S}) = \frac{k}{N} \leq N^{-1/3}$. We bound the Lipschitz constant of f as,

$$|f(\mathcal{S}) - f(\mathcal{S}')| \leq |\,\mathrm{tr}[(\Pi_{\mathcal{S}} - \Pi_{S'}) |\psi\rangle\langle\psi|]| \leq \|\Pi_{\mathcal{S}} - \Pi_{S'}\|_{\mathrm{op}} \| |\psi\rangle\langle\psi| \|_1 = \|\Pi_{\mathcal{S}} - \Pi_{S'}\|_{\mathrm{op}}.$$

Applying Lemma 2.18 with $\eta = N^{-1/3}$ and noting that $\mathbb{E}_S[\mathrm{tr}[\Pi_S |\psi\rangle\langle\psi|]] = \frac{k}{N} \leq N^{-1/3}$, we find that $\Pr_S[\mathrm{tr}[\Pi_{\mathcal{S}} |\psi\rangle\langle\psi|] \geq 2N^{-1/3}] \leq \exp\left(-N^{1/3}\right)$ as desired. □

We apply this lemma to show that quantum circuits cannot distinguish between oracles and their random extensions.

Lemma 2.20 (Quantum oracle indistinguishability)**.** *Suppose $k' \leq N^{2/3}$ and $k \in o(k')$. Let V be an L-query oracle verifier and $|\phi\rangle$ be a quantum state. Let $\mathcal{S}$ be a k-dimensional subspace and $\mathcal{T} \sim \mathcal{Q}_{\mathcal{S}\uparrow,k'}$ be a random k'-dimensional subspace extending $\mathcal{S}$. Then with probability at least $1 - \epsilon$ over the choice of $\mathcal{T}$, it holds that $\|(V^{\mathcal{O}^{\mathcal{S}}} - V^{\mathcal{O}^{\mathcal{T}}}) |\phi\rangle |0\rangle\| \in \mathcal{O}(LN^{-1/3})$, where $\epsilon \in \mathcal{O}(\exp(-N^{1/3}))$.*

Proof. Recalling from Definition 2.9, we define a *hybridized* verifier $V[\ell]$.

$$V[\ell] \triangleq \Pi_{\mathrm{out}} U_L \mathcal{O}^{\mathcal{T}} U_{L-1} \dots U_\ell \mathcal{O}^{\mathcal{S}} U_{\ell-1} \dots U_0 \Pi_{\mathrm{in}},$$

with $\Pi_{\mathrm{out}} = \mathbb{I} \otimes |1\rangle\langle 1|_{\mathsf{R}_{\mathrm{out}}}$ and $\Pi_{\mathrm{in}} = \mathbb{I} \otimes |0\rangle\langle 0|_{\mathsf{R}_{\mathrm{aux}}}$. In other words, the circuit queries the oracle $\mathcal{O}^{\mathcal{S}}$ in the first ℓ queries and then switches to the oracle $\mathcal{O}^{\mathcal{T}}$. Then on the unique witness $|\phi_{\mathcal{S}}\rangle$ of $\mathcal{O}^{\mathcal{S}}$,

$$\begin{aligned}
\|V^{\mathcal{T}}(|\phi_{\mathcal{S}}\rangle |0\rangle) - V^{\mathcal{S}}(|\phi_{\mathcal{S}}\rangle |0\rangle)\|_2 &= \|V[0](|\phi_{\mathcal{S}}\rangle |0\rangle) - V[L](|\phi_{\mathcal{S}}\rangle |0\rangle)\|_2 \\
&\leq \sum_{\ell=0}^{L-1} \|V[\ell](|\phi_{\mathcal{S}}\rangle |0\rangle) - V[\ell+1](|\phi_{\mathcal{S}}\rangle |0\rangle)\|_2 \\
&\leq \sum_{\ell=0}^{L-1} \|\Pi_{\mathrm{out}} U_L \mathcal{O}^{\mathcal{T}} U_{L-1} \dots U_{\ell+1} (\mathcal{O}^{\mathcal{T}} - \mathcal{O}^{\mathcal{S}}) U_\ell \dots U_1 \mathcal{O}^{\mathcal{S}} U_0 \Pi_{\mathrm{in}} (|\phi_{\mathcal{S}}\rangle |0\rangle)\|_2 \\
&\leq \sum_{\ell=0}^{L-1} \|(\mathcal{O}^{\mathcal{T}} - \mathcal{O}^{\mathcal{S}}) |\phi(\ell)\rangle\|_2 \\
&= \sum_{\ell=0}^{L-1} 2\|\Pi_\Delta |\phi(\ell)\rangle\|_2. \qquad (2)
\end{aligned}$$

In the final line, we use that $\mathcal{O}^{\mathcal{T}} - \mathcal{O}^{\mathcal{S}} = 2\Pi_{\mathcal{T}} - 2\Pi_{\mathcal{S}} = 2\Pi_\Delta$. Crucially, $|\phi(\ell)\rangle$ is only a function of the subspace $\mathcal{S}$ and does not depend on $\mathcal{T}$. The distribution over Δ is equivalent to $\mathrm{Haar}_{\mathbb{C}^N \setminus \mathcal{S}, k'-k}$. Thus, by Lemma 2.19,

$$\Pr_{\Delta \sim \mathrm{Haar}_{N-k,k'-k}} [\mathrm{tr}[\Pi_\Delta \phi(\ell)] \geq 2(N-k')^{-1/3}] \leq \exp\left(-(N-k)^{1/3}\right) \leq \exp\left(-N^{1/3}\right)$$

where the last inequality holds as $k \in o(N)$. Thus Equation (2) is at most $\mathcal{O}(L \cdot N^{-1/3})$ with probability at least $1 - L\exp(-N^{1/3})$, which proves the claim. □

Replacing the quantum oracle with classical ones, we get nearly the same proof, except we use Markov's inequality rather than Lemma 2.19. Additionally, we need to set $k' \leq N^{1/4}$ for technical reasons.[6] Applying Markov's inequality over the random variable $r_\ell(T) = \|V^T(|\phi(\ell)\rangle |0\rangle) - V^S(|\phi(\ell)\rangle |0\rangle)\|$ where $\mathbb{E}_T\, r_\ell(T) \leq N^{-3/4}$, we obtain the following lemma.

[6] Later, we'll union bound over all k' elements of T.

Lemma 2.21 (Clasical oracle indistinguishability)**.** *Suppose $k' \leq N^{1/4}$ and $k \in o(k')$. Let V be an L-query oracle verifier and $|\phi\rangle$ be a quantum state. Let S be a k-dimensional classical subspace and T be a random k'-dimensional classical subspace extending S. Then with probability at least $1-\epsilon'$ over the choice of $\mathcal{T}$, it holds that $\|(V^{\mathcal{O}^S} - V^{\mathcal{O}^T})\,|\phi\rangle\,|0\rangle\,\| \in \mathcal{O}(LN^{-3/8})$, where $\epsilon' \in \mathcal{O}(LN^{-3/8})$.*

3 UniqueQMA versus QMA

3.1 Background

We first give a more detailed exposition on the significance of UniqueQMA before stating our results. In a celebrated paper [VV86], Valiant and Vazirani answered the analogous question in the classical setting by giving a randomized reduction from NP to UniqueNP ($\mathsf{NP} \subseteq \mathsf{RP}^{\mathsf{UniqueNP}}$). This was a surprising result to researchers in the 1980's [AB09]. Naturally, this result has led to the UniqueQMA vs QMA question. However, previous works [ABOBS22; JKKS+11] did not succeed in giving a quantum analogue of the Valiant-Vazirani theorem for interesting reasons that seem to be inherently quantum.

Aharonov et al. [ABOBS22] were the first to attempt to quantize Valiant-Vazirani protocol. They pointed out that "the main difficulty in the QMA case is that we do not know in which basis to operate." In particular, the key step (see [AB09, Theorem 17.18]) is introducing an extra test that singles out an accepting witness of the NP circuit. This extra test is a pairwise independent hash function that, with high probability over the hash function family's randomness, uniquely maps some accepting bit string to the all-zeroes string. The work of [ABOBS22] attempted to mimic this random hash function by a random binary POVM measurement ("random" here means some 2-design, to allow for efficient implementation). However, they pointed out that the method fails even when the goal is to single out a single accepting witness out of just two valid witnesses! The reason is due to the concentration of measure in a high-dimensional Hilbert space: the overlaps of two orthogonal vectors with a random projector are both exponentially small in the number of qubits. Fundamentally, this issue arises because unlike in the classical setting, there is a freedom of basis for any high-dimensional subspace, and we do not know in which basis the accepting witnesses is specified.

Jain et al. [JKKS+11] took another approach that shows how to reduce $\mathsf{poly}(n)$ witnesses to one, yielding $\mathsf{FewQMA} \subseteq \mathsf{P}^{\mathsf{UniqueQMA}}$. The idea is that, at least classically, we can achieve this by simply asking the prover to provide a list of all $\mathrm{poly}(n)$ accepting witnesses, sorted in some canonical (e.g. lexicographical) order. In the quantum case, [JKKS+11] notice that a "canonical order" can be imitated by asking the prover to provide a state in the t-fold alternating subspace[7] for the accepting witnesses. When t is taken to be the dimension of accepting witnesses, this subspace has dimension one. Since the alternating subspace can be efficiently tested, this approach evades the need for knowing a basis for the accepting witnesses. However, their strategy requires a witnesses which scales with the number of accepting witnesses and thus restricts their algorithm to the case when there are at most $\mathsf{poly}(n)$ orthogonal accepting witnesses. In fact, their strategy is a black-box transformation and thus cannot resolve the relationship between UniqueQMA and QMA due to our oracle separation result in this section.

One can also think about UniqueQMA from a more physical point of view by the relationship between the existence of a unique witness and inverse polynomial spectral gap. More precisely, as shown in [ABOBS22, Lemma 49], the non-equivalence of UniqueQMA and QMA (under quantum reductions) implies that the local Hamiltonian problem with an inverse polynomial spectral gap promise is *strictly easier* than general local Hamiltonians (this holds even when we restrict the Hamiltonians to be 1D, due to [AGIK09]).

[7] This is the counterpart of the symmetric subspace, which has countless applications in quantum computing [Har13].

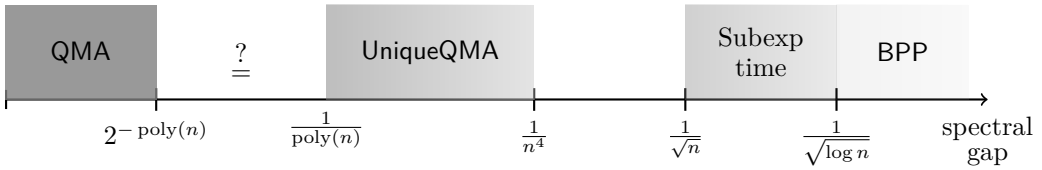

Figure 2: The computational complexity of estimating the ground energy of 1D local Hamiltonians to $1/\mathrm{poly}(n)$ accuracy depends on the spectral gap γ. The QMA-completeness when there is no spectral gap promise is shown in [AGIK09]. Ref. [ABOBS22] showed that the problem is in UniqueQMA when $\gamma = 1/\mathrm{poly}(n)$ and is UniqueQMA-hard for $\gamma = \Omega(1/n^4)$. The subexponential-time and BPP classification for sufficiently large spectral gap is due to [ALVV17], who give a classical randomized algorithm with runtime $n^{\widetilde{O}(\gamma^{-2})}$.

A recent paper [DGF22] gives weak evidence for this statement by studying the 'precise' version of the local Hamiltonian problem, where the goal is to estimate the ground energy to exponential accuracy. Without any promise on the spectral gap, this problem is known to be PSPACE-complete [FL16] (which is a class believed to be vastly more difficult than QMA). However, [DGF22] show that the complexity of the same problem drops to PP under the assumption of an inverse polynomial spectral gap. Since it is believed that PP $\subsetneq$ PSPACE, this means that the inverse polynomial spectral gap assumption makes the problem strictly easier in this high-accuracy regime. Although this doesn't necessarily indicate the same conclusion in the standard polynomial accuracy regime, it is argued in [DGF22] that a spectral gap does seem to indicate easiness as observed in practice. They conjecture that inverse polynomial gapped Hamiltonians admit low-energy states with polynomial circuit complexity, which would in turn imply containment in QCMA.

An inverse polynomial spectral gap of the parent Hamiltonian often implies provably efficient algorithms [STV12; SBE17] in the context of injective tensor networks. For general 1D Hamiltonians, a spectral gap of roughly at least $1/\sqrt{n}$ implies a sub-exponential time algorithm [ALVV17]. See Figure 2 for an incomplete computational complexity phase diagram of 1D Hamiltonians.

3.2 Results overview

In this section, we state the oracle separations involving UniqueQMA and give the proofs for classical oracles in Section 3.3 and for quantum oracles in Section 3.4. As mentioned before, the oracle problem we consider is related to the (Quantum) Approximate Dimension problem of [SY23]. We consider a slightly different oracle promise problem, QApxDimExact with parameters k_1 and k_2 so that the YES case corresponds to subspace dimension exactly k_2, and the NO case corresponds to dimension exactly k_1. Then, the specific problems we consider are,

$$\mathsf{EmptyNonEmpty} \triangleq \mathsf{ApxDimExact}(0,1,N) \cup \mathsf{ApxDimExact}(0,k,N) \text{ and}$$
$$\mathsf{QEmptyNonEmpty} \triangleq \mathsf{QApxDimExact}(0,1,N) \cup \mathsf{QApxDimExact}(0,k,N)\,.$$

where the union is taking within the NO case and YES case oracles individually (so that the NO case is still the identity oracle, and YES case corresponds to subspaces of dimension 1 and k). Since EmptyNonEmpty $\subseteq$ ApxDim$(0,1,N)$ (and QEmptyNonEmpty $\subseteq$ QApxDim$(0,1,N)$), we could've considered the harder problem instead (which would essentially recover Task 1.4), but we use EmptyNonEmpty since this includes all subspace dimensions needed to make the proof go through.

Intuitively, this definition is motivated by the desire to take advantage of UniqueQMA's structural condition on having a unique witness. Restricting to subspaces of a single dimension, say $k = 1$, there is a natural unique verification protocol. Instead, our proof will heavily exploit the interaction between subspaces of different sizes.

3.2.1 Classical oracles

We first restrict ourselves to classical (computational basis) oracles (see, Remark 2.13). We show the following query lower bound.

Theorem 3.1 (UniqueQMA query lower bound, EmptyNonEmpty)**.** *Let V be a L-query, m-qubit verifier which uniquely decides EmptyNonEmpty with error $\varepsilon \leq 2^{-n}$ both on $\mathcal{O}_{No}$ and on at least $(1-2^{-n})$-fraction of YES instances. Then, $L \in \Omega(\sqrt{k})$. In particular, if we take $k \in 2^{\Omega(n)}$, then V requires* exponentially*-many queries in n.*

Combining with standard diagonalization techniques (see Section 3.3.2), these query lower bounds yield the separation $\mathsf{UniqueQMA}^{\mathcal{O}} \subsetneq \mathsf{QMA}^{\mathcal{O}}$.

Remark 3.2 (Size of the Witness Register). The query lower bound holds independently of the witness size (in fact it holds even if the witness size is infinite). But with an arbitrary-size witness, couldn't the prover simply send an exponentially long classical description of the oracle? This doesn't work because it is not uniquely verifiable – the prover could send a different string describing a slightly different oracle. With only a polynomial number of queries, these two are indistinguishable. Applying the Valiant-Vazirani theorem to uniquely single out a bit-string will also not suffice because Valiant-Vazirani is a randomized reduction that succeeds only with probability $1/|\text{witness}|$.

We can immediately strengthen the separation in two directions. First, via a simpler version of the techniques from Theorem 3.6 (essentially, by replacing the application of Levy's Lemma with Markov's inequality, then union bounding over all polynomially-many circuits queried by a P verifier), the lower bound from Theorem 3.1 can be shown to hold for $\mathsf{P}^{\mathsf{UniqueQMA}}$ algorithms as well.[8] Second, since the oracle is classical, EmptyNonEmpty is actually solveable in NP. This yields the following improved result,

Corollary 3.3. *There exists a classical oracle $\mathcal{O}$ such that $\mathsf{UniqueQMA}^{\mathcal{O}} \subsetneq \mathsf{QMA}^{\mathcal{O}}$, and in fact, $\mathsf{P}^{\mathsf{UniqueQMA}^{\mathcal{O}}} \not\subseteq \mathsf{NP}^{\mathcal{O}}$.*

Remark 3.4. This implies that $\mathsf{NP} \not\subseteq \mathsf{P}^{\mathsf{UniqueNP}}$ relative to the same oracle, capturing the open question of *derandomizing* Valiant-Vazirani [BBF98; KVM99], where, given a satisfiable SAT formula, the goal is to deterministically output a list of $\mathsf{poly}(n)$ SAT formulae, at least one of which is uniquely satisfiable [9]. The work of [KVM99] showed that this is possible under a certain circuit complexity assumption. An influential early work [BBF98] also produced an oracle under which $\mathsf{NP} \not\subseteq \mathsf{P}^{\mathsf{UniqueNP}}$. However, our oracle separation proof is arguably more natural and direct than [BBF98]. There, the authors gave an intricate oracle relative to which $\mathsf{P} = \mathsf{UniqueNP}$ and $\mathsf{NP} = \mathsf{EXP}$, and hence $\mathsf{NP} \not\subseteq \mathsf{P} = \mathsf{P}^{\mathsf{UniqueNP}}$ due to the time hierarchy theorem and the self-lowness of P. In contrast, our classical oracle is simpler as it is essentially a random function and allows us to directly show the non-containment $\mathsf{NP} \not\subseteq \mathsf{P}^{\mathsf{UniqueNP}}$ without blowing up the power of NP or collapsing $\mathsf{UniqueNP}$ to P.

3.2.2 Quantum oracles

For our proof of Theorem 3.6, we'll need the same lower bound against the quantum oracles of QEmptyNonEmpty, where we lift the restriction that the subspace is in the computational basis. The simple QMA lower bound for QEmptyNonEmpty follows directly from Theorem 3.1. Suppose a verifier succeeded with probability $\geq 1-2^{-n}$ over quantum oracles. Treat the quantum oracle distribution as drawing a Haar-random unitary U, then conjugating a randomly chosen classical subspace. Thus, this implies that there exists some unitary $U^{\star}$ for which the verifier succeeds with probability $\geq 1-2^{-n}$, and pulling $U^{\star}$ into the verifier's unitaries yields an algorithm for EmptyNonEmpty.

[8] This query lower bound is tight up to quadratic factors; the algorithm of [JKKS+11] gives a $\mathsf{P}^{\mathsf{UniqueQMA}^{\mathcal{O}}}$ verifier solving the above task in $\mathcal{O}(k)$ queries.

[9] This is not to be confused with a much more stringent task called *deterministic isolation*, which requires one to deterministically output a single circuit that is uniquely satisfiable. There is very strong evidence against its possibility [HNOS96; DKVMW13].

Corollary 3.5 (UniqueQMA Query lower bound, QEmptyNonEmpty)**.** *Let V be a L-query, m-qubit verifier which uniquely decides QEmptyNonEmpty with error $\varepsilon \leq 2^{-n}$ both on $\mathcal{O}_{No}$ and on at least $(1-2^{-n})$-fraction of YES instances. Then, $L \in \Omega(\sqrt{k})$.*

Classically, [VV86] shows that approximate counting is contained in $\mathsf{BPP}^{\mathsf{UniqueNP}}$, and without a derandomization assumption it is not known whether approximate counting can be placed in UniqueNP. Therefore, in analogy to the classical case, a more surprising lower bound would be against $\mathsf{BQP}^{\mathsf{UniqueQMA}}$. Indeed, by moving to quantum oracles (and therefore considering the quantum approximate counting problem), we are able to obtain such a lower bound.

Theorem 3.6 (Relativized Oracle Separation)**.** *Let $V^{\mathcal{T}}$ be any L-query BQP verifier with gates implementing controlled calls to the subspace reflection oracle $\mathcal{O}^S$, as well as an oracle for $\textsc{UniqueQCircuitSat}^{\mathcal{O}^S}$. Suppose $V^{\mathcal{T}}$ solves QEmptyNonEmpty, with error at most $\varepsilon \in \exp(-n)$. Then, if $k \geq N^{\alpha}$ for some $\alpha \in \Omega(1)$, then $L \in \Omega(N/k)$.*

Standard techniques yield the desired oracle separation.

Corollary 3.7. *There exists a quantum oracle $\mathcal{O}$ such that $\mathsf{BQP}^{\mathsf{UniqueQMA}^{\mathcal{O}}} \not\subseteq \mathsf{QMA}^{\mathcal{O}}$.*

3.3 Classical oracle separations and UniqueQMA

3.3.1 Query lower bound

In this section, we prove a query lower bound for UniqueQMA algorithms against the problem

$$\mathsf{EmptyNonEmpty} = \mathsf{ApxDimExact}(0,1,N) \cup \mathsf{ApxDimExact}(0,k,N)\,,$$

establishing Theorem 3.1. All oracles in this section are taken to be *classical*. Technically, we prove a lower bound against verifiers with slightly weaker requirements, as this will lead to a simpler statement for the polarization lemma (3.12). However, the lemma directly implies the desired result.

Lemma 3.8. *Let V be a L-query, m-qubit verifier which uniquely decides $\mathsf{ApxDim}(0,k,N)$ and accepts* possibly non-uniquely *for instances corresponding to the YES case of $\mathsf{ApxDim}(0,1,N)$, both with error at most $\varepsilon \in \exp(-n)$ on $(1-2^{-n})$-fraction of instances. Then, $L \in \Omega(\sqrt{k})$.*

For comparison, recall,

Theorem 3.1 (UniqueQMA query lower bound, EmptyNonEmpty)**.** *Let V be a L-query, m-qubit verifier which uniquely decides EmptyNonEmpty with error $\varepsilon \leq 2^{-n}$ both on $\mathcal{O}_{No}$ and on at least $(1-2^{-n})$-fraction of YES instances. Then, $L \in \Omega(\sqrt{k})$. In particular, if we take $k \in 2^{\Omega(n)}$, then V requires* exponentially-*many queries in n.*

The only difference between the statement of Lemma 3.8 and Theorem 3.1 is that in the former the verifier is allowed to accept non-uniquely on subspaces of dimension 1. Any verifier satisfying the conditions of Theorem 3.1 is subject to the same query lower bound. It remains to prove Lemma 3.8.

Proof. We define a hybridized verifier between the NO case oracle $\mathcal{O}^{\emptyset}$ and some YES case oracle $\mathcal{O}^{T}$. Eventually, we will choose T to be a random k-dimensional extension of some 1-dimensional subspace S. For technical reasons, we take $k \in o(\sqrt{N})$.

Let $|\phi_T\rangle$ be the witness for T. By hybridizing as in the proof of Lemma 2.20, we find that

$$\|V^T(|\phi_T\rangle\,|0\rangle) - V^{\emptyset}(|\phi_T\rangle\,|0\rangle)\| \leq \sum_{\ell=0}^{L-1} 2\|\Pi_T\,|\phi(\ell)\rangle\,\|_2\,,$$

where $|\phi(\ell)\rangle = U_\ell ... U_1 \Pi_{\text{in}}(|\phi_T\rangle\,|0\rangle)$.

By assumption that V is a valid UniqueQMA verifier on instances of dimension k with completeness $1-\varepsilon$ and soundness ε, the left-hand side above is at least $1-2\varepsilon$ and thus there is a "critical index" $\ell^\star(T)$ for which

$$\|\Pi_T \left|\phi(\ell^\star(T))\right\rangle\|_2 \geq \frac{1-2\varepsilon}{2L}\,. \tag{3}$$

We write the critical index as a function of T to make the dependence on T explicit. Let $|u\rangle$ be a vector in T. We split this expression into contributions from $S = \mathrm{span}(|u\rangle)$ and the remaining subspace Δ (so that $T = \Delta \oplus S$). Recalling that S is a 1-dimensional subspace spanned by some $|u\rangle$,

$$\|\Pi_T \left|\phi(\ell^\star(T))\right\rangle\|_2^2 = \mathrm{Tr}[\Pi_T \phi(\ell^\star(T))] = \underbrace{\mathrm{Tr}[|u\rangle\langle u| \cdot \phi(\ell^\star(T))]}_{\text{(A)}} + \underbrace{\mathrm{Tr}[\Pi_\Delta \cdot \phi(\ell^\star(T))]}_{\text{(B)}}\,. \tag{4}$$

We will argue that both terms above are small for some choice of T. This will imply L has to be for Equation (3) to be consistent.

Bounding the First Term (A) We establish the following claim.

Claim 3.9. *Suppose $L \in o(\sqrt{k})$. There exists a 1-dimensional subspace $S = \{|u\rangle\}$ such that*

$$\sum_{\ell=0}^{L} \mathrm{Tr}\Big(\Pi_S U_\ell ... U_1 (|\phi_S\rangle\langle\phi_S| \otimes |0\rangle\langle 0|) U_1^\dagger ... U_\ell^\dagger\Big) < \frac{1}{\sqrt{k}}. \tag{5}$$

Proof. We prove this claim by contradiction. Consider any 1-dimensional space S. By assumption, the oracular verifier V^S accepts some state $|\phi_S\rangle$ (possibly non-uniquely!) with probability $\geq 1-\varepsilon$. Then Lemma 2.21 says that for $\gamma, \delta \in \mathcal{O}(LN^{-3/8})$,

$$\Pr_{T \sim \mathcal{C}_{\uparrow S,k}} [\mathrm{acc}(V^T) \geq 1-\varepsilon-\delta] \geq 1-\gamma\,. \tag{6}$$

By uniqueness of T's witness $|\phi_T\rangle$, we can bound the overlap

$$|\,\langle\phi_S|\phi_T\rangle\,| \geq 1-2\varepsilon-\delta\,. \tag{7}$$

Thus for a random classical subspace T and a basis, with probability $\geq 1-k\gamma$, it holds for all $i \in [k]$ that $|\,\langle\phi_{S_i}|\phi_T\rangle\,| \geq 1-2\varepsilon-\delta$.

Summing the contrapositive of Equation (5) over S_i we obtain

$$\begin{aligned}
k \cdot \frac{1}{\sqrt{k}} &\leq \sum_{i=1}^{k}\sum_{\ell=0}^{L} \mathrm{Tr}\Big(\Pi_{S_i} U_\ell ... U_1 (|\phi_{S_i}\rangle\langle\phi_{S_i}| \otimes |0\rangle\langle 0|) U_1^\dagger ... U_\ell^\dagger\Big) && (8)\\
&\leq \sum_{\ell=0}^{L} \mathrm{Tr}\left(\left(\sum_{i=1}^{k}\Pi_{S_i}\right) U_\ell ... U_1 (|\phi_T\rangle\langle\phi_T| \otimes |0\rangle\langle 0|) U_1^\dagger ... U_\ell^\dagger\right) + k\cdot(4\varepsilon+2\delta) && (9)\\
&\leq L + kL\cdot(4\varepsilon+2\delta), && (10)
\end{aligned}$$

where the last inequality uses Π_{S_i} are mutually orthogonal. By choice of parameters, $L \ll \sqrt{k}$ and $\varepsilon, \delta \ll \frac{1}{\sqrt{k}L}$. This yields a contradiction and hence Equation (5) holds for some S_i. □

The reason for taking the sum over ℓ is to remove dependence on T in the claim. Picking S according to this claim and drawing $T \sim \mathcal{Q}_{\uparrow S,k}$ yields a bound on the first term of Equation (4). Again using Equation (7), $|\,\langle\phi_{S_i}|\phi_T\rangle\,| \geq 1-2\varepsilon-\delta$ so,

$$\begin{aligned}
\mathrm{Tr}[\Pi_S \cdot \phi(\ell^\star(T))] &\leq \mathrm{Tr}\Big[\Pi_S U_{\ell^\star(T)} ... U_1 (|\phi_S\rangle\langle\phi_S| \otimes |0\rangle\langle 0|) U_1^\dagger ... U_{\ell^\star(T)}^\dagger\Big] + (4\varepsilon - 2\delta) && (11)\\
&\leq \frac{1}{\sqrt{k}} + 4\varepsilon - 2\delta. && (12)
\end{aligned}$$

Bounding the Second Term (B) Take $S \in \mathcal{S}$ as above, and let T be a random extension of S. Let $\mathcal{A}(T)$ be the event that T has the same witness as some witness of S. Taking γ, δ as above, by Lemma 2.21, $\Pr_T[\mathcal{A}(T)] \geq 1 - \gamma$. Thus, we may condition on this event with at most an additive γ loss in our final bound and treat $|\phi\rangle = |\phi(\ell)\rangle$ as independent of T. Next, Δ is a random dimension subspace of dimension at most $N^{2/3}$ from a $N-1$-dimensional ambient space (the subspace orthogonal to S). Applying Markov's inequality on $r_\Delta = \mathrm{tr}[\Pi_\Delta \phi]$, we have that $\Pr_\Delta[r_\Delta \geq N^{-3/8}] \leq N^{-3/8}$. Thus,

$$\Pr_T[\mathrm{Tr}[\Pi_\Delta \phi(\ell(T))] \geq 2^{-n/3}] \leq 1 - N^{-3/8} + \gamma \leq 1 - 2\gamma\,. \tag{13}$$

Concluding Combining Equation (12) and Equation (13), we conclude that for some $S \in \mathcal{S}$, with probability $1 - 2^{-\Omega(n)}$ over extensions T of S,

$$(4) \leq \frac{1}{\sqrt{k}} + 4\varepsilon - 2\delta + \mathcal{O}(LN^{-3/8}) \in 2^{-\Omega(n)}\,. \tag{14}$$

But recalling Equation (3), (4) is lower bounded by $\frac{1}{2L}(1 - 2\varepsilon)$. Thus, we must have that $L \in \Omega(\sqrt{k})$. □

3.3.2 Diagonalization argument

We now prove Corollary 3.3 using the standard diagonalization argument.

Proof of Corollary 3.3. Let $\mathcal{O} = \{\mathcal{O}_n\}_{n \geq 1}$ be a set of oracles, with $\mathcal{O}_n = \mathcal{O}^{S_n}$, where $S_n \subseteq \mathbb{C}^{2^n}$ and is possibly empty. Define the unary language $L^{\mathcal{O}} \subseteq \{0\}^n$ so that $n \in L^{\mathcal{O}}$ if and only if $S_n \neq \emptyset$ (and thus $\mathcal{O}^{S_n}$ is a positive instance of $\mathsf{EmptyNonEmpty}$). As in [AK07], we assume for simplicity that for a length-n input, any verifier for $L^{\mathcal{O}}$ only queries $\mathcal{O}^n$, and not $\mathcal{O}^m$ for $m \neq n$.[10] For any $\mathcal{O}$, $L^{\mathcal{O}}$ is clearly contained in $\mathsf{QMA}^{\mathcal{O}}$ (and $\mathsf{NP}^{\mathcal{O}}$!). Given input n, and a witness state $|\psi\rangle$, the verifier constructs the state $\frac{1}{\sqrt{2}}(|0\rangle|\psi\rangle + |1\rangle|\psi\rangle)$. Then, apply the oracle $\mathcal{O}_n$ on the second register, conditioned on the first. Finally, the verifier measures the first qubit in the Hadamard basis. If $\mathcal{O}^n = \mathcal{O}_{S_n}$ for some non-empty subspace S_n, the prover can provide a state $|\psi\rangle \in S_n$ such that the prover accepts (with probability 1). Otherwise, no choice of $|\psi\rangle$ will cause the verifier to accept with probability more than $\frac{1}{2}$.

Now we show that there exists an oracle set $\mathcal{O}$ such that $L^{\mathcal{O}} \notin \mathsf{UniqueQMA}^{\mathcal{O}}$. We do this by enumerating all the (countably infinite) $\mathsf{UniqueQMA}$ machines $\mathcal{M}_1, \mathcal{M}_2, \ldots$ and for each $\mathcal{M}_i$ picking some oracle $\mathcal{O}^n$ so that $\mathcal{M}_i$ fails to correctly decide the string 0^n. Here, we say that M_i "correctly decides" some O^n if when $0^n \in L^{\mathcal{O}}$, there exists a unique witness such that $\mathcal{M}_i$ accepts with probability at least $\frac{2}{3}$, and if $O^n \notin L^{\mathcal{O}}$, then $\mathcal{M}_i$ rejects with probability at least $\frac{2}{3}$ for any witness.

Start with $\mathcal{M}_1$. By Theorem 3.1 for any sufficiently large n there is a choice of oracle $\mathcal{O}^n$ such that $\mathcal{M}_1$ fails to correctly decide the string 0^n. Otherwise $\mathcal{M}_1$ would yield a $\mathsf{poly}(n)$-query algorithm for $\mathsf{QApxDim}(k_1, k_2, N)$. Therefore, we pick such an n and fix this choice of $\mathcal{O}^n$. Now for $\mathcal{M}_i$, assume we've fixed $\mathcal{O}^i$, for all $i \leq N$, for some value of N. We can pick some $n > N$ and corresponding oracle $\mathcal{O}^n$ such that $\mathcal{M}_i$ incorrectly decides O^n. Constructing $L^{\mathcal{O}}$ in this way ensures that $L^{\mathcal{O}}$ is not decidable by any $\mathsf{UniqueQMA}$ machine, and thus $L^{\mathcal{O}} \notin \mathsf{UniqueQMA}^{\mathcal{O}}$. □

3.4 Quantum oracle separations and $\mathsf{BQP}^{\mathsf{UniqueQMA}}$

We now extend the previous oracle separation to show that relative to a *random* subspace reflection oracle $\mathcal{O}^{\mathcal{S}}$, $\mathsf{QMA}^{\mathcal{O}^{\mathcal{S}}} \neq \mathsf{BQP}^{\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}}$.[11] In this section, we consider quantum oracles, indicated by the calligraphic $\mathcal{S}$, and study the task

$$\mathsf{QEmptyNonEmpty} \triangleq \mathsf{QApxDimExact}(0, 1, N) \cup \mathsf{QApxDimExact}(0, k, N)\,.$$

[10] However, the same proof would work in the more general setting.

[11] Note that as in, e.g., [INNRY21], when we consider the class $\mathsf{C}_1^{\mathsf{C}_2^{\mathcal{O}}}$, we assume that the C_1 machine has access to not only the oracle for $\mathsf{C}_2^{\mathcal{O}}$, but for $\mathcal{O}$ itself as well. This only makes our lower bound stronger. Additionally, for simplicity we assume the oracle calls are to $\mathcal{O}$, and not its controlled version, but dealing with this case is easily done.

Let $V^{\mathcal{S}}$ be a $\mathsf{BQP}^{\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}}$ verifier. In other words, $V^{\mathcal{S}}$ is a BQP machine with the extra abilities of querying the oracle $\mathcal{O}^{\mathcal{S}}$, and an oracle for solving $\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}$ problems, both at unit cost. We assume the $\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}$ solver oracle is realized by an oracle $\mathcal{A}$ for the $\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}$-complete problem $\textsc{UniqueQCircuitSat}^{\mathcal{O}^{\mathcal{S}}}$. As in Definition 2.14, we take $\mathcal{A}$ to be a oracle extension of $\textsc{UniqueQCircuitSat}$, as this is a promise problem.

In the context of our work, this means that if our goal is to prove that $V^{\mathsf{UniqueQMA}}$ cannot decide some language, it suffices to show that for an adversarially chosen extension $\mathcal{A}$, $V^{\mathcal{A}}$ fails to decide the language. We take a fairly "reasonable" (from our point of view) choice of extension that is similar to the choice in [INNRY21, Section 3].

Definition 3.10 (Decision Oracle for $\textsc{UniqueQCircuitSat}^{\mathcal{O}^{\mathcal{S}}}$)**.** Let $\mathcal{A}^{\mathcal{O}^{\mathcal{S}}}$ be the oracle extension for the promise problem $\textsc{UniqueQCircuitSat}^{\mathcal{O}^{\mathcal{S}}}$ with completeness $1-\varepsilon$ and soundness ε. Then, $\mathcal{A}^{\mathcal{O}^{\mathcal{S}}}$ is defined as follows:

- For oracular $\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}$ circuits $\mathtt{C}$ outside of the promise, outputs $\mathcal{A}^{\mathcal{O}^{\mathcal{S}}}(\mathtt{C}) = \mathcal{A}^{\emptyset}(\mathtt{C})$. In essence, we replace the oracle $\mathcal{O}^{\mathcal{S}}$ by the identity (the empty subspace reflection) on invalid instances, where the behavior of $\mathcal{A}^{\emptyset}(\mathtt{C})$ is defined below,
- $\mathcal{A}^{\emptyset}(\mathtt{C}) = 1$ if the maximum acceptance probability of $\mathtt{C}^{\emptyset}$ is $\geq 1/2$; otherwise $\mathcal{A}^{\emptyset}(\mathtt{C}) = 0$.

Thus, we realize calls to a $\mathsf{UniqueQMA}^{\mathcal{O}^{\mathcal{S}}}$ oracle via $\mathcal{A}^{\mathcal{O}^{\mathcal{S}}}$. We will oftentimes simply write $\mathcal{A}^{\mathcal{S}}$ or $\mathcal{A}^{\emptyset}$, or omit the superscripts $\mathcal{S}, \mathcal{O}^{\mathcal{S}}$ when it is clear from context what subspace reflection we are working with.

To generalize the query lower bound from the previous section, we'll argue the following:

1. $V^{\mathcal{S}}$ only calls $\mathcal{A}^{\mathcal{S}}$ on at most $C(n) = 2^{n^c}$ distinct oracular quantum circuits $\mathtt{C}_1, \ldots, \mathtt{C}_{C(n)}$ for some constant c.
2. For each YES-case subspace $\mathcal{S} \in \Omega_{\text{Yes}}$, the behavior of our $V^{\mathcal{S}}$ must be distinct from $V^{\emptyset}$ and thus we must have that $\mathcal{A}^{\mathcal{S}}(\mathtt{C}_i) \neq \mathcal{A}^{\emptyset}(\mathtt{C}_i)$ for some $i \in [C(n)]$.
3. On the other hand, if we can show *every* $\mathtt{C}_i$ fails to correctly decide $\mathsf{QApxDim}(k_1, k_2, N)$ on all but $\frac{1}{F(n)}$ fraction of subspaces, with $F(n) > C(n)$, this implies (via a union bound) that there must be a subspace $\mathcal{S}$ such that $\mathcal{A}^{\mathcal{S}}(\mathtt{C}_i) = \mathcal{A}^{\emptyset}(\mathtt{C}_i)$ for all $i \in [C(n)]$.
4. Therefore on this $\mathcal{S}$, $V^{\mathcal{S}}$ learns nothing by querying the oracle $\mathcal{A}^{\mathcal{S}}$. Finally, to show that direct calls to $\mathcal{O}^{\mathcal{S}}$ do not help, we can hybridize between $\emptyset$ and a Haar-random $\mathcal{S}$.

The primary difficulty in executing the above plan is showing Item 3. Indeed, the oracle separation from Corollary 3.3 only rules out circuits which solve $\mathsf{QApxDim}(0, k, N)$ on nearly all instances; on the otherhand, a priori $V^{\mathcal{S}}$ could query $\mathcal{A}^{\mathcal{S}}$ on circuits which fulfill the $\mathsf{UniqueQMA}$ promise with probability only half over random $\mathcal{S}$. However, we will circumvent this issue by showing that, for any fixed oracular circuit $\mathtt{C}$, its behavior on a random $\mathcal{S}$ *polarizes*. Recall that $\mathrm{acc}(\mathtt{C}^{\mathcal{S}})$ is the maximum acceptance probability of $\mathtt{C}^{\mathcal{S}}$ over witnesses. A straightforward application of Levy's Lemma shows that if $\mathtt{C}$'s *expected* acceptance probability on a random $\mathcal{S}$ is p, then

$$\Pr_{\mathcal{S}}[|\mathrm{acc}(\mathtt{C}^{\mathcal{S}}) - p| \geq 2^{-\sqrt{n}}] \leq 2^{-\exp(n)}.$$

This will allow us to argue if on larger than $2^{-\sqrt{n}}$-fraction of instances a circuit $\mathtt{C}^{\mathcal{S}}$ accepts a witness with high probability, concentration extends this behavior to nearly all subspaces $\mathcal{S}$.

There is one subtlety, which is that for $\mathsf{UniqueQMA}$, $\mathcal{A}^{\mathcal{S}}(\mathtt{C})$ not only depends on the *maximum* acceptance probability (i.e. the operator norm of the associated POVM $\Pi_{\mathtt{C}}$), but the uniqueness as well. We show that the same concentration above also holds for the Ky Fan 2-norm, defined as the sum of the top two singular values $\Pi_{\mathtt{C}}$.

3.4.1 Circuit polarization

First, we bound the Lipschitz constant of the Ky Fan 2-norm. Treating V as the accept POVM for a quantum verifier.

Lemma 3.11. *For a random k-dimensional subspace $\mathcal{S}$ and verifier $V^{\mathcal{S}}$ with oracle access to $\mathcal{O}^{\mathcal{S}}$, define $F(V^{\mathcal{S}}) = \|V^{\mathcal{S}}\|_{\uparrow 2}$. Then, F is $\mathcal{O}(L)$-Lipschitz.*

The proof is contained in Appendix B. One can also show that the trace norm also has Lipschitz constant $\mathcal{O}(L)$.

Lemma 3.12 (Polarization of UniqueQMA oracle circuits)**.** *Let V be an polynomial-time oracular verifier and $k = N^{\alpha}$, with $\alpha \leq 2/3$. Then for any function $f(n) \in \exp\left(-N^{1/4}\right)$:*

$$\Pr_{\mathcal{T}\sim \mathrm{Haar}_{N,k}}[\mathcal{A}^{\mathcal{T}}(V) \neq \mathcal{A}^{\emptyset}(V)] \leq f(n)\,.$$

Proof. We prove by contradiction. Consider an oracle circuit V such that,

$$\Pr_{\mathcal{T}\sim \mathrm{Haar}_{N,k}}[\mathcal{A}^{\mathcal{T}}(V) \neq \mathcal{A}^{\emptyset}(V)] \geq f(n) \qquad \mathop{\mathbb{E}}_{S\sim \mathrm{Haar}_{N,1}} \|V^{S}\|_{\mathrm{op}} \geq 1 - 3\varepsilon \tag{15}$$

where $\mathcal{A}$ is the UniqueQCircutSat oracle from Definition 3.10.

Case 1: $\|V^{\emptyset}\| < 1/2$ In this case $\mathcal{A}^{\emptyset}(V) = 0$. Consider the expected operator norm of $V^{\mathcal{T}}$ over random choice of $\mathcal{T} \sim \mathrm{Haar}_{N,k}$. First, suppose $\mathbb{E}_{\mathcal{T}} \|V^{\mathcal{T}}\|_{\mathrm{op}} < 1 - 2\varepsilon$, then by Levy's Lemma via Lemma 2.18, $\|V^{\mathcal{T}}\|_{\mathrm{op}} < 1 - \varepsilon$ with probability at least $1 - \gamma$ where $\gamma = \exp\left(-N^{1/3}\right) \ll 1/f(n)$. For any such $\mathcal{T}$, $\mathsf{C}^{\mathcal{T}}$ is either a NO case or falls outside of the UniqueQMA promise. Thus, by Definition 3.10, $\mathcal{A}^{\mathcal{T}}(V) = \mathcal{A}^{\emptyset}(V)$ with probability at least $1 - \gamma$, contradicting Equation (15). Thus, it holds that $\mathbb{E}_{\mathcal{T}} \|V^{\mathcal{T}}\|_{\mathrm{op}} \geq 1 - 2\varepsilon$.

Next, we consider the expected Ky-Fan 2 norm $\mathbb{E}_{\mathcal{T}} \|V^{\mathcal{T}}\|_{\uparrow 2}$. We have $\mathbb{E}_{\mathcal{T}} \|V^{\mathcal{T}}\|_{\uparrow 2} \geq \mathbb{E}_{\mathcal{T}} \|V^{\mathcal{T}}\|_{\mathrm{op}} \geq 1 - 2\varepsilon$. Suppose $\mathbb{E}_{\mathcal{T}} \|V^{\mathcal{T}}\|_{\uparrow 2} > 1 + 2\varepsilon$. Then, with probability $1 - \gamma$, $\lambda_1(V^{\mathcal{T}}) + \lambda_2(V^{\mathcal{T}}) > 1 + \varepsilon$. Using the trivial upper bound $\lambda_1(V^{\mathcal{T}}) \leq 1$,

$$\lambda_2(V^{\mathcal{T}}) > 1 + \varepsilon - \lambda_1(V^{\mathcal{T}}) \geq \varepsilon\,,$$

which again puts $\mathsf{C}^{\mathcal{T}}$ outside of the UniqueQMA promise range and thus $\mathcal{A}^{\mathcal{T}}(V) = \mathcal{A}^{\emptyset}(V)$ with probability at least $1 - \gamma$, contradicting Equation (15). Thus, we conclude that

$$1 - 2\varepsilon \leq \mathop{\mathbb{E}}_{\mathcal{T}\sim \mathrm{Haar}_{N,k}} \|V^{\mathcal{T}}\|_{\uparrow 2} \leq 1 + 2\varepsilon\,.$$

Putting this all together, we have shown that with probability at least $1 - \gamma$, $\lambda_1(V^{\mathcal{T}}) \geq 1 - 3\varepsilon$; and with probability $1 - \gamma$, we get that $\lambda_1(V^{\mathcal{T}}) + \lambda_2(V^{\mathcal{T}}) \leq 1 + 3\varepsilon$. Therefore, assuming Equation (15), then with probability at least $1 - 2\gamma$ over $T \sim \mathrm{Haar}_{N,k}$ it holds that

$$\lambda_1(V^{\mathcal{T}}) \geq 1 - 3\varepsilon \quad \text{and} \quad \lambda_2(V^{\mathcal{T}}) \leq 6\varepsilon\,.$$

But together with the assumption on $\|V^{S}\|_{\mathrm{op}}$, this means that the circuit V solves the $\mathsf{QApxDim}(0, k, N)$ problem with error $\leq 6\varepsilon$ on $1 - 2\gamma$ fraction of instances; this contradicts the lower bound established in Theorem 3.1. Thus, we conclude that one of the inequalities in Equation (15) does not hold.

It's clear that the converse of the first item of Equation (15) yields the desired conclusion. For the second item, if the average is below $1 - 3\varepsilon$, then by Levy's Lemma it concentrates below $1 - 2\varepsilon$. Lemma 2.20 implies that the operator norm over $\mathcal{T}$ also concentrates below $1 - \varepsilon$; but then $\mathcal{T}$ is an invalid instance and $\mathcal{A}^{\mathcal{T}}(V) = \mathcal{A}^{\emptyset}(V)$ as desired.

Case 2: $\|V^\emptyset\| \geq 1/2$ In this case $\mathcal{A}^\emptyset(V) = 1$. Again consider the expected operator norm of $V^\mathcal{T}$ over random choice of $\mathcal{T} \sim \text{Haar}_{N,k}$. First, if $2\varepsilon < \mathbb{E}_\mathcal{T} \|V^\mathcal{T}\|_{\text{op}} < 1 - 2\varepsilon$ then at least $1 - \gamma$ fraction of $\mathcal{T}$ instances are invalid; for these $\mathcal{A}^\mathcal{T}(V) = \mathcal{A}^\emptyset(V)$, yielding a contradiction. If $\mathbb{E}_\mathcal{T} \|V^\mathcal{T}\|_{\text{op}} \geq 1 - 2\varepsilon$, then for $1 - \gamma$ fraction of $\mathcal{T}$'s, $\lambda_1(V^\mathcal{T}) \geq 1 - 3\varepsilon$. Consider any $\mathcal{T}$ such that this condition holds. If either $1 - 3\varepsilon \leq \lambda_1 < 1 - \varepsilon$ or $\lambda_2(V^\mathcal{T}) > \varepsilon$ then this is an invalid instance and $\mathcal{A}^\mathcal{T}(V) = \mathcal{A}^\emptyset(V)$. Otherwise, it is a valid instance and again $\mathcal{A}^\mathcal{T}(V) = \mathcal{A}^\emptyset(V)$.

Thus, $\mathbb{E}_\mathcal{S} \|V^\mathcal{S}\|_{\text{op}} \leq 2\varepsilon$. Then we immediately have that with probability $1 - \gamma$, $\lambda_1(V^\mathcal{S}) \leq 3\varepsilon$. But applying Lemma 2.20 with $k = 0$ and $k' = 1$ contradicts this statement. □

3.4.2 Putting it all together

Finally, we prove the main theorem which shows that an oracular $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ verifier cannot solve the oracle problem QEmptyNonEmpty

Theorem 3.6 (Relativized Oracle Separation)**.** *Let $V^\mathcal{T}$ be any L-query* BQP *verifier with gates implementing controlled calls to the subspace reflection oracle $\mathcal{O}^S$, as well as an oracle for* $\textsc{UniqueQCircuitSat}^{\mathcal{O}^S}$*. Suppose $V^\mathcal{T}$ solves* QEmptyNonEmpty*, with error at most $\varepsilon \in \exp(-n)$. Then, if $k \geq N^\alpha$ for some $\alpha \in \Omega(1)$, then $L \in \Omega(N/k)$.*

Proof. Consider any $\mathsf{BQP}^{\mathsf{UniqueQMA}}$ oracular verifier V and suppose $k \in \omega(\mathsf{poly}(n))$. We'll show that for $\mathcal{T} \sim \text{Haar}_{N,k}$, V fails to distinguish $\mathcal{T}$ from $\emptyset$. When applied to a state $|\psi\rangle$ encoding a superposition of circuits $\{\mathsf{C}_i\}_{i \in p(n)}$, $\mathcal{A}^\mathcal{T}$ acts as,

$$\mathcal{A}^\mathcal{T} |\psi\rangle = \mathcal{A}^\mathcal{T} \sum_{i \in [p(n)]} \alpha_i |\mathsf{C}_i\rangle |b\rangle |\phi_i\rangle = \sum_{i \in [p(n)]} \alpha_i |\mathsf{C}_i\rangle \left|b \oplus \mathcal{A}^\mathcal{T}(\mathsf{C}_i)\right\rangle |\phi_i\rangle \ . \tag{16}$$

We then write verifier $V^\mathcal{T}$ as,

$$V^\mathcal{T} = U_{2L}(\Pi_{\text{control}} \otimes \mathcal{A}^\mathcal{T}) U_{2L-1}(\Pi_{\text{control}} \otimes \mathcal{O}^\mathcal{T}) U_{2(L-1)} \dots U_2(\Pi_{\text{control}} \otimes \mathcal{A}^\mathcal{T}) U_1(\Pi_{\text{control}} \otimes \mathcal{O}^\mathcal{T}) U_0 |0\rangle \ ,$$

with interwoven calls to $\mathcal{A}^\mathcal{T}$ and the oracle $\mathcal{O}^\mathcal{T}$. We can bound $\|V^\mathcal{T} |0\rangle - V^\emptyset |0\rangle\|_2$ via hybridization. Define,

$$V[\ell] = \begin{cases} U_{2L}(\Pi_{\text{control}} \otimes \mathcal{A}^\mathcal{T}) \dots U_t(\Pi_{\text{control}} \otimes \mathcal{A}_\emptyset) U_{t-1}(\Pi_{\text{control}} \otimes \mathcal{O}_\emptyset) U_{2(T-1)} \dots U_0 |0\rangle & \text{if } t \text{ even and} \\ U_{2L}(\Pi_{\text{control}} \otimes \mathcal{A}^\mathcal{T}) \dots U_{t+1}(\Pi_{\text{control}} \otimes \mathcal{A}^\mathcal{T}) U_t(\Pi_{\text{control}} \otimes \mathcal{O}_\emptyset) U_{2(L-1)} \dots U_0 |0\rangle & \text{if } t \text{ odd.} \end{cases}$$

Thus,

$$\|V^\mathcal{T} |0\rangle - V^\emptyset |0\rangle\|_2 = \|V[2L] - V[0]\|_2 \leq \sum_{\ell=0}^{2L-1} \|V[\ell+1] - V[\ell]\|_2$$

Expanding the RHS and using unitary invariance of the ℓ_2-norm yields that,

$$\|V[\ell+1] - V[\ell]\|_2 = \begin{cases} \|(\mathcal{A}^\mathcal{T} - \mathcal{A}^\emptyset) |\phi_\ell\rangle\|_2 & \text{if } \ell \text{ is even and} \\ \|(\mathcal{O}^\mathcal{T} - \mathcal{O}^\emptyset) |\phi_\ell\rangle\|_2 = 2\|\Pi_\mathcal{T} |\phi_\ell\rangle\| & \text{if } t \text{ is odd.} \end{cases}$$

Thus, we have that

$$\|V^\mathcal{T} |0\rangle - V^\emptyset |0\rangle\|_2 \leq \sum_{\ell=0}^{L-1} \|(A^\mathcal{T} - \mathcal{A}^\emptyset) |\phi_{2\ell}\rangle\|_2 + 2\|\Pi_\mathcal{T} |\phi_{2\ell+1}\rangle\|_2 \triangleq \Delta(\mathcal{T}) \,,$$

and we want to lower bound the probability over $\mathcal{T}$ that this is at most $\frac{1}{\exp(n)}$. We bound the second term by Lemma 2.19; each term is at most $\mathcal{O}(k/N)$ with probability $\geq 1 - \gamma$, where $\gamma \in \exp\left(-N^{1/3}\right)$.

For the even case, consider the set $\mathcal{C}$ of circuits with non-zero amplitude for $|\phi_0\rangle, \ldots, |\phi_{2L-1}\rangle$. Since the circuits correspond to polynomial-time quantum verifiers, each makes at most $\mathsf{poly}(n)$-queries to $\mathcal{O}^{\mathcal{T}}$; thus we may apply Lemma 3.12 and conclude that for each $\mathtt{C} \in \mathcal{C}$ that

$$\Pr_{\mathcal{T}\sim \mathrm{Haar}_{N,k}}[\mathcal{A}^{\emptyset}(\mathtt{C}) \neq \mathcal{A}^{\mathcal{T}}(\mathtt{C})] \leq f(n)\,.$$

Moreover, the size of $\mathcal{C}$ can be bounded as $C(n) = |\mathcal{C}| \leq 2L \cdot 2^n$. Let $\mathcal{E}(\mathcal{T})$ be the event that $\mathcal{A}^{\emptyset}(\mathtt{C}) = \mathcal{A}^{\mathcal{T}}(\mathtt{C})$ for every $\mathtt{C} \in \mathcal{C}$. We have,

$$\begin{aligned}
\Pr_{\mathcal{T}}[\mathcal{E}(\mathcal{T})] = \Pr_{\mathcal{T}}[\forall \mathtt{C} \in C,\ \mathcal{A}^{\emptyset}(\mathtt{C}) = \mathcal{A}^{\mathcal{T}}(\mathtt{C})] &= 1 - \Pr_{\mathcal{T}}[\bigvee_{\mathtt{C}\in C} \mathcal{A}^{\emptyset}(\mathtt{C}) \neq \mathcal{A}^{\mathcal{T}}(\mathtt{C})] \\
&\geq 1 - \sum_{\mathtt{C}\in\mathcal{C}} \Pr_{\mathcal{T}}[\mathcal{A}^{\emptyset}(\mathtt{C}) \neq \mathcal{A}^{\mathcal{T}}(\mathtt{C})] \\
&\geq 1 - C(n) f(n) \\
&\geq 1 - \exp\left(-N^{1/4}\right),
\end{aligned}$$

where the second to last inequality holds by Lemma 3.12. If $\mathcal{E}(\mathcal{T})$ holds then $(\mathcal{A}^{\mathcal{T}} - \mathcal{A}^{\emptyset})\,|\phi_\ell\rangle = 0$ for *any* t.

Combining these bounds we have that with probability at least $1 - L(n)\cdot\gamma - \exp\left(-N^{1/4}\right) \geq \exp\left(-N^{1/4}\right)$,

$$\Delta(\mathcal{T}) \leq L(n) \cdot \mathcal{O}(k/N)\,.$$

But by the assumption on error, $\Delta(S) \geq 1 - 2\varepsilon$, and thus this requires $L(n) \in \Omega(N/k)$. □

4 Quantum approximate counting versus QMA

In this section, we use the techniques from Theorem 3.1 to obtain an oracle separation between QMA and quantum approximate counting. We recall our definition,

Definition 2.7 (Quantum approximate counting, QXC)**.** Let V be a m-sized QMA verification circuit which takes a n-qubit state $|\psi\rangle$ as input, with $m \in \mathsf{poly}(n)$. Then, given thresholds $0 < b < a \leq 1$ with $a - b \geq \frac{1}{\mathsf{poly}(n)}$, and an error parameter $\varepsilon \geq \frac{1}{\mathsf{poly}(n)}$, the approximate counting problem QXC is to compute an estimate D such that $(1-\varepsilon)D_a \leq D \leq (1+\varepsilon)D_b$, where D_a is defined as the dimension of the witness subspace for which V accepts with probability at least a, respectively for D_b.

We formalize Task 1.7 as oracle promise problem $\mathsf{ApxDim}(k_1, k_2, N) = (\Omega_{\mathrm{Yes}}, \Omega_{\mathrm{No}})$ with $k_2 \leq N$ and $k_1 \leq k_2(1 - 1/\mathsf{poly}(n))$ where,

- Each $\mathcal{O} \in \Omega_{\mathrm{Yes}}$ corresponds to a reflection over a subspace $S \subseteq \mathbb{C}^N$ of dimension k_1, with $\mathcal{O} = \mathbb{I} - 2\Pi_S$.
- Each $\mathcal{O} \in \Omega_{\mathrm{No}}$ corresponds to a reflection over a subspace $T \subseteq \mathbb{C}^N$ of dimension k_2, with $\mathcal{O} = \mathbb{I} - 2\Pi_T$.

This problem is almost by definition in $\mathsf{QXC}^{\mathcal{O}}$. Define the $\mathsf{QMA}^{\mathcal{O}}$ verifier $V^{\mathcal{O}}$ which applies the oracle $\mathcal{O}$ to the witness controlled on an auxiliary qubit initialized as $|+\rangle$, then measures the auxiliary register in the Hadamard basis. The verifier accepts if it obtains the $|-\rangle$ outcome, and rejects otherwise. Now pick $b = \frac{1}{3}$ and $a = \frac{2}{3}$. For $\mathcal{O} \in \Omega_{\mathrm{YES}}$, the verifier $V^{\mathcal{O}}$ accepts any $|\psi\rangle \in S$ with probability 1, where $\dim(S) = k_1$. Thus, $D_a = D_b = k_1$. Similarly, for $\mathcal{O} \in \Omega_{\mathrm{NO}}$, the verifier accepts with probability 1 when $|\psi\rangle \in T$, where T is k_2-dimensional so $D_a = D_b = k_2$. Thus, we can solve the oracle promise problem $\mathsf{QApxDim}(k_1, k_2, N)$ by setting $\varepsilon = \frac{k_2 - k_1}{2k_2} \geq \frac{1}{\mathsf{poly}(n)}$ and querying a $\mathsf{QXC}^{\mathcal{O}}$ algorithm on $V^{\mathcal{O}}$.

4.1 QMA versus QXC

First we show that by using techniques developed in this work we can easily establish strong query lower bounds for any QMA verifier deciding $\mathsf{ApxDim}(k_1, k_2, N)$. As for UniqueQMA, when working with lower bounds against un-relativized classes, it suffices to consider classical (computational basis) oracles.

Remark 4.1. A similar lower bound where dimension k_1 corresponds to the YES case and dimension k_2 corresponds to the NO case is obtained in [SY23; AKKT20] but through much more intricate methods utilizing a version of the polynomial method for Laurent polynomials. Our proof is much simpler and gets the same quantitative bound, which is shown to be tight in [AKKT20, Theorem 5]. However, their techniques are better in that their lower bound is symmetric (it also works when the YES case is the larger subspace, and the NO case is smaller). Intuitively, this is because their proof goes via the inclusion of QMA in SBQP, the exponentially precise analogue of BQP, and thus is not sensitive to a witness.

Theorem 4.2 (QMA query lower bound for ApxDim)**.** *Let $k_1 \in o(N)$ and $k_2 = 2k_1$. Let $V^{\mathcal{O}}$ be an oracular* QMA *verifier deciding $\mathsf{ApxDim}(k_1, k_2, N)$ with error $\varepsilon \leq \exp(-n)$. Then, $L \in \Omega(\sqrt{N/k_1})$.*

Proof of Theorem 4.2. As in Theorem 3.1, the idea is to consider a pair of oracles $\mathcal{O}^S$ and $\mathcal{O}_T$, which correspond to a "Yes" and "No" case respectively. We choose $T = S \oplus \Delta$, where Δ is a $\delta = k_2 - k_1$-dimensional subspace, orthogonal to S, where we've chosen $k_2 - k_1 = k_1$. Then, we will hybridize between the verifiers corresponding to each oracle. Let V be a L-query verifier deciding $\mathsf{ApxDim}(k_1, k_2, N)$. V can be represented by an alternating sequence of unitaries and calls to an oracle $\mathcal{O}$, with a measurement performed at the end. In particular,

$$V^{\mathcal{O}} \triangleq \Pi_{\text{out}} \mathcal{O} U_{L-1} \dots U_1 \mathcal{O} \Pi_{\text{in}} \,.$$

We write the hybridized inner sequence of unitaries and oracles as,

$$\mathcal{U}_t^{(S,T)} = \mathcal{O}^T U_{L-1} \dots \mathcal{O}^T U_t \mathcal{O}^S U_{t-1} \dots U_1 \mathcal{O}^S \,,$$

where the first t queries are to S and the final $L - t$ are to T. Let $|\phi_S\rangle \triangleq |\psi_S\rangle |0\rangle$ correspond to an optimal YES case witness for $\mathcal{O}^S$. Then, we have the following completeness and soundness guarantees,

$$\text{(Completeness)} \qquad \|V^S(|\phi_S\rangle |0\rangle)\|_2 = \|\Pi_{\text{out}} \mathcal{U}_L^{(S,T)}(|\phi_S\rangle |0\rangle)\|_2 \geq 1 - \varepsilon \tag{17}$$

$$\text{(Soundness)} \qquad \|V^T(|\phi_S\rangle |0\rangle)\|_2 = \|\Pi_{\text{out}} \mathcal{U}_0^{(S,T)}(|\phi_S\rangle |0\rangle)\|_2 \leq \varepsilon \,. \tag{18}$$

This implies that there exists a critical index t such that $\|(U_{t+1}^{(S,T)} - U_\ell^{(S,T)})(|\phi_S\rangle |0\rangle)\|_2 \geq \frac{1-2\varepsilon}{L+1}$. Equivalently,

$$\begin{aligned} \frac{1-2\varepsilon}{L+1} &\leq \|(U_{\ell-1}^{(S,T)} - U_\ell^{(S,T)})(|\phi_S\rangle |0\rangle)\|_2 \\ &= \|(\mathcal{O}_\Delta^\Pi - \mathbb{I})\mathcal{O}^S U_{\ell-1} \dots (|\phi_S\rangle |0\rangle)\|_2 \\ &= 2\|(\Pi \otimes \Pi_\Delta) U_{\ell-1} \dots (|\phi_S\rangle |0\rangle)\|_2 \,, \end{aligned}$$

which implies that $|\widetilde{\phi}_\ell\rangle \triangleq U_{\ell-1} \dots (|\phi_S\rangle |0\rangle)$ has overlap at least $\frac{1-2\varepsilon}{2(L+1)}$ with $\Pi \otimes \Pi_\Delta$. Finally, we can pick $K = \frac{N-k_1}{\delta} = \frac{N-k_1}{k_1}$-many orthogonal δ-dimensional subspaces $\Delta_1, \dots, \Delta_K$. Furthermore, there are at least $K/(L+1)$-many Δ_i's which share the same critical index t. Picking k_1 so that

$$K > \frac{(L+1)^2}{1-2\varepsilon} \tag{19}$$

yields a contradiction as this implies $|\widetilde{\phi}_\ell\rangle$ has norm greater than 1. Therefore, no $\mathsf{QMA}^{\mathcal{O}}$ verifier exists unless $L \geq \sqrt{K(1-2\varepsilon)} \in \Omega(\sqrt{N/k_1})$. □

Combining the above theorem, the $\mathsf{QXC}^{\mathcal{O}}$ algorithm for $\mathsf{ApxDim}(k_1, k_2, N)$, and a standard diagonalization argument (as in the proof of Corollary 3.3), we obtain the following corollary.

Corollary 4.3. *There exists a classical oracle $\mathcal{O}$ such that* $\mathsf{QXC}^{\mathcal{O}} \not\subseteq \mathsf{QMA}^{\mathcal{O}}$.

Before moving to quantum oracles, we state one last strengthening of the above result. Over classical oracles, a variant of Lemma 3.12 holds, but we use Markov's inequality rather than Levy's Lemma and thus we take $f(n) \in \exp(-n)$. Still, this is sufficient to union bound over all at most polynomial circuits queried by a $\mathsf{P}^{\mathsf{QMA}}$ verifier. We obtain the following corollary.

Corollary 4.4. *There exists a classical oracle $\mathcal{O}$ such that* $\mathsf{QXC}^{\mathcal{O}} \not\subseteq \mathsf{P}^{\mathsf{QMA}^{\mathcal{O}}}$, *and in fact* $\mathsf{SBP} \not\subseteq \mathsf{P}^{\mathsf{QMA}}$.

4.2 $\mathsf{QMA}^{\mathsf{QMA}}$ versus QXC

We can do better than the previous section, and by working with quantum oracles, we obtain a separation from between $\mathsf{QMA}^{\mathsf{QMA}}$ and QXC. The key tool is an observation on the polarizing property of QMA circuits, as well as a strengthening of a QMA query lower bound for QApxDim due to [SY23], based on the polynomial method of [BBCMDW01]. One may wonder if these techniques immediately would yield lower bounds against $\mathsf{QMA}^{\mathsf{QMA}}$ (or even constant stacks of QMA). However, the polynomial method by itself is not sufficient to establish a lower bound against verifiers calling oracles for QMA as there is no obvious degree bound on the steps of the algorithm querying the oracle. We circumvent this issue by a similar argument used in the proof of Theorem 3.6: the acceptance probability of QMA circuits over random subspace oracles polarize and we can argue that either these oracle calls do not help *or* we can apply the lower bound of [SY23].

Before showing the polarization result, a couple technical refinements of [SY23] are necessary. First, as stated, their lower bound only holds for algorithms succeeding on all instances. We will need to extend this to work for algorithms succeeding with high probability. Second, we use our standard hybridization argument to complete the lower bound; this works by considering a fixed subspace $\mathcal{S}$ and drawing $T \sim \mathcal{Q}_{\mathcal{S}\uparrow,k}$. However $(\mathcal{S}, T)$ where $T = S \oplus \Delta$ make up a very small fraction of possible pairs, and it is plausible that a QMA verifier succeeds at distinguishing only these very small fraction of instances. Nonetheless, we show that a QMA algorithm succeeds at distinguishing extensions *only if* it succeeds on random pairs $(\mathcal{S}, T)$.

Theorem 4.5 (Probabilistic variant of Theorem 4.2 of [SY23])**.** *Let V be m-qubit* QMA *verifier taking in a $w \leq m$ qubit witness and making L queries to an n-qubit oracle encoding an unknown subspace of $\mathbb{C}^N$. Suppose that V solves the oracle promise problem in the sense that*

$$\Pr_{\substack{\mathcal{S}\sim\mathcal{H}_{N,k_1}\\ T\sim\mathcal{H}_{N,k_2}}}[|\mathrm{acc}(V^{\mathcal{S}}) - \mathrm{acc}(V^{\mathcal{T}})| \geq \varepsilon] \geq 1-\gamma\,,$$

where $\gamma \in \mathcal{O}(\exp(-N))$, ε is a constant smaller than $\frac{1}{2}$, and $k_1 \in \mathcal{O}(k_2)$. Then, $L \in \Omega(\min\{k_2, \sqrt{N/k_2}\})$.

Proof. We begin by removing the dependence of V from the witness. By first applying the in-place amplification of [MW05] to suitably amplify the error ε then replacing the witness with the maximally mixed state, we obtain a L'-query verification algorithm $\widetilde{V}$ which takes no witness and,

- with probability $1-\gamma$ over $\mathcal{S}$, $\mathrm{acc}(V^{\mathcal{S}}) \geq 2^{-w}$ and,
- with probability $1-\gamma$ over $\mathcal{T}$, $\mathrm{acc}(V^{\mathcal{T}}) \leq 2^{-10w}$.

Moreover, $L' \in \mathcal{O}(wL)$. Next, obtain a symmetrization of our query algorithm as a degree $2L'$ polynomial in the *eigenvalues* of the subspace reflection oracle. In our setting, these eigenvalues correspond to strings $(\lambda_1, \ldots, \lambda_N) \in \{-1, 1\}^N$ where $\mathrm{wt}_{=-1}(\vec{\lambda}) \triangleq |\{i \,:\, \lambda_i = -1\}|$ corresponds to the subspace dimension. To obtain this symmetrization, we only require a slight modification of the proof of Theorem 1.3 of [SY23] given in Section 3.1 therein. We sketch this modification briefly. First, given the verifier $\widetilde{V}$, it's acceptance probability can be expressed as a polynomial $p(U, U^\star)$.[12] The symmetrized polynomial q is obtained by taking the Haar average of U, $U^\star$ so that

$$q(U, U^\star) = \mathop{\mathbb{E}}_{g\sim U(d)} p(gUg^\star, gU^\star g^\star)\,. \tag{20}$$

[12]The input $U^\star$ is to allow access to the *inverse* of the oracle; this is not strictly necessary for us, but either way it only increases the degree by a constant multiplicative factor, and thus we retain it for full generality.

Finally, the authors observe that the right-hand side simplifies to a linear combination of terms of the form $\mathrm{Tr}[U^r]$ or $\mathrm{Tr}[(U^\star)^s]$ for some integers r, s. These expressions clearly only depend on the eigenvalues of U and $U^\star$.

In our approximate setting, we do not have the exact equality in Equation (20). However, these expressions still hold with an additive error γ, accounting for the fraction of $\mathcal{S}, T$ for which the verifier may behave poorly. Thus, we obtain a polynomial p for which

$$p(\lambda_1, \ldots, \lambda_N, \lambda_1^\star, \ldots, \lambda_N^\star) \text{ is } \begin{cases} \geq 2^{-w} - \gamma & \text{when } \mathrm{wt}_{=-1}(\vec{\lambda}) \geq k_2 \text{ and,} \\ \leq 2^{-10w} + \gamma & \text{when } \mathrm{wt}_{=-1}(\vec{\lambda}) \leq k_1 \,. \end{cases} \tag{21}$$

At this point, taking the polynomial p in [AKKT20, Theorem 3] to be as defined above yields the desired lower bound. □

Next, we extend the result to a fixed $\mathcal{S}^\star$ and random extensions.

Lemma 4.6 (Query lower bound with fixed $\mathcal{S}^\star$ and extensions $\mathcal{T}$)**.** *Fix a verifier V. Suppose for some constant c with probability at least $p \in \Omega(\exp(-n^c))$ over $S \sim \mathrm{Haar}_{N,k_1}$,*

$$\Pr_{\mathcal{T} \sim \mathcal{Q}_{\mathcal{S}\uparrow,k_2}} [|\mathrm{acc}(V^{\mathcal{S}}) - \mathrm{acc}(V^{\mathcal{T}})| \geq \varepsilon] \geq 1 - \gamma \,, \tag{22}$$

where $\gamma \in \mathcal{O}(\exp(-N^{1/3}))$. Then $L \in \Omega(\sqrt{N/k_1})$.

Proof. Let $\mu = \mathbb{E}_S[\mathrm{acc}(V^{\mathcal{S}})]$. As usual, we apply Lemma 2.18 to bound $\mathrm{acc}(V^{\mathcal{S}})$ within η of μ with probability at least $1 - \gamma'$, with $\eta = N^{-1/3}$ and $\gamma' \in \exp(-N^{1/3})$. Then with probability at least $p - \gamma'$ over $\mathcal{S}$, we have that both $|\mathrm{acc}(V^{\mathcal{S}}) - \mu| \leq \eta$ and Equation (22) holds so that

$$\Pr_{\mathcal{T} \sim \mathcal{Q}_{\mathcal{S}\uparrow,k_2}} [|\mathrm{acc}(V^{\mathcal{T}}) - \mu| \geq \varepsilon + \eta] \geq 1 - \gamma \,.$$

But since this holds for $1 - \gamma$ fraction of T extending $p - \gamma'$-fraction of S, the above statement also holds for $p - \gamma - \gamma'$ fraction of T, drawn from Haar_{N,k_2}, *independent of* $\mathcal{S}$. By choice of p, $p - \gamma - \gamma' \gg \gamma$, and thus $\mathbb{E}_{\mathcal{T} \sim \mathrm{Haar}_{N,k_2}}[\mathrm{acc}(V^{\mathcal{T}})] = \mu'$ with $|\mu' - \mu| \geq \varepsilon - 2\eta$. Once again using concentration of the operator norm, this implies that

$$\Pr_{\mathcal{T} \sim \mathrm{Haar}_{N,k_2}} [|\mathrm{acc}(V^{\mathcal{T}}) - \mu'| \geq \eta] \leq 1 - \gamma \,.$$

Combining these inequalities, we have that with probability $1 - \gamma$ over S and $1 - \gamma$ over T independently, $\mathrm{acc}(V^{\mathcal{T}}) \geq \mu' - \eta$ and $\mathrm{acc}(V^{\mathcal{S}}) \leq \mu + \eta$ so that

$$|\mathrm{acc}(V^{\mathcal{S}}) - \mathrm{acc}(V^{\mathcal{T}})| \geq \varepsilon - 4\eta \,.$$

Since $\varepsilon - 2\eta \in \Omega(1)$, V satisfies the same query lower bound as in Theorem 4.5. □

The important take-away from this lemma is that for verifiers making only polynomially-many queries to the oracle, there exists very few $\mathcal{S}$'s (smaller the inverse exponential) for which Equation (22) holds. Finally, we show our polarization result.

Lemma 4.7 (Polarization of QMA oracle circuits)**.** *Let C be any $n' = \mathsf{poly}(n)$-qubit* QMA *circuit which makes $L \in o(\sqrt{N/k_1})$ queries to a n qubit oracle $\mathcal{O}_{\mathcal{S}} = \mathbb{I} - 2\Pi_{\mathcal{S}}$. Let $\mathcal{S}$ be any dimension k_1 subspace such that Equation (22) fails to hold. Let $\mathcal{T}$ be a $k_2 < N^{2/3}$ dimensional extension of $\mathcal{S}$. Then for some decision-extension $\mathcal{A}$ of an oracle for $\mathsf{QCiruitSAT}(\eta, 1 - \eta)$,*

$$\Pr_{\mathcal{T} \sim \mathcal{Q}_{\mathcal{S}\uparrow,k_2}} [\mathcal{A}(\mathsf{C}^{\mathcal{S}}) \neq \mathcal{A}(\mathsf{C}^{\mathcal{T}})] \leq 2\gamma \,,$$

where $\gamma \leq \exp(-N^{1/3})$.

Proof. Let $\mu_1 = \|\mathsf{C}^{\mathcal{S}}\|_{\mathrm{op}}$. Define $\mathcal{A}$ so that on valid QMA circuits, it outputs correctly (accepting if the operator norm is at least $1 - \eta$, rejecting if at most η). Otherwise, if the input circuit falls outside of the promise gap, it returns $1_{\mu_1 > 1/2}$. Notice that if $2\eta < \mu_2 < 1 - 2\eta$, then with probability $1 - \gamma$, $\mathsf{C}^{\mathcal{T}}$ falls outside of the promise gap and thus $\mathcal{A}(\mathsf{C}^{\mathcal{T}}) = \mathcal{A}(\mathsf{C}^{\mathcal{S}})$ with probability at least $1 - 2\gamma$. We may assume that either $\mu_2 \leq 2\eta$ or $\mu_2 \geq 1 - 2\eta$. We proceed by casing on μ_1.

Case $\mu_1 \le \frac{1}{2}$. In this case $\mathcal{A}(\mathtt{C}^{\mathcal{S}})$ rejects (either as it is invalid or falls in the NO case). Now consider $\mathtt{C}^{\mathcal{T}}$. If $\mu_2 \le 2\eta$ then with probability $1-\gamma$, either the circuit $\mathtt{C}^{\mathcal{T}}$ is invalid ($\|\mathtt{C}^{\mathcal{T}}\| \in (1-\eta, 3\eta)$) or it corresponds to a reject instance. In both cases $\mathcal{A}(C^{\mathcal{T}}) = \mathcal{A}(C^{\mathcal{S}})$. Otherwise, if $\mu_2 \ge 1-2\eta$ then $\mathtt{C}$ solves $\mathsf{QApxDim}(k_1, k_2, N)$ with completeness and soundness parameters $(1-3\eta, \frac{1}{2})$ on all but a 2γ-fraction of instances. Applying the lower bound of [SY23] via Lemma 4.6 yields a contradiction.

Case $\mu_1 > \frac{1}{2}$. This case is symmetric. With probability $1-\gamma$, $\mathcal{A}(\mathtt{C}^{\mathcal{S}})$ accepts. If $\mu_2 \ge 1-2\eta$ then with probability $1-\gamma$, $\mathtt{C}^{\mathcal{T}}$ is either invalid or yields an accepting instance. In both cases, $\mathcal{A}(\mathtt{C}^{\mathcal{T}})$ accepts. Otherwise, if $\mu_2 \le 2\eta$, then $\mathtt{C}$ solves the "small-to-big" variant of $\mathsf{QApxDim}(k_1, k_2, N)$ with completeness/soundness $(\frac{1}{2}, 3\eta)$. This again yields a contradiction. □

Theorem 4.8 (Query lower bound for $\mathsf{QMA}^{\mathsf{QMA}}$)**.** *Let V be any $m = \mathsf{poly}(n)$-qubit, L-query oracular QMA verifier which purports to solve the n-qubit oracle promise problem $\mathsf{QApxDim}(k_1, k_2, N)$ with error $\varepsilon \le \exp(-n)$ and $k_1 < k_2 \le N^{2/3}$. Suppose V has gates implementing controlled calls to a n qubit subspace reflection oracle $\mathcal{O}^{\mathcal{S}}$, as well as an $n_o \le n'$ qubit oracle $\mathcal{A}$ for $\textsc{QCircuitSat}^{\mathcal{O}_S}$. Then $L \in \Omega(N^{1/3})$.*

Proof. We assume each of the L queries to $\mathcal{A}$ are interleaved with queries to $\mathcal{O}$, so that queries to $\mathcal{A}$ happen on even indices ℓ, while queries to $\mathcal{O}$ are on odd indices. Since V queries $\mathcal{A}$ on circuits encoded in at most n_o qubits, this implies that there are at most 2^{n_o} possible queried circuits. Let this set of circuits be $\mathtt{C}$. Since each encoded circuit makes $\mathsf{poly}(n) \in o(\sqrt{N/k_2})$ queries to the oracle, we have that Equation (22) of Lemma 4.6 holds for strictly less than $1/2^{n_o}$ fraction of $\mathcal{S}$'s. Thus, take some $\mathcal{S}^\star$ for which Equation (22) fails for *all* queried circuits $\mathtt{C}$. Letting $\mathcal{T} \sim \mathcal{Q}_{\mathcal{S}^\star\uparrow,k}$, we follow the same hybridization as in the proof of Theorem 3.6 which requires us to bound the quantity,

$$\|V[\ell+1] - V[\ell]\|_2 = \begin{cases} \|(\mathcal{A}^{\mathcal{T}} - \mathcal{A}^{\mathcal{S}^\star})\,|\phi_\ell\rangle\,\|_2 & \text{if } \ell \text{ is even and,} \\ \|(\mathcal{O}^{\mathcal{T}} - \mathcal{O}^{\mathcal{S}^\star})\,|\phi_\ell\rangle\,\|_2 = 2\|\Pi_\Delta\,|\phi_\ell\rangle\,\| & \text{if } \ell \text{ is odd,} \end{cases}$$

where $|\phi_\ell\rangle$ only depends on $\mathcal{S}^\star$ (in particular, $|\phi_1\rangle$ encodes the witness to the outer QMA verifier). We bound each of these over random extensions $\mathcal{T} \sim \mathcal{Q}_{\mathcal{S}^\star\uparrow,k}$.

For the first quantity, we recall from Equation (16) that $\mathcal{A}^{\mathcal{S}^\star}\,|\phi_\ell\rangle = \mathcal{A}^{\mathcal{T}}\,|\phi_\ell\rangle$ if $\mathcal{A}^{\mathcal{S}^\star}(\mathtt{C}) = \mathcal{A}^{\mathcal{T}}(\mathtt{C})$ for each circuit $\mathtt{C}$ encoded by $|\phi_\ell\rangle$. For any fixed $\mathtt{C}$, Lemma 4.7 yields that $\Pr_{\mathcal{T}}[\mathcal{A}^{\mathcal{S}^\star}(\mathtt{C}) = \mathcal{A}^{\mathcal{T}}(\mathtt{C})] \le \exp(-2^{n_o})$. Union bounding over all queried circuits $\mathcal{C}$, we have that the probability that $\mathcal{A}^{\mathcal{S}^\star}\,|\phi_\ell\rangle = \mathcal{A}^{\mathcal{T}}\,|\phi_\ell\rangle$ *for any* t is at most $|\mathtt{C}| \cdot \exp(-2^{n_o}) \le 2^{n_o}\exp(-2^{n_o}) =: p_1$.

For the second, Lemma 2.19 yields

$$\Pr_{\Delta}[\|\Pi_\Delta\,|\phi_\ell\rangle\,\| \ge 2N^{-1/3}] \le \exp\left(-N^{1/3}\right).$$

Thus, we may union bound over all L indices corresponding to this case, yielding that $\|(\mathcal{O}^{\mathcal{T}} - \mathcal{O}^{\mathcal{S}^\star})\,|\phi_\ell\rangle\,\|_2 \ge LN^{-1/3}$ for any ℓ with probability at most $p_2 := \exp\left(-N^{1/3}\right)$.

To complete the hybridization, with probability at least $1 - p_1 - p_2 \ge 1 - \exp\left(-N^{1/3}\right)$ over random $\mathcal{T}$ extending $\mathcal{S}^\star$,

$$\sum_{\ell=1}^{L} \|V[\ell+1] - V[\ell]\|_2 \le \mathcal{O}(LN^{-1/3}).$$

Therefore, to achieve the stated error ε, this requires $L \in \Omega(N^{1/3})$. □

Corollary 4.9. *There exists a quantum oracle $\mathcal{O}$ such that $\mathsf{QXC}^{\mathcal{O}} \not\subseteq \mathsf{QMA}^{\mathsf{QMA}^{\mathcal{O}}}$.*

5 UniqueQMA protocol for chaotic Hamiltonians

The oracle separations in the previous section suggest that we need to use the structure of QMA problems if we want to have a chance to design a unique verifier. In this section, we describe a UniqueQMA verification protocol for local Hamiltonians that are sufficiently chaotic in the sense of the eigenstate thermalization hypothesis (ETH) – an influential hypothesis in quantum statistical mechanics. We start by describing a computational variant of ETH. We then describe the protocol and analyze its performance under this assumption. Finally we discuss implications of the UniqueQMA vs QMA question to the thermalization properties of QMA-hard Hamiltonians.

5.1 Eigenstate thermalization hypothesis

Thermodynamics expects the evolution of a generic many-body system to be ergodic, meaning that the system evolves to eventually explore all possible configurations and thermalize. However, the evolution of a closed quantum system is governed by the Schrodinger's equation which is reversible and also leaves eigenstates invariant. Why do individual energy eigenstates of a large, interacting quantum system look thermal when probed by "simple" observables (like local operators), even though the full wavefunction evolves unitarily without randomness? A widespread explanation to this seeming paradox is the eigenstate thermalization hypothesis (ETH) [Deu91; Sre99], which has produced many predictions in agreements with thermalization on many physical systems [DKPR16]. ETH has many slight variants, but they are typically a set of assumptions about generic Hamiltonians that, informally, predicts that the eigenstates locally resemble a Gibbs state, and furthermore, local observables obey certain random matrix behaviors in the Hamiltonian energy eigenbasis inspired by quantum chaos and random matrix theory.

The first mathematically clean and consistent description of ETH was proposed by Srednicki [Sre99], which is in turn inspired by Wigner's pioneering random matrix theory (RMT) framework of analyzing quantum systems using statistical methods. Srednicki's ETH proposes a "chaotic" behavior of local operators in the eigenbasis of a local Hamiltonian. It is hard to characterize precisely how a local operator A may act on an energy eigenstate $|\nu_i\rangle$, so ETH leverages randomness to state that the matrix elements of a local operator A "looks like" $\langle\nu_j|A|\nu_i\rangle = c_i\delta_{i,j} + e^{-S/2}R_{ij}$, where c_i is some quantity capturing $\langle\nu_i|A|\nu_i\rangle$ and $|\nu_i\rangle, |\nu_j\rangle$ are eigenstates with nearby eigen-energies (a local operator cannot connect far away eigenstates - see [AKL16]) and S is the entropy within this window. Crucially, R_{ij} are chosen to be i.i.d random Gaussians. More formally, we consider the following formulation closely following [CB23, Section IV], which is arguably more mathematically concrete than [Sre99] to work with at a rigorous level.

Let us first introduce some notations. Fix an energy $e_0 > n^{0.99}$ and $\Delta = 1/\text{poly}(n)$ throughout this section[13]. Denote by $\mathcal{I}_\Delta$ the energy window of width $\Delta/2$ around e_0, $\mathcal{I}_\Delta = [e_0 - \frac{\Delta}{2}, e_0 + \frac{\Delta}{2}]$. Denote by Π_Δ the projector onto the subspace of eigenstates with energy $e \in \mathcal{I}_\Delta$. We will always think of the eigenstates as ordered by non-decreasing energies.

Hypothesis 5.1 (Srednicki's ETH ansatzes)**.** *Consider an n-qubit local Hamiltonian H and a narrow energy window of width Δ around an energy e_0. The corresponding "ETH ansatz ensemble", denoted $\mathcal{G}$, is a distribution over Hermitian operators $A \sim \mathcal{G}$ with parameters $\{\mu_\alpha\}_\alpha$ and $\{f_{\alpha,\beta}\}_{\alpha,\beta}$, whose entries within the said window have the form*

$$\langle\nu_\alpha|A|\nu_\beta\rangle = \mu_\alpha\delta_{\alpha,\beta} + \frac{f_{\alpha,\beta}}{\sqrt{\text{Tr}(\Pi_\Delta)}}g_{\alpha,\beta}, \tag{23}$$

whenever the energies of the eigenstates $|\nu_\alpha\rangle$, $|\nu_\beta\rangle$ are in the interval $\mathcal{I}_\Delta$. In the off-diagonal entries, $g_{\alpha,\beta} = g^*_{\beta,\alpha}$ are complex i.i.d. Gaussians with mean 0 and variance 1. The real coefficients $f_{\alpha,\beta}$ represent flexibility in the variance of the Gaussians, subject to $f_{\alpha,\beta} = f_{\beta,\alpha}$ due to the Hermitian condition and $f \le f_{\alpha,\beta} \le 1$, where f is a constant. The diagonal entries $\mu_\alpha = \langle\nu_\alpha|A|\nu_\alpha\rangle$ are non-random, but smoothly vary within the window $\mathcal{I}_\Delta$ (as a function of the eigenenergy). The normalization factor $\text{Tr}(\Pi_\Delta)$ is simply

[13] ETH is usually defined for "extensive" energies, those that scale linearly with the system size. Our threshold $n^{0.99}$ is just a way to capture this without introducing another parameter to set the extensive threshold.

the number of eigenstates within this window, as Π_Δ denotes the projector onto them. All these prescriptions are subject to the (global) constraints $A = A^\dagger$ and $\|A\| = 1$. [14]

We say that H and a collection of local [15] *observables $A_1, \ldots, A_m$ satisfy ETH within the $\mathcal{I}_\Delta$ window around energy e_0 if these operators "look like" randomly drawn from the ETH ensemble.*

Remark 5.2 (On the magnitudes of Δ and f). See [CB23, Section IV] for a discussion on the value of $\Delta = 1/\operatorname{poly}(n)$ for which ETH is expected to hold. Its value can depend on e_0, but we will only consider a fixed energy e_0 throughout and thus ignore such considerations to keep the model simple. The assumption of f being a system size-independent constant, although being quite standard in the ETH literature, could be relaxed. Our arguments will also work if f is inverse polynomially small for a mild polynomial. For example, [DKPR16] numerically suggested that $f = \Omega(1/n^{1/D})$ if the Hamiltonian is on a D-dimensional lattice. The condition $f_{\alpha,\beta} \leq 1$ can be relaxed to a larger upper bound, but we do not consider this to keep track of less variables.

For technical reasons, we will introduce two smoothness properties, which are often not explicitly included in ETH statements, but are expected to hold.

Property 1 (Non-clustering of eigenstates). Consider an energy window $\mathcal{I}_\omega \subset \mathcal{I}_\Delta$, where $\omega = \Delta - q(n)$ and $q(n) = 1/\operatorname{poly}(n)$ is sufficiently small. It holds that $\frac{\operatorname{Tr}(\Pi_\omega)}{\operatorname{Tr}(\Pi_\Delta)} \geq C' := 1 - f^{25}$.

This non-clustering property is used in the proof of Claim 5.10 when we analyze a quantum phase estimation subroutine, which has finite energy resolution. This property is justified since within the small window $\mathcal{I}_\Delta$ we expect the eigenstate density to be smooth. In fact, Property 1 is somewhat implicit by ETH since ETH is supposed to reduce to RMT when the energy interval is small [DKPR16], and RMT ensemble has a smooth density of states.

Similarly, we expect the quantities μ_α do not vary too much within the window $\mathcal{I}_\Delta$.[16] This brings us to our next assumption, which simplifies an analysis in the proof of Proposition 5.7.

Property 2 (Smoothness of diagonal entries). It holds that $|\max_{\alpha\in\mathcal{I}_\Delta} \mu_\alpha - \min_{\alpha\in\mathcal{I}_\Delta} \mu_\alpha| \leq f^7$.

We refer the interested readers to [Deu18] to a recent short overview of ETH that focuses more on the physical motivations, and the much more comprehensive [DKPR16] that further consists of mathematical foundations, random matrix theory background, applications of ETH, the various pieces of evidence, and important references in the vast ETH literature.

A quick objection to the ETH ansatz is - where is the randomness coming from? And what does "look like" mean? Clearly the Hamiltonian is fixed and the statement is being made for all eigenstates in the energy window. The view is that we make a choice of $g_{\alpha\beta}$ by drawing from iid Gaussian and then fix them for the entire calculation (and any predictions are taken "with high probability"). Given that H is a complicated object, the belief is that we would be unable to tell a difference (with high probability over the ETH ensemble) if we did calculations using the precise entries of the A's in the energy eigenbasis. Such behaviors have been experimentally and numerically verified in many physical quantum systems [RDO08], as well as generic, random Hamiltonians, such as the famous SYK model, where ETH is postulated to hold, see [DKPR16, Section 4.3].

In this work, we take this computational indistinguishability viewpoint to design a unique verification protocol for ETH Hamiltonians, but we hope the formulation below could be useful for other applications of ETH in complexity theory.

Hypothesis 5.3 (Computational ETH)**.** *Setup:*

(The Hamiltonian) *Consider an n-qubit local Hamiltonian H and a fixed energy window $\mathcal{I}_\Delta$ around e_0. Suppose that H satisfies Property 1 (non-clustering of eigenstates) in the window $\mathcal{I}_\Delta$.*

[14] This can be guaranteed by the countering fluctuations in the entries outside of the window $\mathcal{I}_\Delta$.

[15] The locality of A_i is typically allowed up to $o(n)$, but A_i should be simple (like Pauli, or efficiently implementable). For this work, it suffices to think of A_i as constant-local Pauli operators.

[16] Indeed, ETH says that an eigenstate locally resembles the Gibbs state. The temperature variation along the energy window $\mathcal{I}_\Delta$ can only be $1/\operatorname{poly}(n)$.

(The Experiment) *Consider a polynomial-time algorithm $\mathcal{N}$ that takes in an input state $|\xi\rangle$ promised to be inside $\mathcal{I}_\Delta$ and runs a polynomial-sized circuit (possibly with ancillas) that ends with a binary POVM $\{\mathsf{Acc}, \mathrm{I} - \mathsf{Acc}\}$. Furthermore, there is a set of $m = \mathrm{poly}(n)$ local Pauli operators $\mathtt{Alg} = \{A_1, \ldots, A_m\}$, such that $\mathcal{N}$ consists of gates that depend directly as a black-box on $\{A_1, \ldots, A_m\}$ and the individual local terms in H (which for examples could be time evolutions of any of these operators, or measurements, etc.)*

(The ETH Ensemble) *Let $\mathcal{N}'$ be an algorithm obtained by the following procedure: (1) iid sample a set of operators $\mathtt{Alg}' = \{A'_1, \ldots, A'_m\}$ from the operator ensemble defined in Hypothesis 5.1, (2) replace the components related to $\mathtt{Alg}$ in $\mathcal{N}$ by their $\mathtt{Alg}'$ versions [17]. More precisely, within the window $\mathcal{I}_\Delta$, each A'_i has the form*

$$\langle \nu_\alpha | A'_i | \nu_\beta \rangle = \mu^i_\alpha \delta_{\alpha,\beta} + \frac{f_{\alpha,\beta}}{\sqrt{\mathrm{Tr}(\Pi_\Delta)}} g^i_{\alpha,\beta}, \tag{24}$$

where the coefficients and random variables are as described in Hypothesis 5.1. Note that the diagonal entries μ^i_α are non-random and thus depend on A_i, so that distinct A'_i is sampled from a distinct distribution. Furthermore, we assume Property 2 holds for each i. The $g^i_{\alpha,\beta}$ are normal complex Gaussian variables sampled iid across different i.

Then, we say H and $\mathcal{N}$ satisfy the computational ETH within the energy window $\mathcal{I}_\Delta$ if it holds, with high probability over the randomness of the ETH ensemble, that

$$\mathrm{Tr}((\mathcal{N}(\xi) - \mathcal{N}'(\xi))\mathsf{Acc}) < \varepsilon_{\mathrm{ETH}} = \mathrm{negl}(n), \qquad \forall \text{ states } \xi \text{ of energy } \in \mathcal{I}_\Delta. \tag{25}$$

The above computational version is an attempt to capture in the computer science language the common physicists' interpretations of ETH. Above, we allow the 'experiment' to be beyond BQP: it can be given a proof (untrusted, except for the fact that its energy is within $\mathcal{I}_\Delta$ which can be efficiently verified). This allows us to work in the context of QMA. The $\varepsilon_{\mathrm{ETH}}$ is similar to the distinguishability advantage encountered in cryptography. Here we take it to be $\mathrm{negl}(n)$, but we actually only need $\varepsilon_{\mathrm{ETH}}$ to be a sufficiently small constant to guarantee the performance of our UniqueQMA verifier in this section.

5.2 Verifying a unique state in an energy window satisfying ETH

Given a local Hamiltonian H on n qubits, in this section, we give a protocol that uniquely accepts a specific state in a known energy window around e_0 satisfying certain properties that will be instantiated by ETH.

The intuition behind the protocol comes from tests of entanglement via quantum expanders [AHLN+14]. More precisely, the EPR state $|\Omega\rangle = \frac{1}{\sqrt{D}} \sum_{i=1}^{D} |ii\rangle$ can be tested locally using unitaries of the form $U \otimes \overline{U}$. We briefly explain the idea: Clearly, $U \otimes \overline{U}|\Omega\rangle = |\Omega\rangle$ for any unitary $U \in SU(D)$ so any averaged operator over the actions $U_i \otimes \overline{U}_i$ for unitaries $U_i \in SU(D), i = 1, \ldots, M$ satisfies $\frac{1}{M} \sum_{i=1}^{M} U_i \otimes \overline{U}_i |\Omega\rangle = |\Omega\rangle$. A quantum expander is roughly a collection of unitaries U_i, such that the second highest eigenvalue of $\frac{1}{M} \sum_{i=1}^{M} U_i \otimes \overline{U}_i$ is bounded away from 1. Such quantum expander families exist with surprisingly small M and constant gap between second highest eigenvalue and 1 (see e.g. [GE07; Has07]). Another common approach, used e.g. for self-testing, is to pick the local Pauli operators, which results in an inverse polynomial gap between the second highest eigenvalue and 1 (see e.g. [NV16]). The expectation value $\frac{1}{M} \sum_{i=1}^{M} \langle \psi | U_i \otimes \overline{U}_i | \psi \rangle$ can be efficiently estimated in many settings and suggests a simple test: If the above expectation is close to 1, then $|\psi\rangle$ is close to $|\Omega\rangle$. Otherwise, $|\psi\rangle$ and $|\Omega\rangle$ are far from each other. In particular, $|\Omega\rangle$ is uniquely accepted by this text. A plausible strategy is therefore to use this principle to uniquely verify the maximally entangled state on two copies of the space spanned by the eigenstates in an energy window. It will turn out that we cannot actually guarantee that the uniquely accepted state is the maximally entangled state on these subspaces. But for our purposes it will suffice that the Perron-Frobenius theorem guarantees the existence of a non-degenerate highest eigenvalue for an eigenstate with non-negative entries (see Claim 5.9). Then, Merlin can simply send this state and Arthur can uniquely verify the ground state energy via a quantum expander.

[17] These A' may require exponential circuit complexity, but we think of them as a black box, and these replacements only appears in our calculations.

Two obstacles prevent us from applying this naively: 1) A general unitary will move not respect the subspace spanned by eigenstates with energies in a window and 2) even if we can apply such unitaries we need to bound the gap of the averaging operator to guarantee uniqueness. We solve 1) by considering unitaries of the form $e^{-\iota\varepsilon A}$ for sufficiently small, but constant $\varepsilon > 0$, which approximately preserve subspaces. We then use phase estimation to prepare measurements of the energy window and effectively apply the projector onto the respective subspaces (see Claim 5.4). In order to approach 2) we invoke the ETH: By choosing the A_i to be local (think local Pauli matrices), their matrix entries in the energy window are indistinguishable from the Gaussians in Hypothesis 5.1. We then use probabilistic techniques to show with high probability over the Gaussians A_i' (see Claim 5.8) that the operators $\mathbb{E}_i e^{-\iota\varepsilon A_i} \otimes e^{\iota\varepsilon A_i}$ are sufficiently gapped (see Claim 5.9). In particular, for the Gaussian A_i' we obtain a gap and, consequently, unique verification. The ETH then guarantees that any quantum algorithm using the correct A_i will have the same behavior.

Finally, we assume there is a known energy e_0 and interval $\mathcal{I}_\Delta$ such that the above properties holds.

Assumption 1 (Existence of "good" Energy)**.** *There exists a known energy $e_0 \in (0,1)$ and value Δ such that Hypothesis 5.3 holds. Furthermore, e_0 can be represented by $O(\log n)$ bits.*

5.2.1 Phase estimation subroutine

In this subsection we describe a way to approximately perform the projector-valued measurement $\{\Pi_\Delta, \mathrm{I} - \Pi_\Delta\}$ using quantum phase estimation (QPE). This will be used in the testing algorithm. We label the system register as $\mathsf{R}_{\mathrm{sys}}$. Introduce an ancilla register $\mathsf{R}_{\mathrm{aux}}$ of dimension L for QPE, where $L \in \mathrm{poly}(n)$ is the resolution of the QPE algorithm and chosen to be large enough such that e_0 can be represented exactly (see Assumption 1). Consider the energy window $\mathcal{I}_\omega$, where $\omega = \Delta - 1/\sqrt{L}$. By choice of L the values $(e_0 \pm \omega/2)L$ are integers exactly represented by the P register. Let $\Pi_{\mathsf{R}_{\mathrm{aux}}}$ be the projector onto the state $|\mu\rangle_{\mathsf{R}_{\mathrm{aux}}} = \frac{1}{\sqrt{L}} \sum_{k=0}^{L-1} |k\rangle_{\mathsf{R}_{\mathrm{aux}}}$. Define the operator U_{QPE} which performs controlled evolution by the Hamiltonian H, followed by the inverse quantum Fourier transform.

$$\begin{aligned} U_{\mathrm{QPE}} &\triangleq \left(\frac{1}{\sqrt{L}} \sum_{m,m'=0}^{L-1} e^{2\pi\iota \cdot \frac{-mm'}{L}} \mathbb{I} \otimes |m\rangle\langle m'|_{\mathsf{R}_{\mathrm{aux}}} \right) \left(\sum_{k=0}^{L-1} e^{2\pi\iota k H} \otimes |k\rangle\langle k|_{\mathsf{R}_{\mathrm{aux}}} \right) \\ &= \frac{1}{\sqrt{L}} \sum_{m,k} e^{2\pi\iota(H - m/L)k} \otimes |m\rangle\langle k|_{\mathsf{R}_{\mathrm{aux}}} . \end{aligned}$$

This allows us to define a projector $\Pi_{\mathsf{R}_{\mathrm{sys}}\mathsf{R}_{\mathrm{aux}}}$, which one should think of as a projector onto eigenstates in $\mathcal{I}_\omega$.

$$\Pi_{\mathsf{R}_{\mathrm{sys}}\mathsf{R}_{\mathrm{aux}}} = U_{\mathrm{QPE}}^{-1} \left(\mathrm{I}_S \otimes \sum_{m=(e_0-\omega/2)L}^{(e_0+\omega/2)L} |m\rangle\langle m| \right) U_{\mathrm{QPE}} = \frac{1}{L} \sum_{k,k'=0}^{L-1} \sum_{m=(e_0-\omega/2)L}^{(e_0+\omega/2)L} e^{2\pi\iota(H-m/L)(k-k')} \otimes |k'\rangle\langle k|_{\mathsf{R}_{\mathrm{aux}}} . \tag{26}$$

Note that

$$\begin{aligned} \Pi_{\mathsf{R}_{\mathrm{aux}}} \Pi_{\mathsf{R}_{\mathrm{sys}}\mathsf{R}_{\mathrm{aux}}} \Pi_{\mathsf{R}_{\mathrm{aux}}} &= \sum_{m=(e_0-\omega/2)L}^{(e_0+\omega/2)L} \left(\frac{1}{L^2} \sum_{k,k'=0}^{L-1} e^{2\pi\iota(H-m/L)(k-k')} \right) \otimes \Pi_{\mathsf{R}_{\mathrm{aux}}} \\ &= \sum_{m=(e_0-\omega/2)L}^{(e_0+\omega/2)L} \left(\mathrm{sinc}_L(H - m/L)\right)^2 \otimes \Pi_{\mathsf{R}_{\mathrm{aux}}} \\ &\triangleq Q(H) \otimes \Pi_{\mathsf{R}_{\mathrm{aux}}}, \end{aligned} \tag{27}$$

with the function $\mathrm{sinc}_L(x)$ defined as $\frac{\sin(\pi L x)}{L \sin(\pi x)}$. This function peaks at $x = 0$ and then sharply goes down to a zero at $1/L$ and then goes up and down at the rate roughly $1/L$. Summing over m would ideally yield the projector Π_ω. Rather than proving this, we show that $Q(H)$ behaves approximately like a projector over a

slightly larger space, given by the full window Π_Δ. This is formalized in the following claim which states that $Q(H)$ is mostly supported on Π_Δ.

Claim 5.4. $\|Q(H) - \Pi_\Delta Q(H)\| \leq \frac{2}{\sqrt{L}}$.

See Appendix A.2 for a proof of this claim. Now, using $Q(H)$ and the two Assumptions, we can define algorithms to test for the maximally entangled state over the subspace Π_Δ.

5.2.2 The testing protocol

We are now in a position to describe our protocol to test for energy subspaces. We first describe the algorithm formally: We reuse some notations from the phase estimation subroutine. `Alg` is described in Algorithm 1, which acts on two copies of S, denoted as $\mathsf{R}^{(1)}_{\text{sys}}, \mathsf{R}^{(2)}_{\text{sys}}$. Note that the projectors $\Pi_{\mathsf{R}_{\text{sys}}\mathsf{R}_{\text{aux}}}$ are implicitly parameterized by e_0.

Algorithm 1: Energy subspace test `Alg`.

Input: state $|\psi\rangle$ on the registers $\mathsf{R}^{(1)}_{\text{sys}}, \mathsf{R}^{(2)}_{\text{sys}}$

1 Prepare and append $|\mu\rangle_{\mathsf{R}^{(1)}_{\text{aux}}} |\mu\rangle_{\mathsf{R}^{(2)}_{\text{aux}}}$ to $|\psi\rangle_{\mathsf{R}^{(1)}_{\text{sys}},\mathsf{R}^{(2)}_{\text{sys}}}$, where $|\mu\rangle = \frac{1}{\sqrt{L}} \sum_{k=0}^{L-1} |k\rangle$.

2 Measure $\{\Pi_{\mathsf{R}^{(1)}_{\text{sys}}\mathsf{R}^{(1)}_{\text{aux}}}, \mathrm{I} - \Pi_{\mathsf{R}^{(1)}_{\text{sys}}\mathsf{R}^{(1)}_{\text{aux}}}\}$, $\{\Pi_{\mathsf{R}^{(2)}_{\text{sys}}\mathsf{R}^{(2)}_{\text{aux}}}, \mathrm{I} - \Pi_{\mathsf{R}^{(2)}_{\text{sys}}\mathsf{R}^{(2)}_{\text{aux}}}\}$ and reject if $\mathrm{I} - \Pi_{\mathsf{R}^{(1)}_{\text{sys}}\mathsf{R}^{(1)}_{\text{aux}}}$ or $\mathrm{I} - \Pi_{\mathsf{R}^{(2)}_{\text{sys}}\mathsf{R}^{(2)}_{\text{aux}}}$ is obtained.

3 Measure $\mathsf{R}^{(1)}_{\text{aux}}$ in the basis $\{\Pi_{\mathsf{R}^{(1)}_{\text{aux}}}, \mathrm{I} - \Pi_{\mathsf{R}^{(1)}_{\text{aux}}}\}$, $\mathsf{R}^{(2)}_{\text{aux}}$ in the basis $\{\Pi_{\mathsf{R}^{(2)}_{\text{aux}}}, \mathrm{I} - \Pi_{\mathsf{R}^{(2)}_{\text{aux}}}\}$ and reject if $\mathrm{I} - \Pi_{\mathsf{R}^{(1)}_{\text{aux}}}$ or $\mathrm{I} - \Pi_{\mathsf{R}^{(2)}_{\text{aux}}}$ is obtained.

4 Prepare a register $\mathsf{R}_{\text{control}}$ and define the state $\frac{1}{\sqrt{2\ell}} \sum_i (|i,0\rangle_{\mathsf{R}_{\text{control}}} + |i,1\rangle_{\mathsf{R}_{\text{control}}})$. Controlled on $|i,0\rangle\langle i,0|_{\mathsf{R}_{\text{control}}}$, apply $e^{\iota\varepsilon A_i}$ on $\mathsf{R}^{(1)}_{\text{sys}}$ and $e^{-\iota\varepsilon\overline{A}_i}$ on $\mathsf{R}^{(2)}_{\text{sys}}$. Controlled on $|i,1\rangle\langle i,1|_T$, apply $e^{-\iota\varepsilon A_i}$ on $\mathsf{R}^{(1)}_{\text{sys}}$ and $e^{\iota\varepsilon\overline{A}_i}$ on $\mathsf{R}^{(2)}_{\text{sys}}$.

5 Repeat Steps 2,3.

6 Measure $\mathsf{R}_{\text{control}}$ to see if the register is in $\frac{1}{\sqrt{2\ell}} \sum_i (|i,0\rangle_{\mathsf{R}_{\text{control}}} + |i,1\rangle_{\mathsf{R}_{\text{control}}})$. Reject if it is not.

Steps 1-3 use the phase estimation procedure explained in the previous section to effectively apply the approximate projector $Q(H)$ to the state $|\psi\rangle$ in registers $S_{1,2}$. Step 4 applies the elements of the quantum expander test conditioned on the register $\mathsf{R}_{\text{control}}$ and Step 5 restricts to the low energy subspace again. Finally, measuring the register $\mathsf{R}_{\text{control}}$ yields a high acceptance probability if and only if the state $|\psi\rangle$ has high overlap with the space Π_Δ spanned by states in the low enery window and sufficiently invariant under $e^{-\iota\varepsilon A_i} \otimes e^{\iota\varepsilon A_i}$. The resulting guarantees are summarized in the following theorem.

Theorem 5.5 (Energy Subspace Test)**.** *Fix a parameter e_0 and consider a Hamiltonian H. There exists an algorithm $\mathtt{Alg} = \mathtt{Alg}(H, e_0, L)$ (where $L = \text{poly}(n)$ is a precision parameter) such that*

1. *If $\lambda_0(H) > e_0 + \frac{3}{L}$ then `Alg` accepts with probability at most $\frac{1}{3}$.*
2. *If $\lambda_0(H) \leq e_0$ and e_0 satisfies Assumption 1 w.r.t. H, then there exists a unique state $|\Psi\rangle$ such that `Alg` accepts on input $|\Psi\rangle$ with probability at least $\frac{2}{3}$, and on states orthogonal to $|\Psi\rangle$, `Alg` accepts with probability at most $\frac{1}{3}$.*

Given a state $|\psi\rangle_{\mathsf{R}^{(1)}_{\text{sys}}\mathsf{R}^{(2)}_{\text{sys}}}$, the acceptance probability is given by

$$\begin{aligned}
&\|\mathbb{E}_i(\Pi_{\mathsf{R}^{(1)}_{\text{aux}}}\otimes\Pi_{\mathsf{R}^{(2)}_{\text{aux}}})(\Pi_{\mathsf{R}^{(1)}_{\text{sys}}\mathsf{R}^{(1)}_{\text{aux}}}\otimes\Pi_{\mathsf{R}^{(2)}_{\text{sys}}\mathsf{R}^{(2)}_{\text{aux}}})(\frac{1}{2}e^{\iota\varepsilon A_i}\otimes\Pi_{\mathsf{R}^{(1)}_{\text{aux}}}\otimes e^{-\iota\varepsilon\overline{A}_i}\otimes\Pi_{\mathsf{R}^{(2)}_{\text{aux}}}+\frac{1}{2}e^{-\iota\varepsilon A_i}\otimes\Pi_{\mathsf{R}^{(1)}_{\text{aux}}}\otimes e^{\iota\varepsilon\overline{A}_i}\otimes\Pi_{\mathsf{R}^{(2)}_{\text{aux}}})\\
&\quad(\Pi_{\mathsf{R}^{(1)}_{\text{sys}}\mathsf{R}^{(1)}_{\text{aux}}}\otimes\Pi_{\mathsf{R}^{(2)}_{\text{sys}}\mathsf{R}^{(2)}_{\text{aux}}})\,|\psi\rangle_{\mathsf{R}^{(1)}_{\text{sys}},\mathsf{R}^{(2)}_{\text{sys}}}|\mu\rangle_{\mathsf{R}^{(1)}_{\text{aux}}}|\mu\rangle_{\mathsf{R}^{(2)}_{\text{aux}}}\|^2\\
&\quad=\|\mathbb{E}_i(Q(H)\otimes Q(H))(\frac{1}{2}e^{\iota\varepsilon A_i}\otimes e^{-\iota\varepsilon\overline{A}_i}+\frac{1}{2}e^{-\iota\varepsilon A_i}\otimes e^{\iota\varepsilon\overline{A}_i})(Q(H)\otimes Q(H))\,|\psi\rangle\|^2\\
&\quad\overset{(a)}{=}\|(Q(H)\otimes Q(H))\big((1-\varepsilon^2)\mathrm{I}\otimes\mathrm{I}+\varepsilon^2\mathbb{E}_i\big(A_i\otimes\overline{A}_i\big)(Q(H)\otimes Q(H))\,|\psi\rangle\|^2+O(\varepsilon^3)
\end{aligned}$$

where (a) follows from a Taylor expanding $e^{\iota\varepsilon A_i}$ and $e^{-\iota\varepsilon A_i}$ and using $A_i^2=\mathrm{I}$. Note that the first order terms in ε cancel due to A_i and $-A_i$.

We can use Claim 5.4 to write the last equation as

$$\|(Q(H)\otimes Q(H))\big((1-\varepsilon^2)\Pi_\Delta\otimes\Pi_\Delta+\varepsilon^2\mathbb{E}_i\Pi_\Delta A_i\Pi_\Delta\otimes\Pi_\Delta\overline{A}_i\Pi_\Delta\big)(Q(H)\otimes Q(H))\,|\psi\rangle\|^2+O(\varepsilon^3+\frac{1}{\sqrt{L}}).$$

Thus, the success probability of the protocol is captured by the operator

$$O_{\text{succ}}:=(Q(H)\otimes Q(H))\big((1-\varepsilon^2)\Pi_\Delta\otimes\Pi_\Delta+\varepsilon^2\mathbb{E}_i\Pi_\Delta A_i\Pi_\Delta\otimes\Pi_\Delta\overline{A}_i\Pi_\Delta\big)(Q(H)\otimes Q(H)). \tag{28}$$

We'll separately analyze O_{succ} for the cases corresponding to Item 1 and Item 2. To prove the former case, we show the following lemma,

Lemma 5.6. *Let H be a Hamiltonian such that $\lambda_0(H)>e_0+\frac{3}{L}$. Then $\|O_{succ}\|\le\frac{1}{3}$.*

Proof. From Equation (28), we see that $O_{\text{succ}}=(Q(H)\otimes Q(H))\widetilde{O}(Q(H)\otimes Q(H))$, where $\|\widetilde{O}\|\le 1$. Therefore, we can bound $\|O_{\text{succ}}\|\le\|Q(H)\otimes Q(H)\|=\max_{|\psi\rangle}\|Q(H)\,|\psi\rangle\|_2^2=\max_\alpha\|Q(H)\,|v_\alpha\rangle\|_2^2$ where in the last equality, we restrict to eigenstates of H since $Q(H)$ is diagonal in H's eigenbasis. And $\lambda_\alpha-e_0\ge\frac{3}{L}$ by assumption, so Item 3 of Lemma A.8 tells us that $\|Q(H)\,|v_\alpha\rangle\|_2^2\le\frac{1}{3}$, which proves the claim. □

For the remainder of this section, we focus on the case where $\lambda(H)\le e_0$ and thus Assumption 1 holds. To show Item 2, it suffices to show that O_{succ} has a unique largest eigenstate and a large enough spectral gap. We will first simplify O_{succ} using the following proposition.

Proposition 5.7. *Let $m\ge n^4$ and let $\mathcal{G}$ be the ETH ensemble as defined in Hypothesis 5.1. With probability $1-e^{-O(m^{1/3})}$ over $\mathcal{G}$, we have*

$$\|O_{\text{succ}}-(Q(H)\otimes Q(H))M(Q(H)\otimes Q(H))\|\le 2\varepsilon^2 f^7+O(\varepsilon^2/m^{1/3})+\varepsilon_{\text{ETH}},$$

where $M:=(1-\varepsilon^2+\varepsilon^2\mathbb{E}_i(\mu^i)^2)\Pi_\Delta\otimes\Pi_\Delta+\varepsilon^2\mathbb{E}_G B\otimes\overline{B}$ and B is drawn from $\mathcal{G}$.

Proof. Since the protocol is a polynomial time quantum computation, we can invoke the computational ETH Hypothesis 5.3 to follow the RMT prescription while only incurring a small error ε_{ETH} in the final result. From Assumption 1, the term $\Pi_\Delta A_i\Pi_\Delta$ can be decomposed into the diagonal terms given by μ^i_α and the remaining terms which follow the RMT prescription. For the latter term, define

$$B_i:=\frac{1}{\sqrt{\mathrm{Tr}(\Pi_\Delta)}}\sum_{\alpha,\beta}f_{\alpha,\beta}g^i_{\alpha,\beta}\,|v_\alpha\rangle\langle v_\beta|$$

For the diagonal terms, we recall from Property 2 that $\mu^i:=\max_\alpha\mu^i_\alpha$ is close to all μ^i_α by f^7 and thus we approximate this term as $\mu^i\Pi_\Delta$. Thus,

$$\|\Pi_\Delta A_i\Pi_\Delta-\mu^i\Pi_\Delta-B_i\|\le f^7,$$

The same decomposition can be done for $\Pi_\Delta \overline{A}_i \Pi_\Delta$, with $\overline{B}_i$ in place of B_i and we get

$$\begin{aligned}&\mathbb{E}_i \Pi_\Delta A_i \Pi_\Delta \otimes \Pi_\Delta \overline{A}_i \Pi_\Delta \\ &\quad = \mathbb{E}_i (\mu^i)^2 \Pi_\Delta \otimes \Pi_\Delta + \Pi_\Delta \otimes \mathbb{E}_i \mu^i \overline{B}_i + \mathbb{E}_i \mu^i B_i \otimes \Pi_\Delta + \mathbb{E}_i B_i \otimes \overline{B}_i + E, \end{aligned} \tag{29}$$

with $\|E\| \leq 2f^7$. We will now repeatedly use the following lemma from Appendix A.1.

Lemma A.4. *Let P be a random Hermitian matrix of dimension $D \times D$ with each entry being an independent complex Gaussian with mean 0 and (non-identical) variance ≤ 1. Then with probability $1 - e^{-D}$, we have $\|P\| \leq 10\sqrt{D}$.*

Now, note that

$$\begin{aligned}\mathbb{E}_i \mu^i B_i &= \frac{1}{\sqrt{\mathrm{Tr}(\Pi_\Delta)}} \sum_{\alpha,\beta} f_{\alpha,\beta} (\mathbb{E}_i \mu^i g^i_{\alpha,\beta}) \, |v_\alpha\rangle\langle v_\beta| \\ &= \frac{\sqrt{\sum_i (\mu^i)^2}}{m\sqrt{\mathrm{Tr}(\Pi_\Delta)}} \sum_{\alpha,\beta} f_{\alpha,\beta} \frac{\sum_i \mu^i g^i_{\alpha,\beta}}{\sqrt{\sum_i (\mu^i)^2}} \, |v_\alpha\rangle\langle v_\beta| := \frac{\sqrt{\sum_i (\mu^i)^2}}{m} B'. \end{aligned}$$

Note that $\sqrt{\sum_i (\mu^i)^2} \leq \sqrt{m}$. By Fact A.3 we see that $\frac{1}{\sqrt{\sum_i (\mu^i)^2}} \sum_i \mu^i g^i_{\alpha,\beta} \sim \mathcal{CN}(0,1)$ and thus $\|B'\| \leq 10$ by Lemma A.4. This means that with probability $1 - e^{-\mathrm{Tr}(\Pi_\Delta)}$, $\|\mathbb{E}_i \, \mu^i B_i\| \leq \frac{1}{\sqrt{m}}$. The same argument works for $\mathbb{E}_i \mu^i \overline{B}_i$. Thus,

$$\mathbb{E}_i \Pi_\Delta A_i \Pi_\Delta \otimes \Pi_\Delta \overline{A}_i \Pi_\Delta = \mathbb{E}_i (\mu^i)^2 \Pi_\Delta \otimes \Pi_\Delta + \mathbb{E}_i B_i \otimes \overline{B}_i + E',$$

with $\|E'\| \leq 2f^6 + O(1/\sqrt{m})$.

Finally, we will show the following claim.

Claim 5.8. *Let $m \geq n^4$ and let $\mathcal{G}$ be the distribution over Gaussians as defined in Hypothesis 5.1. With probability $1 - e^{-\mathcal{O}(m^{1/3})}$ (over the randomness of ETH), we have*

$$\mathbb{E}_i B_i \otimes \overline{B}_i = \mathbb{E}_{\mathcal{G}} B \otimes \overline{B} + \mathcal{O}(m^{-1/3}),$$

where B is drawn from $\mathcal{G}$.

Proof. From Lemma A.4, we know that each B_i has norm $\|B_i\| \leq 10$ with probability $1 - e^{-\mathcal{O}(\mathrm{Tr}(\Pi_\Delta))}$. Each B_i was sampled IID from $\mathcal{G}$. Let $\mathcal{G}'$ be the distribution obtained when we condition $\mathcal{G}$ to Gaussians which ensure $\|B\| \leq 10$. Note that $\|\mathcal{G} - \mathcal{G}'\|_1 \leq e^{-\mathcal{O}(\mathrm{Tr}(\Pi_\Delta))}$. Let $\mathcal{B}(\mathcal{G}) = \mathbb{E}_{\mathcal{G}} B \otimes \overline{B}$ and $\mathcal{B}(\mathcal{G}') = \mathbb{E}_{\mathcal{G}'} B \otimes \overline{B}$. Note that

$$\|\mathcal{B}(\mathcal{G}) - \mathcal{B}(\mathcal{G}')\| \leq 2e^{-\mathcal{O}(\mathrm{Tr}(\Pi_\Delta))} \max_B \|B\|^2 \leq e^{-\mathcal{O}(\mathrm{Tr}(\Pi_\Delta))},$$

since for any B, the operator norm is at most $\mathsf{poly}(\mathrm{Tr}(\Pi_\Delta))$ due to the Greshgorin disc theorem. Thus, for $t := m^{-1/3}$,

$$\begin{aligned}\Pr_{\mathcal{G}}[\|\mathbb{E}_i B_i \otimes \overline{B}_i - \mathcal{B}(\mathcal{G})\| \geq t] &\leq \Pr_{\mathcal{G}'}[\|\mathbb{E}_i B_i \otimes \overline{B}_i - \mathcal{B}(\mathcal{G}')\| \geq t - e^{-\mathcal{O}(\mathrm{Tr}(\Pi_\Delta))}] \\ &\leq \Pr_{\mathcal{G}'}[\|\mathbb{E}_i B_i \otimes \overline{B}_i - \mathcal{B}(\mathcal{G}')\| \geq t/2]\end{aligned}$$

where the subscript refers to the distribution according to which $B_1, \ldots B_m$ are sampled. We will show in Claim 5.9 that $\|\mathcal{B}(\mathcal{G})\| \leq 1$, which means for draws from $\mathcal{G}'$, $\|B_i \otimes \overline{B}_i - \mathcal{B}(\mathcal{G}')\| \leq 12$. Then, we apply the Matrix Chernoff bound [Tro15, Theorem 6.6.1] to the matrices $B \otimes \overline{B} - \mathcal{B}(\mathcal{G}')$, drawn iid from $\mathcal{G}'$. The bound says that

$$\Pr_{\mathcal{G}'}(\|\mathbb{E}_i B_i \otimes \overline{B}_i - \mathcal{B}(\mathcal{G}')\| \geq t/2) \leq 2 \, \mathrm{Tr}(\Pi_\Delta) e^{-t^2 m^2/(8v + 16mt)},$$

where

$$v = \|\sum_i \mathbb{E}_{\mathcal{G}'}(B_i \otimes \overline{B_i} - \mathcal{B}(\mathcal{G}'))^2\| = m\|\mathbb{E}_{\mathcal{G}'}(B \otimes \overline{B})^2 - \mathcal{B}(\mathcal{G}')^2\| \le m\mathbb{E}_{\mathcal{G}'}\|(B \otimes \overline{B} - \mathcal{B}(\mathcal{G}'))^2\| \le 144m.$$

Thus

$$\Pr_{\mathcal{G}'}(\|\mathbb{E}_i B_i \otimes \overline{B}_i - \mathcal{B}(\mathcal{G}')\| \ge t/2) \le 2\operatorname{Tr}(\Pi_\Delta) e^{-t^2m^2/(9\times 2^7 m + 16mt)} \le 2^n e^{-\mathcal{O}(m^{1/3})} = e^{-\mathcal{O}(m^{1/3})}.$$

This completes the proof. □

As a result, we conclude that with probability $1 - e^{-\mathcal{O}(m^{1/3})}$,

$$\mathbb{E}_i \Pi_\Delta A_i \Pi_\Delta \otimes \Pi_\Delta \overline{A}_i \Pi_\Delta = \mathbb{E}_i(\mu^i)^2 \Pi_\Delta \otimes \Pi_\Delta + \mathbb{E}_G B \otimes \overline{B} + E'',$$

with $\|E''\| \le 2f^7 + \mathcal{O}(1/m^{1/3})$. Finally, we add an error term ε_{ETH} from Hypothesis 5.3 (for simplicity we assume the quantifier 'with high probability' in that hypothesis statement is $1 - e^{-\operatorname{poly}(n)} > 1 - e^{-\mathcal{O}(m^{1/3})}$). This completes the proof of Proposition 5.7. □

Next claim helps us understand the spectrum of $\mathbb{E}_G B \otimes \overline{B}$.

Claim 5.9. *The matrix $\mathbb{E}_{\mathcal{G}} B \otimes \overline{B}$ has a unique leading eigenvector $|\Psi\rangle$ with eigenvalue $1 \ge \lambda \ge f^2$ and all other eigenvalues are below $(1-f^4)\lambda$. Furthermore, $|\Psi\rangle$ has all positive entries with ratios bounded between $\frac{1}{f^2}$ and f^2. Thus, the matrix M has unique max eigenvector $|\Psi\rangle$ with eigenvalue $\lambda_0 := 1 - \varepsilon^2 + \varepsilon^2 \mathbb{E}_i(\mu^i)^2 + \varepsilon^2\lambda$ and spectral gap is $\Delta \ge \varepsilon^2 f^4 \lambda$.*

Proof. We have

$$\begin{aligned}
\mathbb{E}_{\mathcal{G}} B \otimes \overline{B} &= \frac{1}{\operatorname{Tr}(\Pi_\Delta)} \sum_{\alpha,\beta,\alpha',\beta'} f_{\alpha,\beta} f_{\alpha',\beta'} \left(\mathbb{E}_{\mathcal{G}} g_{\alpha,\beta} g^*_{\alpha',\beta'}\right) |v_\alpha\rangle\langle v_\beta| \otimes |v_{\alpha'}\rangle\langle v_{\beta'}| \\
&\overset{\text{(b)}}{=} \frac{1}{\operatorname{Tr}(\Pi_\Delta)} \sum_{\alpha,\beta} f^2_{\alpha,\beta} |v_\alpha, v_\alpha\rangle \langle v_\beta, v_\beta|
\end{aligned}$$

where the summations above and onward are over $\alpha, \beta, \alpha', \beta'$ such that $\lambda_\alpha, \lambda_\beta, \lambda_{\alpha'}, \lambda_{\beta'} \in \mathcal{I}_\Delta$ (rather than in $\mathcal{I}_{\Delta_{\text{RMT}}}$). For (b) we use that $\mathbb{E}|g_{\alpha\beta}|^2 = 1$, $\mathbb{E} g^2_{\alpha\beta} = 0$.

Since the matrix has positive entries, by the Perron-Frobenius theorem, its unique leading eigenvector is $|\Psi\rangle = \sum_\alpha x_\alpha |v_\alpha, v_\alpha\rangle$. Let λ be the corresponding eigenvalue. In this proof, we think of $\alpha' s$ as integers between $1, \ldots, N = \operatorname{Tr}(\Pi_\Delta)$, ordered such that x_1 has the largest value and x_N the smallest. The fact that $|\Psi\rangle$ is an eigenvector implies

$$\frac{1}{\operatorname{Tr}(\Pi_\Delta)} \sum_\beta f^2_{\alpha,\beta} x_\beta = \lambda x_\alpha, \qquad \forall \alpha \in \mathcal{I}_\Delta. \tag{30}$$

Choosing $\alpha = N$ in the above and solving for λ yields,

$$\lambda = \frac{1}{\operatorname{Tr}(\Pi_\Delta)} \sum_\beta f^2_{N,\beta} \frac{x_\beta}{x_N} \ge \frac{1}{\operatorname{Tr}(\Pi_\Delta)} \sum_{\beta \in \mathcal{I}_\Delta} f^2_{N,\beta} \cdot 1 \ge f^2,$$

where the last inequality follows from our assumption on $f_{\alpha,\beta}$ from Hypothesis 5.3. We also have the upper bound 1 on λ using the Greshgorin disc theorem.

Now, taking ratios of Equation (30) for $\alpha = 1$ and $\alpha = N$ yields

$$\frac{\sum_\beta f^2_{1,\beta} x_\beta}{\sum_\beta f^2_{N,\beta} x_\beta} = \frac{x_1}{x_N} \implies \frac{1}{f^2} \ge \frac{x_1}{x_N} \ge 1. \tag{31}$$

where we've used that $f \leq f_{\alpha,\beta} \leq 1$. This shows the desired properties of the leading eigenvector $|\Psi\rangle$, except the gap.

Take an eigenstate orthogonal to $|\Psi\rangle$, denoting it $|\Phi\rangle = \sum_\alpha y_\alpha |v_\alpha, \overline{v_\alpha}\rangle$. Note that we can assume WLOG that y_a's are real since M_f is a matrix with real entries in the $|v_\alpha\rangle$ basis. Define the sets $S_+ = \{y_\alpha : y_\alpha > 0\}$ and $S_- = \{y_\alpha : y_\alpha \leq 0\}$, extending to sets of tuples $S_{(-,-)}, S_{(+,+)}, S_{(-,+)}, S_{(+,-)}$ in the natural way. The eigenvalue corresponding to $|\Phi\rangle$ can be written as,

$$\begin{aligned}
\mathrm{Tr}(\Pi_\Delta) \langle\Phi| \mathbb{E}_G B \otimes \overline{B} |\Phi\rangle &= \sum_{\alpha,\beta} f_{\alpha,\beta}^2 y_\alpha y_\beta \\
&= \sum_{\alpha\neq\beta\in S_{(+,+)}} f_{\alpha,\beta}^2 y_\alpha y_\beta + \sum_{\alpha\neq\beta\in S_{(-,-)}} f_{\alpha,\beta}^2 y_\alpha y_\beta \\
&\quad - \sum_{\alpha,\beta\in S_{(+,-)}} f_{\alpha,\beta}^2 y_\alpha |y_\beta| - \sum_{\alpha,\beta\in S_{(-,+)}} f_{\alpha,\beta}^2 |y_\alpha| y_\beta \\
&= \Sigma_{(+,+)} + \Sigma_{(-,-)} - \Sigma_{(+,-)} - \Sigma_{(-,+)} \qquad (32)
\end{aligned}$$

where we've defined $\Sigma_{(\cdot,\cdot)}$ to refer to the (positive) sums over $S_{(\cdot,\cdot)}$. Notice that $\langle\Phi| \mathbb{E}_G B \otimes \overline{B} |\Phi\rangle$ includes the negative terms $-\Sigma_{(+,-)} - \Sigma_{(-,+)}$, whereas the leading eigenvector $|\Psi\rangle$ has non-negative entries (by Perron-Frobenius) and thus no such terms. This is the key aspect which will manifest the spectral gap.

We claim that the ratios between the terms in Equation (32) can be multiplicatively bounded. In particular,

$$\begin{aligned}
\frac{\Sigma_{(+,+)}}{\Sigma_{(-,+)}} = \frac{\sum_{\alpha\neq\beta\in S_{+,+}} f_{\alpha,\beta}^2 y_\alpha y_\beta}{\sum_{\alpha,\beta\in S_{-,+}} f_{\alpha,\beta}^2 |y_\alpha| y_\beta} &\overset{(c)}{\leq} \frac{\sum_{\alpha\neq\beta\in S_{+,+}} y_\alpha y_\beta}{f^2 \sum_{\alpha,\beta\in S_{-,+}} |y_\alpha| y_\beta} \\
&\leq \frac{(\sum_{\alpha\in S_+} y_\alpha)^2}{f^2 (\sum_{\alpha\in S_-} |y_\alpha|)(\sum_{\alpha\in S_+} y_\alpha)} \\
&= \frac{\sum_{\alpha\in S_+} y_\alpha}{f^2 (\sum_{\alpha\in S_-} |y_\alpha|)} \qquad (33)
\end{aligned}$$

where (c) follows by $f^2 \leq f_{\alpha,\beta}^2 \leq 1$. By the assumption that $\langle\Phi|\Psi\rangle = 0$, we have that $\sum_{\alpha\in S_+} y_\alpha x_\alpha = \sum_{\alpha\in S_-} |y_\alpha| x_\alpha$ and using the bounds $x_N \leq x_\alpha \leq x_1$ we obtain

$$x_N \sum_{\alpha\in S_+} y_\alpha \leq \sum_{\alpha\in S_+} y_\alpha x_\alpha \leq x_1 \sum_{\alpha\in S_-} |y_\alpha|. \qquad (34)$$

Combining with Equation (33) yields

$$\frac{\Sigma_{(+,+)}}{\Sigma_{(-,+)}} = \frac{\sum_{\alpha\in S_+} y_\alpha}{f^2 (\sum_{\alpha\in S_-} |y_\alpha|)} \leq \frac{x_1}{f^2 x_N} \leq \frac{1}{f^4}. \qquad (35)$$

Similarly, we can bound $\frac{\Sigma_{(+,-)}}{\Sigma_{(-,+)}}$ as in Equation (33) but instead we get

$$\frac{\Sigma_{(+,-)}}{\Sigma_{(-,+)}} \leq \frac{1}{f^2} \frac{\sum_{\alpha\in S_+} y_\alpha}{\sum_{\alpha\in S_-} |y_\alpha|} \cdot \frac{\sum_{\alpha\in S_-} |y_\alpha|}{\sum_{\alpha\in S_+} y_\alpha} = \frac{1}{f^2} \leq \frac{1}{f^4}.$$

The bounds on $\frac{\Sigma_{(-,-)}}{\Sigma_{(+,-)}}$ and $\frac{\Sigma_{(-,+)}}{\Sigma_{(+,-)}}$ follow identically. This allows us to conclude that,

$$2\Sigma_{(-,+)} \geq f^4 \left(\Sigma_{(+,+)} + \Sigma_{(+,-)}\right), \qquad (36)$$

$$2\Sigma_{(-,+)} \geq f^4 \left(\Sigma_{(-,-)} + \Sigma_{(-,+)}\right), \qquad (37)$$

and therefore

$$\begin{aligned}\Sigma_{(+,+)}+\Sigma_{(-,-)}-\Sigma_{(+,-)}-\Sigma_{(-,+)} &= \Sigma_{(+,+)}+\Sigma_{(-,-)}+\Sigma_{(+,-)}+\Sigma_{(-,+)}\\ &\quad -2\left(\Sigma_{(+,-)}+\Sigma_{(-,+)}\right)\\ &\le \Sigma_{(+,+)}+\Sigma_{(-,-)}+\Sigma_{(+,-)}+\Sigma_{(-,+)}\\ &\quad -f^4\left(\Sigma_{(+,+)}+\Sigma_{(-,-)}+\Sigma_{(+,-)}+-\Sigma_{(-,+)}\right)\\ &= \left(1-f^4\right)\left(\Sigma_{(+,+)}+\Sigma_{(-,-)}+\Sigma_{(+,-)}+\Sigma_{(-,+)}\right). \end{aligned} \tag{38}$$

This leads to an upper bound on Equation (32) as,

$$\begin{aligned}\langle\Phi|\,\mathbb{E}_G B\otimes\overline{B}\,|\Phi\rangle &\le \frac{1}{\mathrm{Tr}(\Pi_\Delta)}\left(1-f^4\right)\left(\Sigma_{(+,+)}+\Sigma_{(-,-)}+\Sigma_{(+,-)}+\Sigma_{(-,+)}\right)\\ &= \left(1-f^4\right)\left(\sum_{\alpha,\beta}(\mathbb{E}_G B\otimes\overline{B})_{\alpha,\beta}|y_\alpha||y_\beta|\right)\\ &\overset{(d)}{\le} (1-f^4)\lambda \end{aligned} \tag{39}$$

where in (d) we've used that the leading eigenvector has positive entries and thus upper bounds the quadratic form $\sum_{\alpha,\beta}(\mathbb{E}_G B\otimes\overline{B})_{\alpha,\beta}|y_\alpha||y_\beta|$. □

Finally, we move to the spectral gap of $(Q\otimes Q)M(Q\otimes Q)$ (with $Q:=Q(H)$), which captures the success probability of the protocol, as shown in Proposition 5.7. The following claim, whose proof is deferred to Appendix, shows that this conjugation does not change $|\Psi\rangle$ much.

Claim 5.10. *We have* $\langle\Psi|\,Q\otimes Q\,|\Psi\rangle \ge 1-\frac{1}{\sqrt{L}}-(1-C')\frac{6}{f^4}=1-\frac{1}{\sqrt{L}}-6f^{21}$.

Theorem 5.11. *Let* $m=n^4$ *and* $\varepsilon=\Theta(f^7)$. *Assume that Assumption 1 holds and* $\lambda_0(H)\le e_0$. *With probability* $1-e^{-O(m^{1/3})}$ *over* G, *the following holds: the protocol Algorithm 1 accepts a unique vector with probability* $p=1-O(f^{14})$ *and accepts any orthogonal vector with probability* $p-\Omega(f^{20})+\varepsilon_{\mathrm{ETH}}+1/\mathrm{poly}(n)$.

Proof. Lets write $M=\lambda_0|\Psi\rangle\langle\Psi|+(\lambda_0-\Delta)N$, for some matrix N orthogonal to $|\Psi\rangle$ with norm $\|N\|\le 1$. We have

$$\langle\Psi|\,(Q\otimes Q)M(Q\otimes Q)\,|\Psi\rangle=\lambda_0|\,\langle\Psi|\,Q\otimes Q\,|\Psi\rangle\,|^2+(\lambda_0-\Delta)\,\langle\Psi|\,Q\otimes QNQ\otimes Q\,|\Psi\rangle\,. \tag{40}$$

To bound the second term, note that

$$\begin{aligned}|\,\langle\Psi|\,(Q\otimes Q)N(Q\otimes Q)\,|\Psi\rangle\,| &= |\,\langle\Psi|\,N\,|\Psi\rangle+\langle\Psi|\,(\mathbb{I}-Q\otimes Q)N\,|\Psi\rangle+\langle\Psi|\,(Q\otimes Q)N(\mathbb{I}-Q\otimes Q)\,|\Psi\rangle\,|\\ &\le \langle\Psi|\,N\,|\Psi\rangle+|\,\langle\Psi|\,(\mathbb{I}-Q\otimes Q)N\,|\Psi\rangle\,|+|\,\langle\Psi|\,(Q\otimes Q)N(\mathbb{I}-Q\otimes Q)\,|\Psi\rangle\,|\\ &\le |\,\langle\Psi|\,(\mathbb{I}-Q\otimes Q)N\,|\Psi\rangle\,|+|\,\langle\Psi|\,(Q\otimes Q)N(\mathbb{I}-Q\otimes Q)\,|\Psi\rangle\,| \quad (\text{Since } N\perp|\psi\rangle)\\ &\le 2\|(\mathbb{I}-Q\otimes Q)\,|\psi\rangle\,\|_2\\ &\le 2\,\langle\Psi|\,(\mathbb{I}-Q\otimes Q)^2\,|\Psi\rangle\\ &\le 2\,\langle\Psi|\,(\mathbb{I}-Q\otimes Q)\,|\Psi\rangle \qquad (\text{Using that } Q\otimes Q\preceq\mathbb{I})\\ &\le 2\left(\frac{1}{\sqrt{L}}+6f^{21}\right) \qquad (\text{By Claim 5.10})\\ &\le \frac{2}{\sqrt{L}}+12f^{21}. \end{aligned}$$

Since $\lambda_0-\Delta\le 1$, combining with Equation (40) yields

$$\langle\Psi|\,(Q\otimes Q)M(Q\otimes Q)\,|\Psi\rangle\ge\lambda_0-12f^{21}-\frac{2}{\sqrt{L}}=\widetilde{\lambda}_0. \tag{41}$$

Since $M \preceq \lambda_0 \mathrm{I}$ and $Q \preceq \mathrm{I}$, we also have that the largest eigenvalue of $(Q \otimes Q)M(Q \otimes Q)$ is at most λ_0. Thus, its largest eigenvalue is in the interval $[\widetilde{\lambda}_0, \lambda_0]$. To show unique largest eigenvector and spectral gap, consider any eigenvector vector $|\phi\rangle$ of $(Q \otimes Q)M(Q \otimes Q)$ such that $\langle\phi| (Q \otimes Q)M(Q \otimes Q) |\phi\rangle \geq \lambda_0 - \delta$, for $\delta = \frac{\Delta}{100}$ (where Δ is the spectral gap of M). For such a $|\phi\rangle$, we have

$$
\begin{aligned}
&\langle\phi| (Q \otimes Q)M(Q \otimes Q) |\phi\rangle \geq \lambda_0 - \delta \\
&\implies \lambda_0 |\langle\phi| Q \otimes Q |\Psi\rangle|^2 + (\lambda_0 - \Delta) \langle\phi| Q \otimes QNQ \otimes Q |\phi\rangle && \geq \lambda_0 - \delta \\
&\implies \lambda_0 |\langle\phi| Q \otimes Q |\Psi\rangle|^2 + (\lambda_0 - \Delta) \langle\phi| Q \otimes Q(\mathrm{I} - |\Psi\rangle\langle\Psi|)Q \otimes Q |\phi\rangle && \geq \lambda_0 - \delta && (N \preceq \mathrm{I} - |\Psi\rangle\langle\Psi|) \\
&\implies \Delta |\langle\phi| Q \otimes Q |\Psi\rangle|^2 + (\lambda_0 - \Delta) \| Q \otimes Q |\phi\rangle \|^2 && \geq \lambda_0 - \delta && (\text{Regrouping}) \\
&\implies \Delta |\langle\phi| Q \otimes Q |\Psi\rangle|^2 + (\lambda_0 - \Delta) && \geq \lambda_0 - \delta && (\|Q\| \leq 1) \\
&\implies |\langle\phi| Q \otimes Q |\Psi\rangle|^2 \geq \frac{\Delta - \delta}{\Delta} = 0.99.
\end{aligned}
$$

Using Claim 5.10, we find that $|\phi\rangle$ has 0.97 overlap with $|\Psi\rangle$. This forces that there be a unique eigenvector above $\lambda_0 - \delta$ (two eigenvectors would not have 0.97 overlap with the same vector). This eigenvector must have eigenvalue at least $\widetilde{\lambda}_0$ by Equation (41). Thus the spectral gap of $(Q \otimes Q)M(Q \otimes Q)$ is at least

$$\widetilde{\lambda}_0 - (\lambda_0 - \delta) = \delta - 12f^{21} - \frac{2}{\sqrt{L}} = \frac{\varepsilon^2 f^4 \lambda}{100} - 12f^{21} - \frac{2}{\sqrt{L}} \geq \frac{\varepsilon^2 f^6}{100} - 12f^{21} - \frac{2}{\sqrt{L}}.$$

From Proposition 5.7, we thus find that a unique vector is accepted by the protocol with probability

$$\frac{\varepsilon^2 f^6}{200} - 2\varepsilon^2 f^7 - 12f^{21} - O(\varepsilon^3) - O(m^{-1/3} + \frac{1}{\sqrt{L}})$$

higher than other vectors orthogonal to this vector. Setting $\varepsilon = \Theta(f^7)$, $m = n^4 = poly(n)$, $L = poly(n)$, and using $C' \geq 1 - f^{25}$, this gives a success probability gap of $\Omega(f^{20}) - \frac{1}{\mathrm{poly}(n)}$. Recalling from Claim 5.9 that $\lambda_0 = 1 - \varepsilon^2 + \varepsilon^2 \mathbb{E}_i(\mu^i)^2 + \varepsilon^2\lambda \geq 1 - \varepsilon^2$, we find that the unique vector is accepted with probability

$$\lambda_0 - 12f^{21} - \frac{2}{\sqrt{L}} \geq 1 - O(f^{14}).$$

Finally, we add the potential error $\varepsilon_{\mathrm{ETH}}$ from Hypothesis 5.3. This proves the theorem. □

Since the above procedure yields a promise gap of $\Omega(f^{20}) - \varepsilon_{\mathrm{ETH}} - \frac{1}{\mathrm{poly}(n)}$ (assuming $\varepsilon_{\mathrm{ETH}} \ll f^{20}$), we can amplify the gap to be any constant via the in-place amplification technique of [MW05]. Not only is the promise gap improved, but this amplification scheme is also spectral gap preserving as shown in [JKKS+11].

Corollary 5.12. *For any choice of constants $1 > a > b > 0$, the in-place amplification of [MW05] applied to Algorithm 1 yields a protocol* `Alg` *which accepts a unique vector with probability $\geq a$ and accepts all orthogonal states with probability $\leq b$.*

Finally, Theorem 5.5 follows by Lemma 5.6 and Corollary 5.12.

5.3 Connections between UniqueQMA vs QMA and thermalization

Finally, we discuss implications and potential interpretations of our main results. Our Theorem 5.5 implies that local Hamiltonians satisfying ETH (as formalized in Assumption 1) at a slightly subextensive energy, e.g. $e_0 = \|H\|^{0.999}$, is in UniqueQMA. However, deciding whether the ground energy is $\leq a\|H\|^{1-\alpha}$ or $\geq b\|H\|^{1-\alpha}$

is QMA-hard for constants $a < b$, and any constant $\alpha > 0$ [Che20]. Put differently, the non-equivalence of UniqueQMA and QMA would imply that QMA-hard local Hamiltonians must violate ETH at near-extensive energy regimes. It is worth noting that this violation is *adversarially robust*, meaning that the perturbed Hamitlonian $H + \delta H$ continue to violate ETH for arbitrary perturbation δV of near-extensive strength, e.g. $\|\delta H\| < \|H\|^{0.999}$.

The previous paragraph can be extended to extensive energy scales under the quantum PCP conjecture.

Conjecture 5.13 (Quantum PCP conjecture, [AAV13])**.** *There are constants $c < d$ such that deciding if the ground energy of a general n-qubit local Hamiltonian H is (yes case) $\leq c\|H\|$ or (no case) $\geq d\|H\|$ is* QMA*-complete.*

Let H be a "yes" case local Hamiltonian witnessing the QMA-hardness as in the quantum PCP conjecture. We can now argue that H violates ETH and robustly so against adversarial extensive perturbations, assuming UniqueQMA and QMA are not equivalent. First, we have the following claim:

Claim 5.14 (ETH Hamiltonians in UniqueQMA)**.** *Fix constants $a < b$. Let H be a local Hamiltonian whose ground energy is either $\leq a\|H\|$ or $\geq b\|H\|$. Furthermore, in the "yes" case, it is promised that H satisfies ETH (as formalized in Assumption 1) for an energy value $e_0 < a\|H\|$. Then there is a* UniqueQMA *protocol that decides if the ground energy is $\leq a\|H\|$ or $\geq b\|H\|$.*

Proof. Take the verifier to be Algorithm 1. The fact that this algorithm corresponds to a UniqueQMA verifier follows directly from Theorem 5.5. □

Let $\delta = \frac{d-c}{4}$ and consider any local Hamiltonian V with $\|V\| \leq \|H\|$. Suppose UniqueQMA $\neq$ QMA, we will show that $H + \delta V$ violates ETH at all energy values $e_0 < c\|H\|$. Suppose for contradiction that $H + \delta V$ satisfies ETH (as formalized in Assumption 1). If the ground energy of H is $\leq c\|H\|$, then the ground energy of $H + \delta V$ is $\leq (c + \delta)\|H\|$ and if the ground energy of H is $\geq d\|H\|$, then the ground energy of $H + \delta V$ is $\geq (d - \delta)\|H\|$. Then we invoke Theorem Claim 5.14 with constants $(c + \delta) < d - \delta$ and get a UniqueQMA protocol to decide which is the case. But the question of determining whether the ground energy of $H + \delta V$ is $\leq (c + \delta)\|H\|$ or $\geq (d - \delta)\|H\|$ is QMA-hard because this tells us if the ground energy of H is $\leq c\|H\|$ or $\geq d\|H\|$. This violates the conjecture that UniqueQMA $\neq$ QMA.

The intuition that qPCP Hamiltonians violate an assumption like ETH is consistent with the widely held belief that any proof of the qPCP conjecture will require extensive use of quantum codes. Such codes and their corresponding code Hamiltonians presumably have a lot of algebraic structure that might prevent thermalization.

Remark 5.15. The results and discussions in this section can be adapted to Hamiltonians satisfying a 'weak' ETH. Instead of assuming a known good energy window like in Assumption 1, a weak ETH Hamiltonian is such that a randomly chosen energy window satisfies Hypothesis 5.3 with high probability. In that case, applying Algorithm 1 to a randomly chosen energy window we find that weak ETH Hamiltonians are in $\mathsf{BPP}^{\mathsf{UniqueQMA}}$. As a result, the non-containment $\mathsf{QMA} \not\subseteq \mathsf{BQP}^{\mathsf{UniqueQMA}}$ would imply QMA-hard Hamiltonians violate weak ETH.

A similar observation can be made about many-body localization (MBL) of QMA-hard Hamiltonians. Many-body localization [NH15; AABS19] is a phenomenon that is experimentally and numerically observed in quantum systems that violate ETH. Some common characterizations of MBL include [HH19; HNO14] slow growth of entanglement in Hamiltonian dynamics, area law for most eigenstates, and quasi-local integrals of motion. Defining and justifying these features is not the scope of this work, but from a complexity-theoretic viewpoint it is perhaps not surprising that a Hamiltonian with the above features are not expected to be QMA-hard. As an example, the quasi-local integrals of motion property roughly states that there is a quasi-local unitary U that diagonalizes H. A quasilocal unitary is a quantum circuit of depth $D = O(\log n)$ which can consist of non-local gates [18]. However, the constraint is that if any gate g acts on k sites, then $\|g - \mathrm{I}\| = e^{-\Omega(k)}$. Thus, we can replace the unitary U with another unitary U', in which any gate acting

[18]The choice $D = O(\log n)$ is made since we are not restricting the geometry, In 1D systems, the depth would be $O(n)$ [EDBAG22]. Regardless, the argument actually stays intact for any $D = \mathrm{poly}(n)$.

on more than $O(\log n)$ qubits is set to identity. The error incurred in this process is $\frac{1}{\text{poly}(n)}$ and we can use the Solovay-Kitaev theorem to compile U' into a polynomial sized circuit. This implies MBL Hamiltonians are in QCMA: we expect a prover to provide us with the description of U and the ground state of the diagonalized Hamiltonian as a bitstring, and then run the circuit to verify the energy. Hence, assuming QCMA $\neq$ QMA [AK07], QMA-hard Hamiltonians are not MBL. Similar to the ETH discussion, this fact can be made robust against adversarial extensive perturbations under the quantum PCP conjecture as follows. Suppose there is a V such that $H + \delta V$ satisfies MBL. Then there is a polynomial-sized circuit U such that $U^\dagger(H + \delta V)U$ diagonal and a classical string $|x\rangle$ that is the ground state in this basis. The prover sends the classical descriptions of V, U, x and then the verifier runs a quantum computer to prepare the ground state and decide if the ground energy of $H + \delta V$ is $\leq (c + \delta)\|H\|$ or $\geq (d - \delta)\|H\|$. This in turn allows to decide if the ground energy of H is $\leq c\|H\|$ or $\geq d\|H\|$, giving a QCMA protocol for the original H. Therefore, under the assumption QCMA $\neq$ QMA, MBL Hamiltonians cannot be expected to be QMA-hard.

The connection to UniqueQMA is more subtle. It is shown in [ABOBS22] that UniqueQCMA is contained in QCMA under randomized reductions, thus QCMA $\subseteq$ BQP$^{\text{UniqueQMA}}$ and the above discussion of MBL also holds under the assumption QMA $\not\subseteq$ BQP$^{\text{UniqueQMA}}$. Then, QMA-hard Hamiltonians are robustly neither ETH or MBL even up to adversarial perturbations of near-extensive strengths (extensive if quantum PCP conjecture is true as shown above).

Finally, we remark that the adversarially robust violation of ETH and MBL described above is impossible on lattice Hamiltonians. Observe that on a D-dimensional lattice of length L, there are local perturbations of total strength $O(\varepsilon^D L^D)$ that disentangle the Hamiltonian into non-interacting D-dimensional cubical patches of size $\frac{1}{\varepsilon} \times \ldots \times \frac{1}{\varepsilon}$. Such non-interacting Hamiltonian is trivially localized.

We find it interesting that a seemingly unrelated complexity theory assumption can lead to such a conclusion, especially that the question of "whether there is only one phase transition, or could there possibly be some sort of intermediate phase that is neither fully localized nor fully thermal?" [NH15] is still a central and active research question in quantum condensed matter physics.

References

[AABS19] Dmitry A Abanin, Ehud Altman, Immanuel Bloch, and Maksym Serbyn. "Colloquium: Many-body localization, thermalization, and entanglement". In: *Reviews of Modern Physics* 91.2 (2019), p. 021001.

[Aar08] Scott Aaronson. "On perfect completeness for QMA". In: *arXiv preprint arXiv:0806.0450* (2008).

[AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. "Guest column: the quantum PCP conjecture". In: *Acm sigact news* 44.2 (2013), pp. 47–79.

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.

[ABD24] Avantika Agarwal and Shalev Ben-David. "Oracle Separations for the Quantum-Classical Polynomial Hierarchy". In: *arXiv preprint arXiv:2410.19062* (2024).

[ABOBS22] Dorit Aharonov, Michael Ben-Or, Fernando G. S. L. Brandão, and Or Sattath. "The Pursuit of Uniqueness: Extending Valiant-Vazirani Theorem to the Probabilistic and Quantum Settings". In: *Quantum* 6 (Mar. 17, 2022), p. 668. DOI: 10.22331/q-2022-03-17-668. URL: https://quantum-journal.org/papers/q-2022-03-17-668/ (visited on 10/26/2023).

[ACKK24] Eric R Anshuetz, Chi-Fang Chen, Bobak T Kiani, and Robbie King. "Strongly interacting fermions are non-trivial yet non-glassy". In: *arXiv preprint arXiv:2408.15699* (2024).

[AGIK09] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. "The Power of Quantum Systems on a Line". In: *Communications in Mathematical Physics* 287.1 (2009), pp. 41–65. ISSN: 1432-0916. DOI: 10.1007/s00220-008-0710-3. URL: https://doi.org/10.1007/s00220-008-0710-3.

[AHLN+14] Dorit Aharonov, Aram W. Harrow, Zeph Landau, Daniel Nagaj, Mario Szegedy, and Umesh Vazirani. "Local Tests of Global Entanglement and a Counterexample to the Generalized Area Law". In: *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. 2014 IEEE 55th Annual Symposium on Foundations of Computer Science. Oct. 2014, pp. 246–255. DOI: 10.1109/FOCS.2014.34. URL: https://ieeexplore.ieee.org/document/6979009 (visited on 11/29/2023).

[AIK22] Scott Aaronson, DeVon Ingram, and William Kretschmer. "The Acrobatics of BQP". In: *37th Computational Complexity Conference*. 2022.

[AK07] Scott Aaronson and Greg Kuperberg. "Quantum versus Classical Proofs and Advice". In: *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*. Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07). June 2007, pp. 115–128. DOI: 10.1109/CCC.2007.27. URL: https://ieeexplore.ieee.org/abstract/document/4262757 (visited on 06/19/2024).

[AK25] Avantika Agarwal and Srijita Kundu. "A Cautionary Note on Quantum Oracles". In: *arXiv preprint arXiv:2504.19470* (2025).

[AKKT20] Scott Aaronson, Robin Kothari, William Kretschmer, and Justin Thaler. "Quantum Lower Bounds for Approximate Counting via Laurent Polynomials". In: *LIPIcs, Volume 169, CCC 2020* 169 (2020), 7:1–7:47. ISSN: 1868-8969. DOI: 10.4230/LIPIcs.CCC.2020.7. arXiv: 1904.08914 [quant-ph]. URL: http://arxiv.org/abs/1904.08914 (visited on 06/18/2024).

[AKL16] Itai Arad, Tomotaka Kuwahara, and Zeph Landau. "Connecting Global and Local Energy Distributions in Quantum Spin Models on a Lattice". In: *Journal of Statistical Mechanics: Theory and Experiment* 2016.3 (Mar. 2016), p. 033301. ISSN: 1742-5468. DOI: 10.1088/1742-5468/2016/03/033301. URL: https://dx.doi.org/10.1088/1742-5468/2016/03/033301 (visited on 01/03/2024).

[AKLT04] Ian Affleck, Tom Kennedy, Elliott H Lieb, and Hal Tasaki. "Rigorous results on valence-bond ground states in antiferromagnets". In: *Condensed Matter Physics and Exactly Soluble Models: Selecta of Elliott H. Lieb* (2004), pp. 249–252.

[ALVV17] Itai Arad, Zeph Landau, Umesh Vazirani, and Thomas Vidick. "Rigorous RG algorithms and area laws for low energy eigenstates in 1D". In: *Communications in Mathematical Physics* 356 (2017), pp. 65–105.

[BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. "Strengths and Weaknesses of Quantum Computing". In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1510–1523. ISSN: 0097-5397, 1095-7111. DOI: 10.1137/S0097539796300933. arXiv: quant-ph/9701001. URL: http://arxiv.org/abs/quant-ph/9701001 (visited on 06/06/2024).

[BBCMDW01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald De Wolf. "Quantum lower bounds by polynomials". In: *Journal of the ACM (JACM)* 48.4 (2001), pp. 778–797.

[BBF98] Richard Beigel, Harry Buhrman, and Lance Fortnow. "NP might not be as easy as detecting unique solutions". In: *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. 1998, pp. 203–208.

[BCGHZ24] Sergey Bravyi, Anirban Chowdhury, David Gosset, Vojtech Havlicek, and Guanyu Zhu. "Quantum complexity of the Kronecker coefficients". In: *PRX Quantum* 5.1 (2024), p. 010329.

[BCGW22] Sergey Bravyi, Anirban Chowdhury, David Gosset, and Pawel Wocjan. "On the Complexity of Quantum Partition Functions". In: *Nature Physics* 18.11 (Nov. 2022), pp. 1367–1370. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/s41567-022-01742-5. arXiv: 2110.15466 [quant-ph]. URL: http://arxiv.org/abs/2110.15466 (visited on 05/12/2024).

[BDCG89] Shai Ben-David, Benny Chor, and Oded Goldreich. "On the theory of average case complexity". In: *Proceedings of the twenty-first annual ACM symposium on Theory of computing*. 1989, pp. 204–216.

[BFS11] Brielin Brown, Steven T Flammia, and Norbert Schuch. "Computational difficulty of computing the density of states". In: *Physical review letters* 107.4 (2011), p. 040501.

[BS08] Salman Beigi and Peter W. Shor. *On the Complexity of Computing Zero-Error and Holevo Capacity of Quantum Channels*. 2008. arXiv: 0709.2090 [quant-ph]. URL: https://arxiv.org/abs/0709.2090.

[CB23] Chi-Fang Chen and Fernando G. S. L. Brandão. *Fast Thermalization from the Eigenstate Thermalization Hypothesis*. Mar. 29, 2023. DOI: 10.48550/arXiv.2112.07646. arXiv: 2112.07646 [cond-mat, physics:math-ph, physics:quant-ph]. URL: http://arxiv.org/abs/2112.07646 (visited on 01/14/2024). preprint.

[CFHL19] William Cottrell, Ben Freivogel, Diego M. Hofman, and Sagar F. Lokhande. "How to build the thermofield double state". In: *Journal of High Energy Physics* 2019.2 (2019), p. 58. ISSN: 1029-8479. DOI: 10.1007/JHEP02(2019)058. URL: https://doi.org/10.1007/JHEP02(2019)058.

[Che20] Lijie Chen. "Unpublished work". 2020.

[Deu18] Joshua M Deutsch. "Eigenstate thermalization hypothesis". In: *Reports on Progress in Physics* 81.8 (2018), p. 082001.

[Deu91] Josh M Deutsch. "Quantum statistical mechanics in a closed system". In: *Physical review a* 43.4 (1991), p. 2046.

[DGF22] Abhinav Deshpande, Alexey V Gorshkov, and Bill Fefferman. "Importance of the Spectral gap in Estimating Ground-State Energies". In: *PRX Quantum* 3.4 (2022), p. 040327.

[DKPR16] Luca D'Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol. "From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics". In: *Advances in Physics* 65.3 (2016), pp. 239–362.

[DKVMW13] Holger Dell, Valentine Kabanets, Dieter Van Melkebeek, and Osamu Watanabe. "Is Valiant–Vazirani's isolation probability improvable?" In: *computational complexity* 22.2 (2013), pp. 345–383.

[EDBAG22] Adam Ehrenberg, Abhinav Deshpande, Christopher L Baldwin, Dmitry A Abanin, and Alexey V Gorshkov. "Simulation Complexity of Many-Body Localized Systems". In: *arXiv preprint arXiv:2205.12967* (2022).

[ES14] László Erdős and Dominik Schröder. "Phase transition in the density of states of quantum spin glasses". In: *Mathematical Physics, Analysis and Geometry* 17.3 (2014), pp. 441–464.

[FL16] Bill Fefferman and Cedric Yen-Yu Lin. "A complete characterization of unitary quantum space". In: *arXiv preprint arXiv:1604.01384* (2016).

[fSV02] given-i=AYu family=Kitaev given=A. Yu., A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. USA: American Mathematical Society, June 2002. 257 pp. ISBN: 978-0-8218-3229-5.

[GE07] David Gross and Jens Eisert. "Quantum margulis expanders". In: *arXiv preprint arXiv:0710.0651* (2007).

[GS23] Friedrich Götze and Holger Sambale. "Higher order concentration on Stiefel and Grassmann manifolds". In: *Electronic Journal of Probability* 28 (2023), pp. 1–30.

[GS88] Joachim Grollmann and Alan L Selman. "Complexity measures for public-key cryptosystems". In: *SIAM Journal on Computing* 17.2 (1988), pp. 309–335.

[Har13] Aram W. Harrow. *The Church of the Symmetric Subspace*. Aug. 29, 2013. DOI: 10.48550/arXiv.1308.6595. arXiv: 1308.6595 [quant-ph]. URL: http://arxiv.org/abs/1308.6595 (visited on 10/27/2023). preprint.

[Has07] Matthew B Hastings. "Random unitaries give quantum expanders". In: *Physical Review A—Atomic, Molecular, and Optical Physics* 76.3 (2007), p. 032315.

[HE23] Dominik Hangleiter and Jens Eisert. "Computational advantage of quantum random sampling". In: *Reviews of Modern Physics* 95.3 (2023), p. 035001.

[HH19] Yichen Huang and Aram W Harrow. "Instability of localization in translation-invariant systems". In: *arXiv preprint arXiv:1907.13392* (2019).

[HNO14] David A Huse, Rahul Nandkishore, and Vadim Oganesyan. "Phenomenology of fully many-body-localized systems". In: *Physical Review B* 90.17 (2014), p. 174202.

[HNOS96] Lane A Hemaspaandra, Ashish V Naik, Mitsunori Ogihara, and Alan L Selman. "Computing solutions uniquely collapses the polynomial hierarchy". In: *SIAM Journal on Computing* 25.4 (1996), pp. 697–708.

[HO22] Matthew B Hastings and Ryan O'Donnell. "Optimizing strongly interacting fermionic Hamiltonians". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 776–789.

[INNRY21] Sandy Irani, Anand Natarajan, Chinmay Nirkhe, Sujit Rao, and Henry Yuen. "Quantum search-to-decision reductions and the state synthesis problem". In: *arXiv preprint arXiv:2111.02999* (2021).

[JKKS+11] Rahul Jain, Iordanis Kerenidis, Greg Kuperberg, Miklos Santha, Or Sattath, and Shengyu Zhang. *On the Power of a Unique Quantum Witness*. Jan. 18, 2011. DOI: 10.48550/arXiv.0906.4425. arXiv: 0906.4425 [quant-ph]. URL: http://arxiv.org/abs/0906.4425 (visited on 05/01/2024). preprint.

[JVV86] Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani. "Random generation of combinatorial structures from a uniform distribution". In: *Theoretical computer science* 43 (1986), pp. 169–188.

[Kit15a] Alexei Kitaev. "A simple model of quantum holography (part 1)". In: *Entanglement in strongly-correlated quantum matter* (2015), p. 38.

[Kit15b] Alexei Kitaev. "A simple model of quantum holography (part 2)". In: *Entanglement in strongly-correlated quantum matter* (2015), p. 38.

[KN97] Tohru Koma and Bruno Nachtergaele. "The spectral gap of the ferromagnetic XXZ-chain". In: *Letters in Mathematical Physics* 40.1 (1997), pp. 1–16.

[Kup09] Greg Kuperberg. "How hard is it to approximate the Jones polynomial?" In: *arXiv preprint arXiv:0908.0512* (2009).

[KVM99] Adam R Klivans and Dieter Van Melkebeek. "Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses". In: *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*. 1999, pp. 659–667.

[LVV15] Zeph Landau, Umesh Vazirani, and Thomas Vidick. "A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians". In: *Nature Physics* 11.7 (2015), pp. 566–569.

[Mon21] Andrea Montanari. "Optimization of the Sherrington–Kirkpatrick Hamiltonian". In: *SIAM Journal on Computing* 0 (2021), FOCS19–1.

[MS14] Ramis Movassagh and Peter W Shor. "Power law violation of the area law in quantum spin chains". In: *arXiv preprint arXiv:1408.1657* (2014).

[MS97] VD Milman and G Schechtman. "Global vs. local asymptotic theories of finite dimensional normed spaces". In: *Duke Math. J* 90 (1997), pp. 73–93.

[MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. "Matching is as easy as matrix inversion". In: *Combinatorica* 7.1 (1987), pp. 105–113. ISSN: 1439-6912. DOI: 10.1007/BF02579206. URL: https://doi.org/10.1007/BF02579206.

[MW05] Chris Marriott and John Watrous. *Quantum Arthur-Merlin Games*. June 15, 2005. DOI: 10.48550/arXiv.cs/0506068. arXiv: cs/0506068. URL: http://arxiv.org/abs/cs/0506068 (visited on 02/29/2024). preprint.

[NH15] Rahul Nandkishore and David A Huse. "Many-body localization and thermalization in quantum statistical mechanics". In: *Annu. Rev. Condens. Matter Phys.* 6.1 (2015), pp. 15–38.

[NN24] Anand Natarajan and Chinmay Nirkhe. "A distribution testing oracle separation between QMA and QCMA". In: *Quantum* 8 (2024), p. 1377.

[NV16] Anand Natarajan and Thomas Vidick. "Robust self-testing of many-qubit states". In: *arXiv preprint arXiv:1610.03574* (2016).

[OOLS+] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain. *NIST Digital Library of Mathematical Functions*. https://dlmf.nist.gov/, Release 1.2.0 of 2024-03-15. URL: https://dlmf.nist.gov/.

[Orú14] Román Orús. "A practical introduction to tensor networks: Matrix product states and projected entangled pair states". In: *Annals of physics* 349 (2014), pp. 117–158.

[RDO08] Marcos Rigol, Vanja Dunjko, and Maxim Olshanii. "Thermalization and its mechanism for generic isolated quantum systems". In: *Nature* 452.7189 (2008), pp. 854–858.

[SBE17] Martin Schwarz, Olivier Buerschaper, and Jens Eisert. "Approximating local observables on projected entangled pair states". In: *Physical Review A* 95.6 (2017), p. 060102.

[SK75] David Sherrington and Scott Kirkpatrick. "Solvable model of a spin-glass". In: *Physical review letters* 35.26 (1975), p. 1792.

[SM23] Oles Shtanko and Ramis Movassagh. *Preparing thermal states on noiseless and noisy programmable quantum processors*. 2023. arXiv: 2112.14688 [quant-ph].

[Sre99] Mark Srednicki. "The approach to thermal equilibrium in quantized chaotic systems". In: *Journal of Physics A: Mathematical and General* 32.7 (1999), p. 1163.

[Sto83] Larry Stockmeyer. "The complexity of approximate counting". In: *Proceedings of the fifteenth annual ACM symposium on Theory of computing*. 1983, pp. 118–126.

[STV12] Martin Schwarz, Kristan Temme, and Frank Verstraete. "Preparing projected entangled pair states on a quantum computer". In: *Physical review letters* 108.11 (2012), p. 110502.

[SV17] Julian Sonner and Manuel Vielma. "Eigenstate thermalization in the Sachdev-Ye-Kitaev model". In: *Journal of High Energy Physics* 2017.11 (2017), p. 149. ISSN: 1029-8479. DOI: 10.1007/JHEP11(2017)149. URL: https://doi.org/10.1007/JHEP11(2017)149.

[SWVC07] Norbert Schuch, Michael M. Wolf, Frank Verstraete, and J. Ignacio Cirac. "Computational Complexity of Projected Entangled Pair States". In: *Phys. Rev. Lett.* 98 (14 2007), p. 140506. DOI: 10.1103/PhysRevLett.98.140506. URL: https://link.aps.org/doi/10.1103/PhysRevLett.98.140506.

[SY23] Adrian She and Henry Yuen. "Unitary Property Testing Lower Bounds by Polynomials". In: *Leibniz International Proceedings in Informatics (LIPIcs): 14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. 2023.

[SY93] Subir Sachdev and Jinwu Ye. "Gapless spin-fluid ground state in a random quantum Heisenberg magnet". In: *Physical review letters* 70.21 (1993), p. 3339.

[SZ09] Yaoyun Shi and Shengyu Zhang. "Note on quantum counting classes". In: *URL: http://www. cse. cuhk. edu. hk/˜ syzhang/papers/SharpBQP. pdf* (2009).

[Tao23] Terence Tao. *Topics in random matrix theory*. Vol. 132. American Mathematical Society, 2023.

[Tod91] Seinosuke Toda. "PP is as hard as the polynomial-time hierarchy". In: *SIAM Journal on Computing* 20.5 (1991), pp. 865–877.

[Tro15] Joel A. Tropp. *An Introduction to Matrix Concentration Inequalities*. 2015. arXiv: 1501.01571 [math.PR].

[Ver10] Roman Vershynin. "Introduction to the non-asymptotic analysis of random matrices". In: *arXiv preprint arXiv:1011.3027* (2010).

[VV86] L. G. Valiant and V. V. Vazirani. "NP Is as Easy as Detecting Unique Solutions". In: *Theoretical Computer Science* 47 (Jan. 1, 1986), pp. 85–93. ISSN: 0304-3975. DOI: 10.1016/0304-3975(86)90135-0. URL: https://www.sciencedirect.com/science/article/pii/0304397586901350 (visited on 10/26/2023).

[Zha24] Mark Zhandry. *Toward Separating QMA from QCMA with a Classical Oracle*. arXiv:2411.01718 [quant-ph]. Nov. 2024. DOI: 10.48550/arXiv.2411.01718. URL: http://arxiv.org/abs/2411.01718 (visited on 06/11/2025).

A Miscellaneous proofs

A.1 Probabilistic tools

In this section, we use **bold face** to distinguish random variables from constants.

Fact A.1 (Bernstein's inequality)**.** *Let $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_t$ be independent random variables with $\mathbb{E}\,\boldsymbol{x}_i = 0$, satisfying*

1. *$\sum_i \mathbb{E}\,\boldsymbol{x}_i^2 = v$*
2. *For all $k > 2$, $\mathbb{E}\,\boldsymbol{x}_i^k \leq \frac{\mathbb{E}\,\boldsymbol{x}_i^2}{2} k! L^{k-2}$*

Then,

$$\Pr\left[\sum X_i \geq \varepsilon\right] \leq exp\left[\frac{-\varepsilon^2/2}{v + L\varepsilon/3}\right]$$

Fact A.2 (Distribution of squares of complex Gaussians)**.** *Let $g \sim \mathcal{CN}(0,1)$. Then, $|\boldsymbol{g}|^2 \equiv \frac{1}{2}\boldsymbol{h}$, where $\boldsymbol{h}$ is drawn from a Chi-squared distribution with 2 degrees of freedom, i.e. $\boldsymbol{h} \sim \chi^2(2)$.*

Proof. Writing $\boldsymbol{g} = \boldsymbol{a} + \iota\boldsymbol{b}$ with $\boldsymbol{a}, \boldsymbol{b} \sim \mathcal{N}(0, 1/2)$, we have that $|\boldsymbol{g}|^2 = \boldsymbol{a}^2 + \boldsymbol{b}^2$, i.e. the sum of the squares of two Gaussians with variance $\frac{1}{2}$. Renormalizing such that each of their variances is 1 yields the claim. □

Fact A.3 (Linear combination of complex Gaussians)**.** *Let $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_n$ be a collection of complex Gaussians, $\boldsymbol{g}_i \sim \mathcal{CN}(0,1)$. Then, for any $u \in \mathbb{C}^n$,*

$$\widetilde{\boldsymbol{g}} \triangleq \sum_i u_i \boldsymbol{g}_i \sim \mathcal{CN}(0, \|u\|_2^2)$$

Proof. Each $\boldsymbol{g}_i \equiv \boldsymbol{a}_i + \iota\boldsymbol{b}_i$ with $\boldsymbol{a}_i, \boldsymbol{b}_i \sim \mathcal{N}(0, 1/2)$. The claim will follow by showing that,

1. If $\boldsymbol{g} \sim \mathcal{CN}(0,1)$ then $u\boldsymbol{g} \sim \mathcal{CN}(0, |u|^2)$
2. If $\boldsymbol{g}, \boldsymbol{h} \sim \mathcal{CN}(0,1)$ then $\boldsymbol{g} + \boldsymbol{h} \sim \mathcal{CN}(0, 2)$.

The claim follows easily then, since each $u_i\boldsymbol{g}_i \sim \mathcal{CN}(0, |u_i|^2)$ and $\sum_i u_i\boldsymbol{g}_i \sim \mathcal{CN}(0, \sum_i |u_i|^2) = \mathcal{CN}(0, \|u_i\|_2^2)$.

First we show the first property. Let $u = v + \iota w$, we have that

$$\begin{aligned} u\boldsymbol{g} &= (v\boldsymbol{a} - w\boldsymbol{b}) + \iota(w\boldsymbol{a} + v\boldsymbol{b}) \\ &= \boldsymbol{a}' + \iota\boldsymbol{b}' \end{aligned}$$

Clearly $\boldsymbol{a}'$ and $\boldsymbol{b}'$ are each Gaussians; moreover, they can be verified to be independent. Additionally, $\mathbb{E}\,\boldsymbol{a}' = 0$ and $\mathbb{E}(\boldsymbol{a}')^2 = v^2\,\mathbb{E}\,\boldsymbol{a}^2 + w^2\,\mathbb{E}\,\boldsymbol{b}^2 = \frac{1}{2}|u|$, so $\boldsymbol{a}' \equiv |u|\boldsymbol{a}''$ with $\boldsymbol{a}'' \sim \mathcal{CN}(0, 1/2)$. Similarly $\boldsymbol{b}' \equiv |u|\boldsymbol{b}''$. So,

$$u\boldsymbol{g} = (\boldsymbol{a}' + \iota\boldsymbol{b}') \equiv |u|(\boldsymbol{a}'' + \iota\boldsymbol{b}'') \sim \mathcal{CN}(0, |u|) \tag{42}$$

Now consider the second property. If $\boldsymbol{g} = \boldsymbol{a} + \iota\boldsymbol{b}$ and $\boldsymbol{h} = \boldsymbol{c} + \iota\boldsymbol{d}$, then clearly $\boldsymbol{g} + \boldsymbol{h} = (\boldsymbol{a}+\boldsymbol{c}) + \iota(\boldsymbol{b}+\boldsymbol{d}) =: \boldsymbol{a}' + \iota\boldsymbol{b}'$. Clearly $\boldsymbol{a}', \boldsymbol{b}'$ are independent and each distributed as $\mathcal{N}(0,1)$. Pulling out the scaling factor as in Equation (42) yields that $\boldsymbol{g} + \boldsymbol{h} \sim \mathcal{CN}(0, 2)$. □

Next, we prove Lemma A.4

Lemma A.4. *Let P be a random Hermitian matrix of dimension $D \times D$ with each entry being an independent complex Gaussian with mean 0 and (non-identical) variance ≤ 1. Then with probability $1 - e^{-D}$, we have $\|P\| \leq 10\sqrt{D}$.*

Proof. This proof closely follows the proof of Fact D.1 in [CB23], which in turn is based on argument from Section 2.3 of [Tao23] (with the slightly modification that in [Tao23], they consider bounded random variables). Write $P = \sum_{\alpha,\beta} \boldsymbol{g}_{\alpha,\beta} |v_\alpha\rangle\langle v_\beta|$. Recalling that $\|P\| = \max_{|\psi\rangle} \|P|\psi\rangle\|_2$, we first analyze the concentration we get for a single fixed $|\psi\rangle = \sum_\alpha \psi_\alpha |v_\alpha\rangle$. We have,

$$P|\psi\rangle = \sum_\alpha \langle(\boldsymbol{g}_{\alpha,\beta})_\beta, (\psi_\beta)_\beta\rangle |v_\alpha\rangle =: \sum_\alpha \boldsymbol{P}_\alpha |v_\alpha\rangle$$

By Fact A.3, $\boldsymbol{P}_\alpha = \langle (\boldsymbol{g}_{\alpha,\beta})_\beta, (\psi_\beta)_\beta \rangle \sim \mathcal{CN}(0, \|(\psi_\beta)_\beta\|_2^2)$. Then,

$$\Pr\Big[\|P\,|\psi\rangle\,\|_2 \geq c\sqrt{D}\Big] = \Pr\left[\sum_\alpha |\boldsymbol{P}_\alpha|^2 \geq c^2 D\right]$$

At this point, we'd like to apply Fact A.1 with the random variables $\boldsymbol{x}'_\alpha = |\boldsymbol{P}_\alpha|^2$. By Fact A.2 and using that $\|\,|\psi\rangle\,\|_2 = 1$, each $\boldsymbol{x}'_\alpha$ is distributed as $\frac{1}{2}\|(\psi_\beta)_\beta\|_2^2 \cdot \boldsymbol{y}_\alpha = \frac{1}{2}\boldsymbol{y}_\alpha$ where $\boldsymbol{y}_\alpha \sim \chi^2(2)$. A mild issue is that $\mathbb{E}\,\boldsymbol{x}'_\alpha = 1$, not 0 as required in Bernstein's inequality. Thus, we pick $\boldsymbol{x}_\alpha = \boldsymbol{x}'_\alpha - 1$. We also need to bound $\mathbb{E}\,\boldsymbol{x}_\alpha^k$, compared to $\mathbb{E}\,\boldsymbol{x}_\alpha^2$. We have that,

$$\begin{aligned}
\mathbb{E}\,\boldsymbol{x}_\alpha^k &= \mathbb{E}\left(\frac{1}{2}\boldsymbol{y}_\alpha - 1\right)^k && (\boldsymbol{y}_\alpha \sim \chi^2(2)) \\
&= \sum_{\ell=0}^k \binom{k}{\ell} \frac{1}{2^\ell}(-1)^{k-\ell}\,\mathbb{E}\,\boldsymbol{y}_\alpha^\ell \\
&= \sum_{\ell=0}^k \binom{k}{\ell} \frac{1}{2^\ell}(-1)^{k-\ell}(2^\ell\Gamma(\ell+1)) && \text{Moment of Chi-square Distribution} \\
&\leq \sum_{\ell=0}^k \binom{k}{\ell}(-1)^{k-\ell}\frac{(2\ell)!!}{2^\ell} && (\Gamma(\ell+1) \leq \tfrac{(2k!!)}{2^{k-1}}) \\
&\leq (2k)!! \sum_{\ell=0}^k \binom{k}{\ell}\frac{1}{2^\ell} \\
&\leq (2k)!!2^k \\
&= (2^k k!)2^k = 4^k k!
\end{aligned}$$

Noticing that $\mathbb{E}\,\boldsymbol{x}_\alpha^2 = 1$, it suffices to take $L = 8$ in the statement of Fact A.1. Finally, we have that $\mathbb{E}\sum_\alpha \boldsymbol{x}_\alpha^2 = D$. And thus, we have

$$\begin{aligned}
\Pr\Big[\|P\,|\psi\rangle\,\|_2 \geq c\sqrt{D}\Big] &= \Pr\left[\sum_\alpha |\boldsymbol{P}_\alpha|^2 - 1 \geq (c^2-1)D\right] \\
&= \Pr\left[\sum_\alpha \boldsymbol{x}_\alpha \geq (c^2-1)D\right] \\
&\leq \exp\left[\frac{-\frac{(c^2-1)^2}{2}D^2}{D + \frac{8}{3}(c^2-1)D}\right] && \text{(By Fact A.1)} \\
&\leq e^{-\frac{(c^2-1)}{7}D} && (43)
\end{aligned}$$

This is a bound for a single $|\psi\rangle$. In order to lift this to a bound on the operator norm, we use Lemma 2.3.2 of [Tao23] as well as a standard bound on the size of epsilon-nets.

Lemma A.5 (Lemma 2.3.2 of [Tao23])**.** *Let Σ be a maximal 1/2-net of the sphere S. Then, for any $D \times D$ matrix with complex coefficients and any $\lambda > 0$, we have*

$$\Pr[\|M\| \geq \lambda] \leq \Pr\left[\bigcup_{y\in\Sigma} \|My\|_2 \geq \frac{\lambda}{2}\right]$$

Lemma A.6 (Lemma 5.2 of [Ver10])**.** *Let $0 < \varepsilon < 1$, and let Σ be a 1/2-net of the sphere S. Then Σ has cardinality at most 5^D.*

Now, setting $\lambda = 10\sqrt{D}$ in Lemma A.5 and $c = 5$ in Equation (43), we find that $\Pr\left[\|P\| \geq 10\sqrt{D}\right] \leq e^{-3D}$. Taking the union bound over all states in the epsilon net (of which there are at most $5^D \leq e^{2D}$, we obtain a bound of

$$\Pr\left[\|P\| \geq 10\sqrt{D}\right] \leq e^{-D}$$

□

A.2 Properties of the Q operator

Let $e_0 \in (0,1)$ be some which energy value and $\mathcal{I}_\omega = [e_0 - \omega/2, e_0 + \omega/2]$ be an interval of width ω around e_0. Let $\mathsf{disc}_L(\mathcal{I}_\omega)$ be the discretization of $\mathcal{I}_\omega$ consisting of multiples of $1/L$ contained in $\mathcal{I}_\omega$. Then $L \cdot \mathsf{disc}_L(\mathcal{I}_\omega)$ is the element-wise multiplication of $m \in \mathsf{disc}_L(\mathcal{I}_\omega)$ by L such that set is now a set of consecutive integers. Then, given a Hamiltonian H, $Q(H)$ is defined as

$$Q(H) \triangleq \sum_{m \in L \cdot \mathsf{disc}_L(\mathcal{I}_\omega)} \left(\mathrm{sinc}_L(H - m/L)\right)^2. \tag{44}$$

In this section we'll work often with sums of the form $\sum_x \frac{1}{x^2}$. These can be computed by the following lemma,

Lemma A.7. *It holds that*

$$\sum_{x=a}^{b} \frac{1}{x^2} = \Gamma^{(1)}(a) - \Gamma^{(1)}(b+1),$$

where $\Gamma^{(n)}$ is known as the polygamma *function and is the n^{th} derivative of the gamma function.*

This identity follows from the series expansion of $\Gamma^{(n)}(x)$ which can be found in [OOLS+, (5.15.1)]. Now, we establish basic properties of the Q operator.

Lemma A.8. *Let Q be defined as in Equation (44). Let $\{|v_\alpha\rangle\}_{\alpha \in [N]}$ be the eigenbasis of H. Then,*

1. *Q is diagonal in the basis $\{|v_\alpha\rangle\}_\alpha$,*
2. *(Almost Identity) When $\lambda_\alpha \in [e_0 - \omega/2 + \frac{c}{L}, e_0 + \omega/2 - \frac{c}{L}] \subseteq \mathcal{I}_\omega$, then $\|Q(H)|v_\alpha\rangle\|_2^2 \geq 1 - \frac{1}{c}$,*
3. *(Almost Zero) When $|\lambda_\alpha - e| \geq \frac{c}{L}$ for all $e \in \mathcal{I}_\omega$, then $\|Q(H)|v_\alpha\rangle\|_2^2 \leq \frac{1}{c}$.*

Proof. The first property is clear from the definition of $Q(H)$. To show Item 2, take $\lambda_\alpha \in [e_0 - \omega/2 + \frac{c}{L}, e_0 + \omega/2 - \frac{c}{L}]$. Since $Q(H)$ is diagonal in H's eigenbasis,

$$\begin{aligned} Q(H)|v_\alpha\rangle &= \sum_{m \in L \cdot \mathsf{disc}_L(\mathcal{I}_\omega)} \mathrm{sinc}_L(\lambda_\alpha - m/L)^2 |v_\alpha\rangle \\ \implies \|Q(H)|v_\alpha\rangle\|_2 &= \sum_{m \in L \cdot \mathsf{disc}_L(\mathcal{I}_\omega)} \mathrm{sinc}_L(\lambda_\alpha - m/L)^2 \end{aligned} \tag{45}$$

Write $\lambda_\alpha = \frac{1}{L}\lfloor L \cdot \lambda_\alpha \rceil_L + \delta$, where $\lfloor L \cdot \lambda_\alpha \rceil_L$ is $L \cdot \lambda_\alpha$ rounded to the *nearest* integer and $\delta \in [0, \frac{1}{2L})$ is the remainder[19]. Each summand in Equation (45) is then $\mathrm{sinc}_L(\frac{1}{L}(\lfloor L \cdot \lambda_\alpha \rceil - m) + \delta)$. Since we've picked λ_α to be in the interval $[e_0 - \frac{\omega}{2} + \frac{c}{L}, e_0 + \frac{\omega}{2} - \frac{c}{L}]$, there exists m^* in the sum such that $m^* = \lfloor L \cdot \lambda_\alpha \rceil$. Furthermore, there are at least $c-1$-values of m both above and below m^*.

[19] Rounding to the nearest integer rather than the smaller one is important to ensure the worst case is $\mathrm{sinc}_L(1/2)$.

Thus, we have

$$\begin{aligned}
\sum_{m\in L\cdot \mathsf{disc}_L(\mathcal{I}_\omega)} \mathrm{sinc}_L(\lambda_\alpha - m/L)^2 &\geq \sum_{k=-c+1}^{c-1} \mathrm{sinc}_L(\tfrac{k}{L}+\delta) \\
&= \sum_{k=-c+1}^{c-1} \left(\frac{\sin(\pi(k+L\delta))}{L\sin\big(\pi(\frac{k}{L}+\delta)\big)}\right)^2 \\
&\geq \sum_{k=-c+1}^{c-1} \left(\frac{\sin(\pi(k+L\delta))}{\pi(k+L\delta)}\right)^2 && \text{(Since } \sin x \leq x\text{)} \\
&= \frac{\sin(\pi L\delta)^2}{\pi^2} \sum_{k=-c+1}^{c-1} \frac{1}{(k+L\delta)^2} \\
&= 1 - \frac{\sin(\pi L\delta)^2(\Gamma^{(1)}(c-L\delta)+\Gamma^{(1)}(c+L\delta))}{\pi^2} && (46) \\
&\geq 1 - \frac{\sin(\pi L\delta)^2}{\pi^2}\left(\frac{1}{c-L\delta}+\frac{1}{c+L\delta}\right) && (\Gamma^{(1)}(x) \leq \tfrac{x-1}{x^2})
\end{aligned}$$

where Equation (46) uses Lemma A.7 and that $\Gamma^{(1)}(1-x) = \frac{\pi^2}{\sin(\pi x)^2} - \Gamma^{(1)}(x)$. Recall that $\delta \in (0, \frac{1}{2L})$ and we can upper bound $\frac{1}{c-L\delta} + \frac{1}{c+L\delta}$ by $\frac{2}{c}$. Additionally, $\sin(\pi L\delta)^2 \leq \frac{1}{2}\pi^2$ and so the above can be lower bounded by $1 - \frac{1}{c}$.

Next, we show Item 3. We continue from Equation (45). We will use the following upper bound on $\mathrm{sinc}_L(x)$

$$\mathrm{sinc}_L(x) = \left(\frac{\sin(\pi Lx)}{L\sin(\pi x)}\right)^2 \leq \left(\frac{1}{L\sin(\pi x)}\right)^2 \leq \left(\frac{1}{L(\pi x - \frac{1}{3}(\pi x)^3)}\right)^2 = \left(\frac{1}{L\pi x}\cdot\frac{1}{1-\frac{1}{3}(\pi x)^2}\right)^2 \leq \frac{4}{(L\pi x)^2}$$

where the last inequality holds as long as $x \leq \frac{1}{\pi}$. Plugging this back into Equation (45), we see that

$$\|Q(H)\,|v_\alpha\rangle\,\|_2 \leq \sum_{m\in L\cdot \mathsf{disc}_L(\mathcal{I}_\omega)} \frac{4}{(L\pi(\lambda_\alpha - m/L))^2} = \frac{4}{\pi^2}\sum_{m\in L\cdot \mathsf{disc}_L(\mathcal{I}_\omega)} \frac{1}{(L\cdot\lambda_\alpha - m)^2}.$$

By assumption, all $m \in L\cdot \mathsf{disc}_L(\mathcal{I}_\omega)$ satisfies $|\lambda_\alpha - \frac{m}{L}| \geq \frac{c}{L}$, and in particular $|L\cdot\lambda_\alpha - m| \geq c$. Let $x^* = \min|L\cdot\lambda_\alpha - m| \geq c$. Observing that $L\cdot\mathsf{disc}_L(\mathcal{I}_\omega)$ consists of consecutive integers, we rewrite the above sum as,

$$\begin{aligned}
\frac{4}{\pi^2}\sum_{x=x^*}^{x^*+|\mathsf{disc}_L(\mathcal{I}_\omega)|}\frac{1}{x^2} &= \frac{4}{\pi^2}(\Gamma^{(1)}(x^*) - \Gamma^{(1)}(1+x^*+|\mathsf{disc}_L(\mathcal{I}_\omega)|)) && \text{(Lemma A.7)} \\
&\leq \frac{4}{\pi^2}\Gamma^{(1)}(x^*) && (\Gamma^{(1)}(x) \geq 0 \text{ for } x>0) \\
&\leq \frac{4}{\pi^2 x^*}. && (\Gamma^{(1)}(x) \leq \tfrac{x-1}{x^2})
\end{aligned}$$

Recalling that $x^* \geq c$, we may upper bound Equation (45) as

$$\|Q(H)\,|v_\alpha\rangle\,\|_2 \leq \frac{4}{\pi^2 c} \leq \frac{1}{c}. \tag{47}$$

□

With the above lemmas, we can prove the promised claims.

Claim 5.4. $\|Q(H) - \Pi_\Delta Q(H)\| \le \frac{2}{\sqrt{L}}$.

Proof of Claim 5.4. Since Π_Δ and $Q(H)$ are both diagonal in H's eigenbasis, it suffices to consider each eigenstate of H in turn. For $|v_\alpha\rangle$ such that $\lambda_\alpha \in \mathcal{I}_\Delta$, Π_Δ acts as the identity and clearly there is no difference between $Q(H)$ and $\Pi_\Delta Q(H)$. For $|v_\alpha\rangle$ with $\lambda_\alpha \notin \mathcal{I}_\Delta$, $\Pi_\Delta |v_\alpha\rangle = 0$ and thus a bound on $\|Q(H)|v_\alpha\rangle\|_2$ yields a bound on $\|Q(H) - \Pi_\Delta Q(H)\|$.

By definition of $Q(H)$, we can write

$$Q(H)|v_\alpha\rangle = \sum_{m=(e_0-\omega/2)L}^{(e_0+\omega/2)L} \mathrm{sinc}_L(\lambda_\alpha - m/L)^2 |v_\alpha\rangle \tag{48}$$

$$\implies \|Q(H)|v_\alpha\rangle\|_2 = \sum_{m=(e_0-\omega/2)L}^{(e_0+\omega/2)L} \mathrm{sinc}_L(\lambda_\alpha - m/L)^2. \tag{49}$$

We defined $\omega = \Delta - \frac{1}{\sqrt{L}}$ and since $\lambda_\alpha \notin \mathcal{I}_\Delta$, $|\lambda_\alpha - \epsilon_0| \ge \Delta/2 \ge \omega/2 + \frac{1}{2\sqrt{L}}$. As a result, $|\lambda_\alpha - \frac{m}{L}| \ge \frac{1}{2\sqrt{L}}$. Using Item 3 of Lemma A.8 with $c = \frac{\sqrt{L}}{2}$ yields that $\|Q(H)|v_\alpha\rangle\|_2 \le \frac{2}{\sqrt{L}}$. □

Claim 5.10. *We have* $\langle\Psi| Q \otimes Q |\Psi\rangle \ge 1 - \frac{1}{\sqrt{L}} - (1-C')\frac{6}{f^4} = 1 - \frac{1}{\sqrt{L}} - 6f^{21}$.

Proof of Claim 5.10. Let $\mathcal{Q} \triangleq Q(H) \otimes Q(H)$. Note that $|\Psi\rangle = \sum_{\alpha\in\mathcal{I}_\Delta} x_\alpha |v_\alpha, v_\alpha\rangle$ is such that all coefficients are positive and by Claim 5.9 are within a factor of $\frac{f^4}{6}$ of each other. Define

$$S = \{\alpha : \lambda_\alpha \in [e_0 - \tfrac{\omega}{2} + \tfrac{c}{L}, e_0 + \tfrac{\omega}{2} - \tfrac{c}{L}]\}$$

and

$$\overline{S} = \{\alpha : \lambda_\alpha \in [e_0 - \tfrac{\Delta}{2}, e_0 - \tfrac{\omega}{2} + \tfrac{c}{L}) \cup (e_0 + \tfrac{\omega}{2} - \tfrac{c}{L}, e_0 + \tfrac{\Delta}{2}]$$

Thus,

$$\begin{aligned}
\langle\Psi|\mathcal{Q}|\Psi\rangle &= \sum_\alpha x_\alpha^2 \cdot \|Q\otimes Q|v_\alpha, v_\alpha\rangle\|_2 = \sum_\alpha x_\alpha^2 \cdot \|Q|v_\alpha\rangle\|_2^2 && (Q \text{ is diagonal w.r.t. } \{|v_\alpha\rangle\}_\alpha)\\
&\ge \left(1 - \frac{1}{c}\right)\sum_{\alpha\in S} x_\alpha^2 && (\text{Item 2 of Lemma A.8})\\
&= \left(1 - \frac{1}{c}\right)\left(1 - \sum_{\alpha\in\overline{S}} x_\alpha^2\right)\\
&\ge \left(1 - \frac{1}{c}\right)(1 - |\overline{S}| x_{\max}^2) && (x_{\max} \triangleq \max_\alpha x_\alpha.)
\end{aligned}$$

However, $\sum_\alpha x_\alpha^2 = 1$, which means $x_{\max}^2(|S| + |\overline{S}|) \le \frac{6}{f^4}$. Thus, $\langle\Psi|Q|\Psi\rangle \ge (1 - 1/c)\big(1 - \frac{6}{f^4}\frac{|\overline{S}|}{|S|+|\overline{S}|}\big)$. By Property 1 we have that for $c = \sqrt{L}$, $\frac{|\overline{S}|}{(|S|+|\overline{S}|)} \le 1 - C'$, which implies

$$\langle\Psi|Q|\Psi\rangle \ge \left(1 - \frac{1}{\sqrt{L}}\right)\left(1 - \frac{6(1-C')}{f^4}\right) \ge 1 - \frac{1}{\sqrt{L}} - \frac{6(1-C')}{f^4}.$$

This completes the proof. □

B Lipschitz property of the Ky-Fan 2 Norm

Recall that V^S is a (POVM for a) L-query oracular verifier, so that,

$$V^S = (|0\rangle\langle 0| \otimes \mathbb{I})U_0^\dagger(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_1^\dagger \dots U_{L-1}^\dagger(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_L^\dagger \cdot \\ (\mathbb{I} \otimes |1\rangle\langle 1|)U_L\mathcal{O}^S U_{L-1} \dots U_1(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_0(|0\rangle\langle 0| \otimes \mathbb{I})\,. \tag{50}$$

The Ky-Fan k norm is defined as the 1-norm of the top k singular values of an operator and naturally arises when considering operators V^S which correspond to UniqueQMA verifiers. We show that the Lipschitz constant of these operators is bounded.

Lemma 3.11. *For a random k-dimensional subspace $\mathcal{S}$ and verifier $V^{\mathcal{S}}$ with oracle access to $\mathcal{O}^{\mathcal{S}}$, define $F(V^{\mathcal{S}}) = \|V^{\mathcal{S}}\|_{\uparrow 2}$. Then, F is $\mathcal{O}(L)$-Lipschitz.*

Proof. To bound $|F(\Pi_S) - F(\Pi_{S'})|$, we will hybridize between S and S'. Define "forward" and "backward" hybridized circuits,

$$C^\dagger[\ell] \triangleq (|0\rangle\langle 0| \otimes \mathbb{I})U_0^\dagger(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_1^\dagger \dots U_{L-\ell}^\dagger(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_{L-\ell+1}^\dagger(\Pi_{\text{control}} \otimes \mathcal{O}^{S'}) \dots U_L^\dagger(\mathbb{I} \otimes |1\rangle\langle 1|)$$
$$C[\ell] \triangleq (\mathbb{I} \otimes |1\rangle\langle 1|)U_L(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_{L-1} \dots U_\ell(\Pi_{\text{control}} \otimes \mathcal{O}^S)U_{\ell-1}(\Pi_{\text{control}} \otimes \mathcal{O}^{S'}) \dots U_0(|0\rangle\langle 0| \otimes \mathbb{I})$$

where for both $C[\ell]$ and $C^\dagger[\ell]$, ℓ denotes the number of oracle calls from the right which are S' (the remaining being to S). Therefore, to hybridize V^S_{accept} and $V^{S'}_{\text{accept}}$, we proceed as,

$$V^S_{\text{accept}} = C^\dagger[0]C[0] \to C^\dagger[0]C[1] \to \dots \to C^\dagger[0]C[L] \to C^\dagger[1]C[L] \to \dots \to C^\dagger[L]C[L] = V^{S'}_{\text{accept}}$$

Then,

$$\begin{aligned} |F(V^S) - F(V^{S'})| &= |\|V^S_{\text{accept}}\|_{\uparrow 2} - \|V^{S'}_{\text{accept}}\|_{\uparrow 2}| \\ &\le \|V^S_{\text{accept}} - V^{S'}_{\text{accept}}\|_{\uparrow 2} \\ &\le \left\| \sum_{\ell=0}^{L-1} (C^\dagger[0]C[\ell] - C^\dagger[0]C[\ell+1]) + \sum_{\ell=0}^{L-1} (C^\dagger[\ell]C[L] - C^\dagger[\ell+1]C[L]) \right\|_{\uparrow 2} \\ &\le \sum_{\ell=0}^{L-1} \|C^\dagger[0]C[\ell] - C^\dagger[0]C[\ell+1]\|_{\uparrow 2} + \sum_{\ell=0}^{L-1} \|C^\dagger[\ell]C[L] - C^\dagger[\ell+1]C[L]\|_{\uparrow 2} \end{aligned} \tag{51}$$

Consider the first summand. Note that besides the projectors $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ at the beginning and end of $C^\dagger[0]$, the remainder of the terms are unitary. Write $C^\dagger[0] = \Pi_0 U \Pi_1$. Using sub-multiplicativity we write,

$$\|\Pi_0 U \Pi_1(C[\ell] - C[\ell+1])\|_{\uparrow 2} \le \|\Pi_0\|_{\uparrow 2} \cdot \|U\Pi_1(C[\ell] - C[\ell+1])\|_{\uparrow 2} \le 2\|U\Pi_1(C[\ell] - C[\ell+1])\|_{\uparrow 2}\,,$$

where we used that since Π_0 is a projector, $\|\Pi_0\|_{\uparrow 2} \le 2$. Next, since the Ky Fan norm is unitarily invariant, U does not affect the norm, and we can accrue another factor of 2 to remove the Π_1 term as well. Thus, for any ℓ

$$\|C^\dagger[0]C[\ell] - C^\dagger[0]C[\ell+1]\|_{\uparrow 2} \le 4\|C[\ell] - C[\ell+1]\|_{\uparrow 2}\,.$$

For this final term, we again remove the leading and trailing projector yielding a multiplicative penalty of 2, then we are left with

$$\|U_L\mathcal{O}^S U_{L-1} \dots U_\ell(\Pi_{\text{control}} \otimes (\mathcal{O}^S - \mathcal{O}^{S'}))U_{\ell-1}\mathcal{O}^{S'}U_{\ell-2} \dots U_0\|_{\uparrow 2}\,.$$

Everything except the inner $\Pi_{\text{control}} \otimes (\mathcal{O}^S - \mathcal{O}^{S'})$ terms can be discarded, as they are unitary and it suffices to bound $\|\Pi_{\text{control}} \otimes (\mathcal{O}^S - \mathcal{O}^{S'})\|_{\uparrow 2}$. Note that since Π_{control} is a projector, it simply copies the singular values

of $\mathcal{O}^S - \mathcal{O}^{S'}$ $\dim(\Pi_{\text{control}})$-times. But we're only taking the top 2 eigenvalues, so this yields a multiplicative penalty of 2. In conclusion, we see that

$$\|C^\dagger[0]C[\ell] - C^\dagger[0]C[\ell+1]\|_{\uparrow 2} \leq 16\|\mathcal{O}^S - \mathcal{O}^{S'}\|_{\uparrow 2} = 32\|\Pi^S - \Pi^{S'}\|_{\uparrow 2}$$

The final step is to relate the Ky Fan 2-norm to the Frobenius norm. Recall each of the norms' definitions in terms of singular values, we note that,

$$\sigma_1 + \sigma_2 \leq \sqrt{2}\sqrt{\sigma_1^2 + \sigma_2^2} \leq \sqrt{2}\sqrt{\sum_i \sigma_i^2} = \sqrt{2}\|\Pi_S - \Pi_{S'}\|_2$$

Applying this argument to each of the L terms, and an analogous argument to the second summand in Equation (51) yields a final bound of

$$|F(\Pi_S) - F(\Pi_{S'})| \leq 64\sqrt{2}L\|\Pi_S - \Pi_{S'}\|_2$$

□

C On the number of efficiently verifiable states

In this appendix we argue that there are comparably few states that can be uniquely verified efficiently. More precisely, we show the following statement:

Proposition C.1. *Let $\delta > 0$ and $p(n)$ a polynomial in n. Further, denote by S_n the set of all states $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$ that are verifiable by a quantum circuits with $p(n)$ 2-local gates. Then, S_n can be covered by $O_\delta(e^{np(n)})$ many balls of radius δ.*

Proof. A state $|\psi\rangle$ is verifiable if and only if a quantum circuit C accepts $|\psi\rangle$ with probability $\geq \frac{2}{3}$ and rejects all state orthogonal to $|\psi\rangle$ with probability $\geq \frac{2}{3}$. We can amplify this probability with majority vote to $1 - e^{-\Omega(k)}$ by increasing the gate count by a factor of k.

By standard arguments [MS97] an ε-net N_ε on $SU(4)$ of size $O((5/\varepsilon)^{15})$ exists. We pick $\varepsilon := 2^{-n}$. Now consider any state $|\psi\rangle \in S_n$ and its verifying circuit C_ψ with at most $kp(n)$ gates. Then, by a standard telescope argument [BBBV97], there is a circuit C'_ψ with gates from N_ε, which generates states that are $kp(n)2^{-n}$ close in infidelity. Then, C'_ψ accepts $|\psi\rangle$ with probability at least $1 - e^{-\Omega(k)} - kp(n)2^{-n}$ and rejects every orthogonal state with probability at least $1 - e^{-\Omega(k)} - kp(n)2^{-n}$. For every such verifying circuit with gates in N_ε we pick one state that is verified with high probability. The resulting set is called $T_{\varepsilon,k}$ By the above argument, any state in S_n needs to have an overlap of $1 - O(e^{-\Omega(k)} + kp(n)2^{-n})$ with a state from $T_{\varepsilon,k}$. It then suffices to upper bound

$$|T_{n,k}| \leq |N_\varepsilon|^{kp(n)} \leq e^{O(nkp(n))}. \tag{52}$$

Picking k a sufficiently large constant completes the argument. □