Quantum State Search with Classical Interactive Advice

Mohammed Barhoush

IonQ Inc.

Abstract. We study whether a quantum verifier can find a hidden quantum state by exchanging classical messages with an unbounded prover who knows the state. The verifier has access to a reflection oracle about the target. For a Haar-random n-qubit target, we prove that achieving constant average fidelity with the target requires

$$(q+1)^2(B+1) = \Omega(2^n),$$

where q is the number of verifier queries and B is the total prover-to-verifier communication. The bound allows arbitrary classical interaction and unrestricted computation between queries. Our proof extends the lower bound of Aaronson and Kuperberg (ToC'07) from a single classical proof to an arbitrary classical interaction.

As a result, we obtain a quantum-oracle separation of QMA from QCIP. In particular, we give a unitary oracle $\mathcal{U}$ such that there exists a language $L \in \mathsf{QMA}^{\mathcal{U}}$ but $L \notin \mathsf{QCIP}^{\mathcal{U}}$. Our separation holds even for $\mathsf{UniqueQMA}$, the restriction of QMA to verification with a unique quantum witness, which is known to be strictly weaker than QMA in the quantum black-box setting.

We further give an oracle construction for a notion known as QMA–$\mathsf{QCIP}[2]$ gap problem. This gives average case hardness for a QMA language under classical access to a $\mathsf{QCIP}[2]$ decision oracle and is used in establishing negative results on witness encryption and quantum indistinguishability obfuscation.

1 Introduction

Unstructured search is one of the basic problems in computer science. For a search space of size 2^n, Grover's algorithm uses $O(2^{n/2})$ queries, and this dependence is optimal [10,24]. A prover who knows the marked item can remove this query requirement by sending its n-bit address to the verifier. Does classical interaction provide a similar advantage when the hidden object is a quantum state?

Aaronson and Kuperberg [3] studied this question for a single classical message. They used an oracle that marks an unknown n-qubit state $|\psi\rangle$ by applying the reflection

$$R_\psi = \mathbb{I} - 2|\psi\rangle\langle\psi|.$$

They proved that polynomial classical advice is not sufficient to find the hidden state and distinguish this oracle from the identity, given only polynomial queries. On the other hand, quantum advice easily allows distinguishing this oracle from the identity using ψ as advice.

These findings have important implications for the relationship between the classes QMA and QCMA. Informally, both these classes describe problems whose solutions can be verified by a polynomial-time quantum verifier, using a quantum witness in the case of QMA and a classical witness in the case of QCMA. The relation between these classes captures a fundamental question in quantum complexity theory: can quantum witnesses enable more powerful verification than classical witnesses?

Aaronson and Kuperberg's result shows that a classical witness cannot do what a quantum witness can do in this oracle model, thus demonstrating the first oracle separation between QMA and QCMA.

We ask whether the same advantage holds against classical *interactive* advice. In our setting, the verifier interacts with an unbounded prover that knows ψ and tries to help the verifier find the hidden state through classical interaction. The verifier may retain quantum memory and interleave classical messages with oracle queries. Such a protocol cannot simply be replaced by one short classical advice: there is not a necessarily "optimal" transcript of interactive classical messages that can be sent in one-shot to the verifier since the success of a transcript may depend on the internal state of the verifier and quantum computations cannot be derandomized.

Our main result extends the Aaronson–Kuperberg lower bound to arbitrary classical interaction.

Theorem 1 (Hidden-state search with classical interaction, informal). *Let $|\psi\rangle$ be a Haar-random n-qubit state. If the verifier makes q oracle queries to the oracle R_ψ and the prover sends at most B classical bits in total, then producing an output with constant average fidelity with $|\psi\rangle$ requires*

$$(q+1)^2(B+1) = \Omega(2^n).$$

Fidelity means $\langle\psi|\rho|\psi\rangle$ for the verifier's output ρ, averaged over the hidden state and the protocol's randomness. For polynomial prover-to-verifier communication, the theorem gives an exponential query requirement. For $2n \leq B = o(2^n)$, it matches the one-message upper bound of Aaronson and Kuperberg [3]. Classical interaction therefore does not aid in performing unstructured search for a hidden Haar random state.

The Aaronson–Kuperberg separation actually separates UniqueQMA from QCMA, since the QMA language built in their oracle model has one unique witness up to global phase. The class UniqueQMA restricts QMA to yes-instances with a one-dimensional accepting witness subspace: a designated witness is accepted with high probability, while every orthogonal witness is accepted with low probability [27].

Our oracle construction generalizes the Aaronson–Kuperberg separation by extending the QCMA to the class QCIP, which is the complexity class with a quantum polynomial-time verifier and a computationally–unbounded prover that can interact with each other through classical messages. In particular, as a direct consequence of Theorem 1 we obtain the following result.

Theorem 2. *There is a quantum oracle $\mathcal{U}$ such that*

$$\mathsf{UniqueQMA}^{\mathcal{U}} \not\subseteq \mathsf{QCIP}^{\mathcal{U}}.$$

The class QCIP was introduced by Buhrman, Le Gall, and Weggemans [16] and was further studied by Grewal and Kretschmer [23]. Notice that $\mathsf{QCIP}[1] = \mathsf{QCMA}$[1]. This separation does not assert an unrelativized separation, since polynomial-round QCIP contains QMA in the plain model [16]. It is also consistent with computationally sound classical verification of quantum computations [33], whose soundness under the learning-with-errors assumption applies only against efficient cheating provers.

In concurrent and independent work, Çakan [17] separates QMA from polynomial-round QCIP using a different approach than our own. Çakan's separation has the advantage of being relative to a classical oracle which is a more restricted model. In comparison, our approach provides a simpler and perhaps more intuitive separation. Moreover, our result separates UniqueQMA rather than QMA, and thus provides a more fine-tuned separation. Çakan's separation does not provide a unique witness to the constructed QMA language so it does not separate UniqueQMA from QCIP. Note that uniqueness is a substantive restriction: there exists black-box separations between QMA and UniqueQMA [5].

An average-case gap problem. Next, we use our lower bound to build a primitive called a QMA–QCIP[2] gap problem, introduced by Barhoush, Morimae, and Movassagh [7]. Assuming this gap problem, they present a restricted black-box

[1] For an integer k, we let $\mathsf{QCIP}[k]$ denote the QCIP class restricted to k classical messages.

separation between witness encryption with classical ciphertexts and any falsifiable assumption. This class of assumptions essentially includes any assumption where the challenger can efficiently check whether the adversary has won the game and includes one-way functions and LWE. Their barrier also implies a separation between classically encoded quantum indistinguishability obfuscation and falsifiable assumptions. We provide an oracle construction of this gap problem which supports these separations.

Theorem 3 (Oracle gap problem, informal). *There exist a quantum unitary oracle $\mathcal{T}$ and a language $G \in \mathsf{QMA}^{\mathcal{T}}$ with yes- and no-instances efficiently samplable relative to $\mathcal{T}$. Nevertheless, the classical yes- and no-instance distributions remain computationally indistinguishable to quantum polynomial-time algorithms with access to $\mathcal{T}$ and classical access to a $\mathsf{QCIP}[2]^{\mathcal{T}}$ decision oracle.*

Barhoush, Morimae, Nishimaki, and Yamakawa [8] introduce a similar notion of QMA–QCMA gap assumption, which they build relative to a quantum oracle. Assuming such a gap, they obtain a quantum black-box barrier for non-interactive classical verification of quantum computation from falsifiable assumptions. We adapt their oracle construction to permit QCIP[2] decision access.

We will now discuss how our results relate to various areas in quantum cryptography.

State preparation with verification access. The task of transmitting a quantum state using classical interaction is known as remote state preparation. This area has well-established communication–entanglement tradeoffs [12,13], including lower bounds [30]. These models do not provide the verifier with a reflection oracle for the state to be prepared, which allows the ability to verify a state. Our contribution is to include this additional resource and quantify its tradeoff with classical communication. Note that verification ability can help restore a state when a sufficiently large subsystem is already available [20]. The task of preparing state given the ability to verify is also closely related to publicly verifiable quantum money [2,1].

Comparison with quantum no-telegraphing and quantum fire. Nehoran and Zhandry [35] establish quantum no-telegraphing even when the sender is computationally unbounded: a receiver with bounded oracle access cannot reconstruct the state from a polynomial-length classical description. Their result concerns one-way communication, whereas our lower bound permits arbitrary classical interaction.

Çakan, Goyal, and Shmueli [18] establish interactive untelegraphability relative to a classical oracle. Their model allows unbounded computation between oracle queries and even unrestricted classical communication, but bounds the oracle queries of *both* the sender and the receiver. This restriction on the sender is essential when comparing their result with ours: their security guarantee does not cover classical messages obtained through unrestricted oracle computation. In our setting, the prover has complete knowledge of the hidden state $|\psi\rangle$ and

may perform arbitrary computation and oracle queries. We bound only the total prover-to-verifier communication and the verifier's oracle queries. Thus, our lower bound does not rely on the prover's inability to learn the state or obtain a useful description of it.

Casper, Nehoran, and Sattath [19] construct conversable quantum fire in the standard model under cryptographic assumptions. Their states resist one-way classical transfer by quantum polynomial-time parties but admit transfer through classical interaction.

1.1 Acknowledgments

The concepts and proof ideas were mainly generated by the authors. We used ChatGPT 6 Astra to help write the proofs, improve presentation and check for mistakes. The authors take responsibility for any mistakes or issues.

1.2 Technical overview

We now give a simplified description of the proof of our separation.

The unitary oracle and UniqueQMA language we construct is the same as the Aaronson–Kuperberg separation. Their oracle hides an unknown Haar random n-qubit pure state $|\psi\rangle$ through the reflection

$$R_\psi = \mathbb{I} - 2|\psi\rangle\langle\psi|.$$

The promise is that the oracle is either R_ψ, for some hidden state $|\psi\rangle$, or the identity. A UniqueQMA prover can send $|\psi\rangle$ itself, after which one controlled oracle query distinguishes the two cases perfectly. Aaronson and Kuperberg proved that any verifier that distinguishes R_ψ from $\mathbb{I}$ with bounded error on average for a Haar random n-qubit state ψ, even when given an m-bit classical witness chosen as an arbitrary function of ψ, must make $\Omega(\sqrt{2^n/(m+1)})$ queries.

We need to generalize this argument to QCIP provers. We first show that the prover in a QCIP protocol can be made deterministic. In particular, for each hidden state ψ, we fix one optimal prover that may depend arbitrarily on the hidden state. The idea is that given any transcript of messages and the state ψ, there must exist a response by the prover that maximizes the probability of the verifier accepting. We fix such a maximizing response for each transcript to obtain a deterministic prover.

The transcript-information lemma. Set $D = 2^n$, and let μ_D be the Haar probability measure on pure states in $\mathbb{C}^D$. For each ψ, let P_ψ be the deterministic prover fixed above. A public transcript h is the ordered list of classical messages exchanged so far. Let $C(h)$ be the set of hidden states whose provers send all the replies recorded in h, when given the verifier messages recorded there. For a positive-probability transcript h, define

$$L(h) := \ln \frac{1}{\mu_D(C(h))}.$$

The quantity $\mu_D(C(h))$ is the probability that a Haar-random state belongs to $C(h)$. We now consider the experiment in which the verifier's access to the reflection oracle is replaced with access to an identity oracle. In this case, the verifier only has knowledge of the state ψ through the messages in the transcript h. In other words, the probability of generating the verifier messages in h is the same for every state in $C(h)$. Thus, conditioned on h, the hidden state has Haar measure restricted to $C(h)$, and the verifier's conditional private state depends only on h. Thus $L(h)$ measures how much the transcript has narrowed the possible hidden states relative to the verifier. Initially,

$$C(\varnothing) = \mathbb{CP}^{D-1} \qquad \text{and} \qquad L(\varnothing) = 0.$$

A verifier message does not change this quantity. Indeed, conditioned on the earlier transcript h, the verifier's state is independent of the particular hidden state in $C(h)$. Therefore, if the verifier next sends a, then

$$C(ha) = C(h) \quad \text{and} \quad L(ha) = L(h).$$

Now suppose that the prover sends a reply z. Since P_ψ is deterministic, this reply restricts the compatible set to

$$C(hz) = \{\psi \in C(h) : P_\psi(h) = z\}.$$

The sets $C(hz)$, as z ranges over the possible replies, partition $C(h)$. Define

$$p_z := \frac{\mu_D(C(hz))}{\mu_D(C(h))}.$$

Note that p_z is exactly the conditional probability that the prover sends z, given the earlier transcript h.

If reply z is observed, the increase in L is

$$\begin{aligned} L(hz) - L(h) &= \ln \frac{1}{\mu_D(C(hz))} - \ln \frac{1}{\mu_D(C(h))} \\ &= \ln \frac{\mu_D(C(h))}{\mu_D(C(hz))} = \ln \frac{1}{p_z}. \end{aligned}$$

Therefore, averaging over the possible replies, the expected increase is

$$\mathbb{E}[L(hZ) - L(h) \mid h] = \sum_z p_z \ln \frac{1}{p_z}.$$

Here the expectation is over the distribution of ψ. The right-hand side is the entropy of the prover's reply, measured using natural logarithms. If the reply contains b bits, it has at most 2^b possible values, and hence

$$\sum_z p_z \ln \frac{1}{p_z} \le \ln(2^b) = b \ln 2.$$

Thus, a b-bit prover message can increase the expected value of L by at most $b \ln 2$, while verifier messages do not increase it at all.

Finally, let B_t be the total number of prover bits sent before query t and let T_t denote the public transcript immediately before that query. Summing the preceding bound over all prover messages before that query gives

$$\mathbb{E}[L(T_t)] = \mathbb{E}\left[\ln \frac{1}{\mu_D(C(T_t))}\right] \leq B_t \ln 2. \tag{II}$$

This expectation correctly weights each transcript by its probability: transcripts with very small compatible sets may have large $L(h)$, but they are correspondingly unlikely to occur.

From information to query complexity. The bound on $\mathbb{E}[L]$ allows us to bound the information the verifier has before each oracle query. In particular, it allows us to show that before any query, the query register has expected overlap at most $O((B+1)/D)$ with the hidden state. Replacing that identity query by a reflection changes the state by at most twice the square root of this overlap in trace distance i.e. $O(\sqrt{(B+1)/D})$. We now replace all q queries one at a time until we return to the standard experiment where all queries by the verifier are answered with the reflection oracle. The triangle inequality gives a total expected change of at most

$$O\left(q\sqrt{\frac{B+1}{D}}\right).$$

This is where q enters: each query contributes to the total change. Constant distinguishing advantage consequently requires $(q+1)^2(B+1) = \Omega(D)$.

The QMA–QCIP[2] *gap problem.* For the gap problem, we follow a similar construction to that of Barhoush, Morimae, Nishimaki, and Yamakawa [8].

At length n, choose a random sparse set L_n of yes-labels and an independent Haar-random state $|\psi_x\rangle$ for each label x. The construction has three oracle components:

- **Testing oracle.** Given a label x and a quantum state, this oracle applies the reflection about $|\psi_x\rangle$ if $x \in L_n$, and the identity otherwise. Thus, a verifier given $|\psi_x\rangle$ as a quantum witness can test whether x is a yes-instance using one controlled query.
- **Yes-sampling oracle.** Starting from a blank state, this oracle prepares a uniform superposition of yes-labels, each accompanied by a requested number of copies of its corresponding witness. Measuring the label register gives a uniformly random $x \in L_n$, together with copies of $|\psi_x\rangle$.
- **No-sampling oracle.** Starting from a blank state, this oracle prepares a uniform superposition of labels outside L_n. Measuring the label register gives a uniformly random no-instance.

To demonstrate the existence of a QMA–QCIP gap problem relative to these oracles, we need to establish that no QPT adversary can distinguish a randomly sampled yes-instance from a randomly sampled no-instance, even given access to these oracles as well as a QCIP[2] oracle.

Suppose an adversary receives an instance $y \in L$. Since the part of the oracle generating yes instances only generates them in a uniform superposition, it is easy to carry out a hybrid argument where y is deleted from the yes-sampling oracle and replaced into the no-sampling oracle. These changes only have negligible difference on the output of the oracles given that the state generated are in uniform superposition over exponentially many elements.

However, the adversary still has access to a QCIP oracle, which seems stronger than a QCIP verifier: the adversary can adaptively query the oracle to decide QCIP instances. Nevertheless, a change of perspective makes this situation less daunting. The adversary's interaction with the oracle resembles that of a QCIP verifier with a prover, since both involve classical interaction. By utilizing this idea, we show that our lower-bound arguments for QCIP verifiers also apply to the adversary, establishing that it cannot find the hidden state associated with y. This completes the indistinguishability argument for the QMA–QCIP gap problem.

Finally, a short second-moment argument followed by Borel–Cantelli fixes one base oracle that work for all efficient distinguishers. The formal construction and proof appear in Section 6.

2 Model and Preliminaries

All logarithms written as ln are natural logarithms. For density operators ρ and σ, let

$$\mathrm{D}(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\|_1$$

denote trace distance. Trace distance is contractive under quantum channels and bounds the difference between the probabilities of every measurement outcome.

Let μ_D be Haar measure on the pure states in $\mathbb{C}^D$, or equivalently on $\mathbb{CP}^{D-1}$. We write $\Psi \leftarrow \mu_D$ for a Haar-random pure state.

A QPT algorithm means a uniform polynomial-time quantum oracle algorithm.

2.1 The marked-state reflection oracle

For every unit vector $|\psi\rangle \in \mathbb{C}^D$, define

$$R_\psi = \mathbb{I} - 2|\psi\rangle\langle\psi|.$$

The two cases of the promise problem are

$$\mathsf{marked} : R_\psi \text{ for some } |\psi\rangle, \qquad \mathsf{unmarked} : \mathbb{I}.$$

As in Aaronson and Kuperberg [3], controlled oracle calls are allowed. This only strengthens the verifier.

2.2 Classical-message quantum interactive proofs

We make the communication model explicit. The verifier and prover start in a fixed product state; in particular, they have no shared classical or quantum setup and no shared entanglement. They exchange only classical strings. Formally, every outgoing message register is measured, or equivalently completely dephased, in the computational basis before it is transmitted. The verifier and prover may retain arbitrary private quantum registers. A *public transcript* is the round-indexed list of all transmitted classical messages. It does not include either party's private measurement outcomes.

Messages alternate, and the prover sends the last counted message. More explicitly, for $r = 2p - 1$ the order is

$$Z_1 \longrightarrow A_1 \longrightarrow Z_2 \longrightarrow \cdots \longrightarrow A_{p-1} \longrightarrow Z_p,$$

while for $r = 2p$ it is

$$A_1 \longrightarrow Z_1 \longrightarrow A_2 \longrightarrow \cdots \longrightarrow A_p \longrightarrow Z_p.$$

Here A_i is sent by the verifier and Z_i by the prover. Thus $\mathsf{QCIP}[1] = \mathsf{QCMA}$, $\mathsf{QCIP}[2]$ is verifier-first, and $\mathsf{QCIP}[3]$ has order $P \to V \to P$. Padding by ignored messages makes equivalent conventions harmless. Following Buhrman, Le Gall, and Weggemans, we use QCIP without a bracketed parameter for the union of polynomial-time protocols having at most n^d classical messages, over all fixed integers $d \geq 1$. Their total communication is also polynomial.

On a fixed input, a *deterministic classical prover* simply specifies one classical reply after every possible public transcript. We write this reply as $P(h)$, where h is the transcript. No efficiency or uniformity condition is imposed on P. When the oracle is W, let $\mathrm{Acc}_W(V, P)$ be the verifier's acceptance probability when interacting with P, and define

$$\mathrm{val}_W(V) = \max_P \mathrm{Acc}_W(V, P),$$

where the maximum is over deterministic classical provers.

Lemma 1 (Deterministic classical provers suffice). *Fix an oracle W and a protocol with classical communication and no shared initial state. The maximum acceptance probability over all quantum provers is attained by a deterministic classical prover.*

Proof. Fix any quantum prover. Conditioned on a public transcript, the parties' private states are a product state. Indeed, they start with no shared state; local computations preserve this property; and conditioning on a transmitted classical message changes only the sender's private state. Thus, once the transcript is fixed, the verifier's conditional state does not depend on the prover's private quantum memory.

We now fix the prover's replies, starting with its last message and working backwards. At a transcript where the prover sends its last message, there is a

probability distribution over the possible classical replies. Its acceptance probability is therefore an average of the acceptance probabilities obtained by fixing the reply. Some fixed reply is at least as good as this average, so choose such a reply for that transcript. At a transcript of probability zero, choose any reply. Do this for every transcript at the last prover turn. Then repeat the same argument at the preceding prover turn, with all later replies already fixed, and continue backwards. This produces a deterministic reply for every public transcript, and the resulting deterministic classical prover accepts at least as often as the original quantum prover. Finally, the protocol has finitely many turns and finite message registers, so there are only finitely many deterministic classical provers. One of them therefore attains the largest acceptance probability. ■

When $W = R_\psi$, we allow the deterministic prover P to depend arbitrarily on ψ. This already grants the prover everything it could learn from unlimited oracle access.

To define oracle-query slots unambiguously, divide the verifier's computation into phases separated by transmitted messages. Within each phase, pad every adaptive query path to the phase's maximum number of queries, using oracle calls controlled on a zero flag. Every execution then has the same chronologically ordered list of query slots. For slot t, let B_t be the sum of the padded lengths of the prover messages in earlier communication slots. This is a deterministic bound, independent of the branch taken. Let B be the total padded prover-to-verifier communication in the protocol. Private measurements may be deferred, while the measurements that create transmitted classical messages are kept explicit.

2.3 Relativized model

A quantum oracle is a sequence $\mathcal{U} = (U_1, U_2, \ldots)$, where U_n is a unitary on a specified polynomial number of qubits. We use the length-respecting convention of Aaronson and Kuperberg: on input 0^n, a machine queries only U_n. This makes the diagonal construction transparent.

2.4 Geometric Lemma

We state an important geometric lemma used in Aaronson–Kuperberg separation [3].

Definition 1. *A probability measure ν on $\mathbb{CP}^{D-1}$ is* p-uniform*, for $0 < p \leq 1$, if $p\nu(G) \leq \mu_D(G)$ for every measurable set G. Equivalently, we write $p\nu \leq \mu_D$ as an inequality of measures.*

Lemma 2 (Haar overlap under conditioning (Lemma 3.2 in [3])). *Let ν be p-uniform and let ρ be a density operator on $\mathbb{C}^D$, where $D \geq 2$. Then*

$$\mathbb{E}_{\Psi \leftarrow \nu}[\langle \Psi | \rho | \Psi \rangle] \leq \frac{1 + \ln(1/p)}{D}.$$

3 What a Classical Transcript Reveals

We first describe the experiment used throughout this section. Fix a verifier V with q padded query slots, fix $t \in [q]$, and fix a measurable choice $\psi \mapsto P_\psi$, for $\psi \in \mathbb{CP}^{D-1}$, of deterministic classical provers. The measurable choice needed for the lower bound is constructed explicitly in Theorem 4. Sample $\Psi \leftarrow \mu_D$, let the prover be P_Ψ, and replace every verifier oracle call before slot t by the identity. The prover is already represented by the fixed prover P_Ψ, so it makes no oracle calls inside this experiment. We call this the *identity-prefix execution*. Let T_t be its public classical transcript immediately before query t. The randomness in T_t comes from the Haar state and from the verifier's measurements and private randomness. When we say that the verifier's state is fixed by a transcript, we mean that its conditional density operator is fixed; it may be mixed because private outcomes are averaged.

For a possible transcript h, define its *compatible set*

$$C(h) = \{\psi \in \mathbb{CP}^{D-1} : P_\psi \text{ sends every prover message recorded in } h\}.$$

Thus $C(h)$ contains exactly the hidden states consistent with the prover's part of h. Verifier messages still appear in h, because they determine which reply the prover sends. Their own probabilities impose no further restriction on ψ once the preceding transcript is fixed. Measurability of the choice $\psi \mapsto P_\psi$ makes every set $C(h)$ Haar measurable.

Warm-up: three messages. Consider a transcript

$$Y \longrightarrow A \longrightarrow Z,$$

where Y, Z are prover messages and A is the verifier's challenge. Capital letters denote random messages, and lower-case letters denote their realized values. After observing $Y = y$, the compatible states form $C(y)$. Once y is fixed, the identity-prefix verifier's private state is fixed as well. Hence the distribution of A is the same for every $\psi \in C(y)$, and observing $A = a$ does not shrink the set:

$$C(y, a) = C(y).$$

The final prover message can shrink it:

$$C(y, a, z) = \{\psi \in C(y) : P_\psi(y, a) = z\}.$$

Thus Y and Z can reveal information about Ψ, but A cannot reveal new information once Y is known. If query t occurs before Z, then only the bits in Y are charged.

Lemma 3 (Query overlap after a classical transcript). *In the identity-prefix execution, let $\rho_{t,h}$ be the state of the oracle-input register before query t, conditioned on transcript $T_t = h$. If at most B_t padded prover bits precede this query, then*

$$\mathbb{E}_{\Psi, T_t}[\langle\Psi|\rho_{t,T_t}|\Psi\rangle] \le \frac{1 + B_t \ln 2}{D}. \tag{1}$$

Proof. We first show that, conditioned on any possible transcript h, the hidden state is Haar distributed over $C(h)$, and the verifier's private state depends only on h, not on the particular $\psi \in C(h)$.

Both statements hold at the empty transcript. Now suppose first that the verifier sends the next message. Given the preceding transcript, its private state is fixed because all earlier verifier oracle calls are identities. Therefore the distribution of its message, and its residual post-measurement state, do not depend on ψ. The compatible set is unchanged. Similarly, at any point in the protocol, when the verifier sends a message, it does not change the compatible set.

Suppose instead that the prover sends the next message z. Since P_ψ is deterministic, the new compatible set is exactly

$$C(hz) = \{\psi \in C(h) : P_\psi(h) = z\},$$

where hz denotes h extended by z. Conditioning on this reply restricts the hidden state to $C(hz)$. This proves the induction.

We next bound the expected logarithmic inverse mass of the compatible set. At the empty transcript this quantity is zero, and a verifier message leaves it unchanged. A padded b-bit prover message partitions $C(h)$ into the sets $C(hz)$. Let

$$p(z \mid h) = \frac{\mu_D(C(hz))}{\mu_D(C(h))}.$$

Therefore, the expected logarithmic mass of the compatible set is:

$$\begin{aligned}\sum_z p(z \mid h) \ln \frac{1}{\mu_D(C(hz))} &= \ln \frac{1}{\mu_D(C(h))} + \sum_z p(z \mid h) \ln \frac{1}{p(z \mid h)} \\ &\le \ln \frac{1}{\mu_D(C(h))} + b \ln 2.\end{aligned}$$

The inequality holds because a b-bit message has at most 2^b possible values, so its conditional entropy is at most $b \ln 2$. Average over the preceding transcript and repeat for every prover message before query t. Their padded lengths total at most B_t, giving

$$\mathbb{E}\left[\ln \frac{1}{\mu_D(C(T_t))}\right] \le B_t \ln 2. \tag{2}$$

For each possible transcript h, the query state $\rho_{t,h}$ is fixed, whereas the hidden state has Haar measure conditioned on $C(h)$. Lemma 2 therefore gives

$$\mathbb{E}[\langle \Psi | \rho_{t,h} | \Psi \rangle \mid T_t = h] \le \frac{1 + \ln(1/\mu_D(C(h)))}{D}.$$

Averaging over T_t and applying Eq. (2) proves the claim. ■

4 The General Query Lower Bound

We combine Lemma 3 with the standard quantum-query hybrid method [10]. The following theorem also allows access to a fixed background oracle independent of Ψ, which we suppress from the notation.

Theorem 4 (Classical-interaction query bound). *Let V exchange only classical messages, make at most q queries to the distinguished D-dimensional oracle, and receive at most B padded prover bits. Then*

$$\mathbb{E}_{\Psi}\Big[\max_{P}\left|\mathrm{Acc}_{R_\Psi}(V,P)-\mathrm{Acc}_{\mathbb{I}}(V,P)\right|\Big]\le 2q\sqrt{\frac{1+B\ln 2}{D}}, \tag{3}$$

where the maximum is over deterministic classical provers. The same bound holds for $\mathbb{E}_{\Psi}[|\mathrm{val}_{R_\Psi}(V)-\mathrm{val}_{\mathbb{I}}(V)|]$.

Proof. For every ψ, choose a deterministic classical prover P_ψ maximizing the absolute difference in Eq. (3). At each fixed input length, there are finitely many possible deterministic response functions on the protocol's transcript tree. Order them and choose the first maximizer. Since acceptance probabilities for each fixed prover are continuous in ψ, this choice is Borel measurable.

Use the same prover throughout the hybrid. Only the verifier makes oracle calls as the prover has full knowledge of ψ. Order the verifier's oracle calls chronologically. For $t=0,\dots,q$, let Ω_t^{ψ} be the final state of the verifier and public transcript when the first t calls use $\mathbb{I}$ and the remaining calls use R_ψ.

To bound the effect of one replacement, purify the state before the query to $|\Phi\rangle$, and let u be the overlap of its oracle-input register with $|\psi\rangle$. The two post-query purifications have inner product $1-2u$, so their trace distance is

$$\sqrt{1-|1-2u|^2}=2\sqrt{u(1-u)}\le 2\sqrt{u}.$$

Discarding the purifying register and applying later operations cannot increase trace distance. For a controlled query, u is replaced by the weight on $|1\rangle\otimes|\psi\rangle$, which is at most u, so the same bound holds.

Conditioned on transcript h before query t, let $\rho_{t,h}$ be the reduced density operator of the oracle-input register, obtained by tracing out all other registers; the overlap defined above is therefore $u=\langle\psi|\rho_{t,h}|\psi\rangle$. Convexity of trace distance, concavity of the square root, and Lemma 3 give

$$\begin{aligned}\mathbb{E}_{\Psi}[\mathrm{D}(\Omega_{t-1}^{\Psi},\Omega_t^{\Psi})] &\le 2\mathbb{E}_{\Psi,T_t}\left[\sqrt{\langle\Psi|\rho_{t,T_t}|\Psi\rangle}\right]\\ &\le 2\sqrt{\frac{1+B\ln 2}{D}}.\end{aligned}$$

We can telescope over the q queries to obtain Eq. (3) for the chosen maximizing provers.

We now compare the optimal acceptance probabilities. Suppose first that $\mathrm{val}_{R_\psi}(V)\ge\mathrm{val}_{\mathbb{I}}(V)$, and choose a deterministic prover $P^\star$ that achieves $\mathrm{val}_{R_\psi}(V)$. This prover specifies a classical reply to every possible transcript. We can use exactly the same reply rules when the verifier instead queries the identity oracle. Since $\mathrm{val}_{\mathbb{I}}(V)$ is the maximum acceptance probability over all prover strategies, it satisfies

$$\mathrm{val}_{\mathbb{I}}(V)\ge\mathrm{Acc}_{\mathbb{I}}(V,P^\star).$$

Together with $\mathrm{val}_{R_\psi}(V) = \mathrm{Acc}_{R_\psi}(V, P^\star)$, this gives

$$\mathrm{val}_{R_\psi}(V) - \mathrm{val}_{\mathbb{I}}(V) \le \mathrm{Acc}_{R_\psi}(V, P^\star) - \mathrm{Acc}_{\mathbb{I}}(V, P^\star).$$

Thus, the difference between the optimal values is at most the difference obtained by keeping this one prover fixed. If the identity-oracle value is larger, choose a prover optimal for the identity oracle and reverse the argument. In either case,

$$\big|\mathrm{val}_{R_\psi}(V) - \mathrm{val}_{\mathbb{I}}(V)\big| \le \max_P \big|\mathrm{Acc}_{R_\psi}(V, P) - \mathrm{Acc}_{\mathbb{I}}(V, P)\big|.$$

Averaging over the hidden state and applying Eq. (3) proves the second part of the theorem. In particular, a constant average absolute gap between the two values requires

$$q = \Omega\left(\sqrt{\frac{D}{B+1}}\right).$$

■

5 One Oracle Separating Every Message Count

For each $n \ge 1$, choose a bit $b_n \in \{0,1\}$ and a Haar random n-qubit state $|\psi_n\rangle$. Define the n-qubit unitary

$$U_n = \mathbb{I}_{2^n} - 2b_n|\psi_n\rangle\langle\psi_n|,$$

where $\mathbb{I}_{2^n}$ is the identity on n qubits. Thus $U_n = \mathbb{I}_{2^n}$ when $b_n = 0$, while $U_n = R_{\psi_n}$ when $b_n = 1$. The quantum oracle is the family $\mathcal{U} = (U_n)_{n\ge1}$. The construction below specifies the bits b_n and states $|\psi_n\rangle$.

For each length n, define

$$L_n^{\mathcal{U}} = \begin{cases} \{0^n\}, & b_n = 1, \\ \varnothing, & b_n = 0. \end{cases}$$

The associated unary language is

$$L^{\mathcal{U}} = \bigcup_{n\ge1} L_n^{\mathcal{U}} = \{0^n : b_n = 1\}.$$

Thus deciding whether $0^n \in L^{\mathcal{U}}$ means determining whether U_n is the identity or a reflection.

Theorem 5 (Simultaneous separation). *There exists a quantum oracle $\mathcal{U}$ such that*

$$L^{\mathcal{U}} \in \mathsf{UniqueQMA}^{\mathcal{U}} \qquad \text{and} \qquad L^{\mathcal{U}} \notin \mathsf{QCIP}^{\mathcal{U}}.$$

In particular, the same oracle satisfies $\mathsf{UniqueQMA}^{\mathcal{U}} \not\subseteq \mathsf{QCIP}[r]^{\mathcal{U}}$ for every fixed integer $r \ge 1$.

Proof. First, we show that $L^{\mathcal{U}} \in \mathsf{UniqueQMA}^{\mathcal{U}}$ for every oracle of this form. Reject non-unary inputs. On input 0^n, the verifier receives an n-qubit witness $|\phi\rangle$, prepares a control qubit in $|+\rangle$, applies controlled-U_n, applies a Hadamard gate to the control, and accepts exactly when the outcome is 1. If $U_n = R_{\psi_n}$, the witness $|\psi_n\rangle$ is accepted with probability one. If $U_n = \mathbb{I}$, every witness is rejected.

Enumerate all polynomial-time uniform quantum verifiers with classical communication as $V_1, V_2, \ldots$, using a fixed finite universal gate set. Let $q_i(n)$ and $B_i(n)$ be polynomial bounds on the oracle queries and total prover-to-verifier communication of V_i. This enumeration includes every verifier with polynomially many messages.

At stage i, choose a fresh input length n_i, larger than every length used earlier, such that the right-hand side of Eq. (3), with $q = q_i(n_i)$, $B = B_i(n_i)$, and $D = 2^{n_i}$, is less than $1/3$. This is possible because q_i and B_i are polynomials whereas 2^{n_i} is exponential.

Consider V_i on input 0^{n_i}, and let $v_0 = \mathrm{val}_{\mathbb{I}}(V_i)$ when the length-n_i component is the identity. There are two cases.

- If $v_0 > 1/3$, set $U_{n_i} = \mathbb{I}$. Then $0^{n_i} \notin L^{\mathcal{U}}$, but V_i violates soundness.
- If $v_0 \leq 1/3$, Theorem 4 gives

$$\mathbb{E}_{\Psi}[\mathrm{val}_{R_\Psi}(V_i)] \leq v_0 + \mathbb{E}_{\Psi}[|\mathrm{val}_{R_\Psi}(V_i) - v_0|] < \frac{2}{3},$$

  where Ψ is Haar distributed on n_i qubits. Hence some $|\psi_{n_i}\rangle$ satisfies $\mathrm{val}_{R_{\psi_{n_i}}}(V_i) < 2/3$. Set $U_{n_i} = R_{\psi_{n_i}}$. Then $0^{n_i} \in L^{\mathcal{U}}$, but V_i violates completeness.

Set all unused components to the identity. The length-respecting restriction on verifier oracle access ensures that later stages do not change V_i's behavior on 0^{n_i}. The value val already maximizes over every deterministic prover, so future oracle components cannot give the prover an additional advantage omitted by the bound. Thus every polynomial-round QCIP verifier fails on some input, while the same language lies in $\mathsf{UniqueQMA}^{\mathcal{U}}$. ■

6 An Average-Case QMA–QCIP[2] Gap Problem

The unary language used in Section 5 gives a worst-case separation, but it does not provide efficiently samplable instances and witnesses. We now give a sampler-augmented oracle satisfying the average-case definition of Barhoush, Morimae, and Movassagh [7]. The construction and the seven-hybrid proof follow the QMA–QCMA gap oracle of Barhoush, Morimae, Nishimaki, and Yamakawa [8]. The new ingredient is the threshold-coupling argument inside the proof of Theorem 6, which controls adaptive queries to a $\mathsf{QCIP}[2]$ decision oracle.

6.1 The relativized gap definition

We use explicit descriptions over a fixed finite universal gate set. A query to a $\mathsf{QCIP}[2]$ decision oracle is a string $z = (V_1, V_2, x)$. On input x, the first verifier circuit produces a classical challenge a and a private quantum state that it passes to the second verifier. The prover returns a classical answer depending on a, and the second verifier circuit decides whether to accept. Let $\mathrm{val}_{\mathcal{T}}(z)$ be the largest acceptance probability over all deterministic classical provers when both verifier circuits have access to the quantum oracle $\mathcal{T}$.

We use the convention in [7] to define a $\mathsf{QCIP}[2]$ oracle. Note that for a $\mathsf{QCIP}[2]$ language some inputs do not belong to yes or no instances, so it is not clear how the oracle should respond to such inputs. We deal with this as in [7] by defining a whole class of oracles where any function can belong to the class as long as it accepts yes instances and rejects no instances. Then, any adversary that uses a $\mathsf{QCIP}[2]$ oracle to solve some task should be able to succeed regardless of which function in the class it is given oracle access to. We call a function in this class a completion of the $\mathsf{QCIP}[2]$ oracle.

Definition 2 (Relativized two-message decision oracle). *The set $\mathcal{O}^{\mathcal{T}}_{\mathsf{QCIP}[2]}$ consists of all total Boolean functions $\mathcal{O}$ satisfying*

$$\mathcal{O}(z) = 1 \quad \text{if } \mathrm{val}_{\mathcal{T}}(z) \geq \frac{2}{3}, \qquad \mathcal{O}(z) = 0 \quad \text{if } \mathrm{val}_{\mathcal{T}}(z) \leq \frac{1}{3}.$$

When the value lies strictly between $1/3$ and $2/3$, the answer is unrestricted. Oracle calls are classical. In this oracle construction, a completion is an arbitrary total Boolean function satisfying these two constraints.

The quantifier over completions prevents an algorithm from relying on specially chosen answers to ambiguous queries.

Definition 3 (Relativized QMA–$\mathsf{QCIP}[2]$). *gap problem] Let $G^{\mathcal{T}} \in \mathsf{QMA}^{\mathcal{T}}$, let $Q^{\mathcal{T}}$ be its verifier, and let $\mathcal{Y}_n$ and $\mathcal{N}_n$ be distributions supported on $G^{\mathcal{T}} \cap \{0,1\}^n$ and $\{0,1\}^n \setminus G^{\mathcal{T}}$, respectively. The pair $(\mathcal{Y}, \mathcal{N})$ is a gap problem if:*

1. *for every polynomial t, a QPT oracle algorithm on input $(1^n, 1^{t(n)})$ samples $(x, |w_x\rangle^{\otimes t(n)})$, where x has distribution $\mathcal{Y}_n$ and $|w_x\rangle$ is accepted by $Q^{\mathcal{T}}$ with probability at least $2/3$;*
2. *a possibly inefficient oracle algorithm samples $\mathcal{N}_n$; and*
3. *for every polynomial s and every QPT A, there is a completion $\mathcal{O} \in \mathcal{O}^{\mathcal{T}}_{\mathsf{QCIP}[2]}$ such that*

$$\Big| \Pr_{X_1,\ldots,X_s \leftarrow \mathcal{Y}_n}[A^{\mathcal{T},\mathcal{O}}(X_1,\ldots,X_s) = 1] - \Pr_{X_1,\ldots,X_s \leftarrow \mathcal{N}_n}[A^{\mathcal{T},\mathcal{O}}(X_1,\ldots,X_s) = 1] \Big| \leq \mathrm{negl}(n).$$

The distinguisher has quantum query access to $\mathcal{T}$ and classical access to $\mathcal{O}$.

6.2 The random-subset oracle

Fix $n \geq 4$, let $N = 2^n$, and put

$$K = \left\lfloor 2^{n/2} \right\rfloor .$$

Choose L_n uniformly among all K-element subsets of $\{0,1\}^n$. Independently choose an n-qubit Haar-random state $|\psi_x\rangle$ for every $x \in \{0,1\}^n$.

The oracle has one testing component and two sampling components. The testing component W_n has a control qubit, an n-bit label register, and an n-qubit state register. On computational-basis control and label registers, it acts as

$$W_n|c,x\rangle|\phi\rangle = |c,x\rangle\, R_{\psi_x}^{c\,\mathbf{1}[x\in L_n]}|\phi\rangle. \tag{8}$$

Thus a controlled query at $x \in L_n$ applies the reflection about $|\psi_x\rangle$, while a query at $x \notin L_n$ applies the identity.

For a unit vector $|a\rangle$ orthogonal to $|0\rangle$, write

$$S(a) = |a\rangle\langle 0| + |0\rangle\langle a| + \mathbb{I} - |0\rangle\langle 0| - |a\rangle\langle a|. \tag{9}$$

This unitary swaps $|0\rangle$ and $|a\rangle$ and fixes their common orthogonal complement. For every integer $t \geq 1$, define

$$|Y_{n,t}\rangle = \frac{1}{\sqrt{K}} \sum_{x\in L_n} |x\rangle|\psi_x\rangle^{\otimes t}|1\rangle, \qquad |N_n\rangle = \frac{1}{\sqrt{N-K}} \sum_{x\notin L_n} |x\rangle|1\rangle. \tag{10}$$

The oracle contains the blocks

$$\mathsf{GenYes}_{n,t} = S(Y_{n,t}) \quad \text{for every } t \geq 1, \qquad \mathsf{GenNo}_n = S(N_n).$$

The integer t is a classical tag. Equivalently, all these blocks form one tagged block-diagonal unitary. A polynomial-time machine can address only tags and registers of polynomial length.

Let $\mathcal{T}$ denote the collection of these independently sampled components over all input lengths, and define

$$G^{\mathcal{T}} = \bigcup_{n\geq 4} L_n.$$

Membership at the finitely many smaller lengths is defined arbitrarily.

Lemma 4 (Verification and sampling). *For every realization of $\mathcal{T}$, the language $G^{\mathcal{T}}$ belongs to $\mathsf{QMA}^{\mathcal{T}}$ with perfect completeness and perfect soundness. Moreover, the uniform distributions on L_n and on its complement satisfy the first two conditions of Definition 3.*

Proof. On input x, the QMA verifier performs the Hadamard test for the controlled testing block W_n. If $x \in L_n$, the witness $|\psi_x\rangle$ is a -1 eigenvector and is accepted with probability one. If $x \notin L_n$, the addressed block is the identity for every witness, so the verifier rejects with probability one.

On input $(1^n, 1^t)$, the yes sampler applies $\mathsf{GenYes}_{n,t}$ to $|0\rangle$, measures the label register, and returns the remaining witness registers. Its output is a uniform $x \in L_n$ together with $|\psi_x\rangle^{\otimes t}$. Applying GenNo_n to $|0\rangle$ and measuring the label gives a uniform element of the complement. Each sampler makes one oracle query. ■

6.3 Stability of the oracle hybrids

The seven-hybrid proof changes one sampler term or one reflection at a time. The sampler transitions use the elementary perturbation estimate below. The reflection transitions will be handled directly inside the seven-hybrid proof by applying Theorem 4 to a suitable meta-verifier. Throughout this subsection, "exponentially small" means at most 2^{-cn} for some constant $c > 0$ and all sufficiently large n. The constant and the starting point may depend on the fixed algorithm and its polynomial resource bounds. We freely absorb polynomial factors into this notation.

Lemma 5 (Puncturing a coherent sampler). *Let* $|A\rangle = M^{-1/2} \sum_{j=1}^{M} |a_j\rangle$, *where the* $|a_j\rangle$*'s are orthonormal, and let*

$$|A^{-i}\rangle = \frac{1}{\sqrt{M-1}} \sum_{j \neq i} |a_j\rangle.$$

Then

$$\big\| S(A) - S(A^{-i}) \big\| = O(M^{-1/2}).$$

The same estimate holds when one orthogonal term is added to an M*-term superposition.*

Proof. The two target states have inner product $\sqrt{(M-1)/M}$, and hence

$$\big\| |A\rangle - |A^{-i}\rangle \big\|^2 = 2 - 2\sqrt{1 - \frac{1}{M}} = O(M^{-1}).$$

Expanding the definition of S in Eq. (9) and using $\| |a\rangle\langle a| - |b\rangle\langle b| \| \leq 2 \, \| |a\rangle - |b\rangle \|$ gives the claimed operator-norm bound. The calculation for adding a term is identical. For every fixed deterministic prover, the ordinary query hybrid changes the acceptance probability by at most twice the number of queries times the operator-norm change. Taking the maximum over deterministic provers preserves this bound.

For the tagged family $\{\mathsf{GenYes}_{n,t}\}_t$, puncturing the same label changes every tagged block by the same $O(M^{-1/2})$ amount. The norm of the tagged direct-sum difference is the maximum of the block norms, so the estimate is uniform in t. ■

We now choose the answers on ambiguous $\mathsf{QCIP}[2]$ queries. For every binary string z, independently sample

$$\Theta_z \leftarrow [2/5, 3/5]$$

once and keep it fixed. Write $\Theta = (\Theta_z)_z$ for this threshold table. Given a base oracle $\mathcal{T}$, define

$$\mathcal{O}_{\mathcal{T},\Theta}(z) = \mathbf{1}[\mathrm{val}_{\mathcal{T}}(z) \geq \Theta_z]. \tag{11}$$

The random table is only a probabilistic-method device. Once it is fixed, the completion is deterministic and every repeated query receives the same answer. Every such function is a valid deterministic completion: values at least $2/3$ are always accepted and values at most $1/3$ are always rejected. If two base oracles give values v and v' to the same description, then, using the same threshold in both worlds,

$$\Pr_{\Theta}[\mathbf{1}[v \geq \Theta_z] \neq \mathbf{1}[v' \geq \Theta_z]] \leq 5|v - v'|. \tag{12}$$

Indeed, an answer differs only when the uniform threshold falls between v and v', and the interval $[2/5, 3/5]$ has length $1/5$.

6.4 The gap problem

We follow the seven-hybrid argument of Barhoush, Morimae, Nishimaki, and Yamakawa [8, Theorem 8], using Theorem 4 to handle classical interaction.

Theorem 6 (Oracle gap problem). *For some realization $\mathcal{T}$ of the construction above, the uniform distributions on L_n and its complement form a* QMA–QCIP[2] *gap problem against uniform* QPT *distinguishers.*

Proof. QMA membership and the sampling requirements follow from Lemma 4. We prove indistinguishability first for one input label. Fix a uniform QPT algorithm A, and let $Q = Q(n)$ be a polynomial bounding its queries and the description lengths and resources of the protocols it can submit to the completion oracle, under the stated encoding conventions.

Sample the base oracle $\mathcal{T}$, a uniform $y \in L_n$, and a uniform $\bar{y} \notin L_n$. Independently sample the threshold table Θ and use it throughout. Whenever a hybrid changes the base oracle, its completion is changed accordingly, using this same table.

Consider the following hybrids, each of which runs A and outputs its result.

H_0. Use the original oracle and give A the input y.
H_1. Remove the term indexed by y from every $\mathsf{GenYes}_{n,t}$.
H_2. Replace the reflection at address y by the identity.
H_3. Add y to GenNo_n.
H_4. Keep this modified oracle, but replace the input y by $\bar{y}$.
H_5. Remove y from the no sampler.
H_6. Restore the reflection at address y.
H_7. Restore the term indexed by y in every yes sampler.

Each hybrid changes only the operation specified, leaving everything else as in the preceding hybrid. Thus H_7 uses the original oracle and a uniform no-instance.

We first explain how to account for the completion oracle when comparing two neighboring base oracles $\mathcal{T}$ and $\mathcal{T}'$. Change the completion while keeping the

base oracle fixed, then change the base oracle while keeping the completion fixed. For the first change, couple the executions until their first different completion answer. On a new query z, the threshold Θ_z is independent of the preceding history, and Eq. (12) bounds the disagreement probability by $5|\mathrm{val}_{\mathcal{T}}(z) - \mathrm{val}_{\mathcal{T}'}(z)|$. A repeated query cannot cause the first disagreement. Hence a union bound over the at most Q queries bounds this change by

$$5Q \max_z |\mathrm{val}_{\mathcal{T}}(z) - \mathrm{val}_{\mathcal{T}'}(z)|,$$

where the maximum is over the descriptions allowed by the resource bound. This maximum also covers adaptively chosen descriptions.

Sampler changes. The transitions H_0 to H_1, H_2 to H_3, H_4 to H_5, and H_6 to H_7 each add or remove one sampler term. By Lemma 5, the operator-norm change is exponentially small, uniformly over all tags. The ordinary query hybrid therefore bounds both the change in every protocol value and the direct change seen by A. The preceding completion bound preserves exponential smallness.

Reflection changes. For H_1 to H_2 and H_5 to H_6, the yes samplers no longer contain y. Fix the labels and all oracle data other than ψ_y. The remaining background is independent of ψ_y. Apply Theorem 4 to a verifier that first receives a protocol description z from its prover and then executes that protocol. The description and subsequent prover messages have polynomial total length, and the verifier makes polynomially many queries. Its prover may choose both z and the later replies as a function of ψ_y. For each z, the difference between optimal values is at most the maximum acceptance-probability difference for the same deterministic prover. The bound on this maximum in Theorem 4 therefore gives

$$\mathbb{E}_{\psi_y}\left[\max_z |\mathrm{val}_{\mathcal{T}}(z) - \mathrm{val}_{\mathcal{T}'}(z)|\right] \le 2^{-\Omega(n)}.$$

Thus the completion change is exponentially small. For the direct oracle change, reverse the comparison if necessary and keep the completion belonging to the world in which the reflection has been removed. This completion and the punctured samplers are independent of ψ_y, so Theorem 4 with no prover messages also bounds the direct change by $2^{-\Omega(n)}$.

The middle switch. Put $S = L_n \setminus \{y\}$. Conditioned on S, the ordered pair $(y, \bar{y})$ is uniform over distinct elements of S^c. In H_3 and H_4, the testing and yes-sampling components use exactly the labels in S, and the no sampler is uniform over S^c. The modified oracle and its completion therefore do not depend on which element of the pair was originally in L_n. Exchanging y and $\bar{y}$ leaves their joint distribution unchanged, so H_3 and H_4 have identical output distributions.

The triangle inequality now gives

$$|\Pr[H_0 = 1] - \Pr[H_7 = 1]| \le 2^{-\Omega(n)}. \tag{4}$$

These probabilities include the choice of the base oracle and threshold table. To obtain a fixed realization, write p_Y and p_N for A's yes- and no-input acceptance probabilities for fixed $(\mathcal{T}, \Theta)$. Apply Eq. (4) to an algorithm that runs A on its

challenge and independently on an internally sampled yes-label, accepting only if both runs accept. Repeat with an internally sampled no-label. This bounds the absolute values of $\mathbb{E}[p_Y^2 - p_Y p_N]$ and $\mathbb{E}[p_Y p_N - p_N^2]$ by $2^{-\Omega(n)}$. Subtracting gives

$$\mathbb{E}_{\mathcal{T},\Theta}[(p_Y - p_N)^2] \leq 2^{-\Omega(n)}.$$

Markov's inequality and the Borel–Cantelli lemma imply that $|p_Y - p_N|$ is negligible with probability one over $(\mathcal{T}, \Theta)$. There are countably many uniform QPT algorithms, so we can fix a realization for which this holds for every such algorithm.

Finally, polynomially many independent samples follow by the usual hybrid over the input coordinates. A single-challenge distinguisher chooses a uniformly random coordinate for its challenge and samples the preceding coordinates from the yes distribution and the following coordinates from the no distribution, using the public sampling oracles. Its advantage is the original distinguishing advantage divided by the number of samples. Thus a non-negligible advantage on polynomially many samples would contradict the single-sample bound for the fixed realization. This proves the remaining requirement of Definition 3. ■

References

1. Aaronson, S.: Quantum copy-protection and quantum money. In: 24th Annual IEEE Conference on Computational Complexity (CCC 2009), pp. 229–242. IEEE (2009). https://doi.org/10.1109/CCC.2009.42
2. Aaronson, S., Christiano, P.: Quantum money from hidden subspaces. Theory of Computing **9**, 349–401 (2013). https://doi.org/10.4086/toc.2013.v009a009
3. Aaronson, S., Kuperberg, G.: Quantum versus classical proofs and advice. Theory of Computing **3**, 129–157 (2007). https://doi.org/10.4086/toc.2007.v003a007
4. Aharonov, D., Naveh, T.: Quantum NP—a survey. arXiv:quant-ph/0210077 (2002). https://arxiv.org/abs/quant-ph/0210077
5. Anshu, A., Haferkamp, J., Hwang, Y., Nguyen, Q.T.: On the complexity of unique quantum witnesses and quantum approximate counting. In: 17th Innovations in Theoretical Computer Science Conference (ITCS 2026), LIPIcs, vol. 362, pp. 10:1–10:20. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2026). https://doi.org/10.4230/LIPIcs.ITCS.2026.10
6. Badertscher, C., Cojocaru, A., Colisson, L., Kashefi, E., Leichtle, D., Mantri, A., Wallden, P.: Security limitations of classical-client delegated quantum computing. In: Advances in Cryptology—ASIACRYPT 2020, LNCS, vol. 12492, pp. 667–696. Springer (2020). https://doi.org/10.1007/978-3-030-64834-3_23
7. Barhoush, M., Morimae, T., Movassagh, R.: Separating quantum indistinguishability obfuscation from falsifiable assumptions. arXiv:2608.25195 (2026). https://arxiv.org/abs/2608.25195
8. Barhoush, M., Morimae, T., Nishimaki, R., Yamakawa, T.: Separating non-interactive classical verification of quantum computation from falsifiable assumptions. In: Advances in Cryptology—CRYPTO 2026, LNCS, vol. 16804, pp. 279–310. Springer (2026). https://doi.org/10.1007/978-3-032-35409-9_10
9. Ben-David, S., Kundu, S.: Oracle separation of QMA and QCMA with bounded adaptivity. In: 51st International Colloquium on Automata, Languages,

and Programming (ICALP 2024), LIPIcs, vol. 297, pp. 21:1–21:18. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2024). https://doi.org/10.4230/LIPIcs.ICALP.2024.21
10. Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.: Strengths and weaknesses of quantum computing. SIAM Journal on Computing **26**(5), 1510–1523 (1997). https://doi.org/10.1137/S0097539796300933
11. Bennett, C.H., Brassard, G., Crépeau, C., Jozsa, R., Peres, A., Wootters, W.K.: Teleporting an unknown quantum state via dual classical and Einstein–Podolsky–Rosen channels. Physical Review Letters **70**(13), 1895–1899 (1993). https://doi.org/10.1103/PhysRevLett.70.1895
12. Bennett, C.H., DiVincenzo, D.P., Shor, P.W., Smolin, J.A., Terhal, B.M., Wootters, W.K.: Remote state preparation. Physical Review Letters **87**(7), 077902 (2001). https://doi.org/10.1103/PhysRevLett.87.077902
13. Bennett, C.H., Hayden, P., Leung, D.W., Shor, P.W., Winter, A.: Remote preparation of quantum states. IEEE Transactions on Information Theory **51**(1), 56–74 (2005). https://doi.org/10.1109/TIT.2004.839476
14. Bostanci, J., Haferkamp, J., Nirkhe, C., Zhandry, M.: Separating QMA from QCMA with a classical oracle. In: Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026), pp. 595–606. ACM (2026). https://doi.org/10.1145/3798129.3800776
15. Bostanci, J., Huang, A., Vaikuntanathan, V.: Separating quantum and classical advice with good codes. arXiv:2602.09385 (2026). https://arxiv.org/abs/2602.09385
16. Buhrman, H., Le Gall, F., Weggemans, J.: Classical versus quantum queries in quantum PCPs with classical proofs. arXiv:2411.00946 (2024). https://arxiv.org/abs/2411.00946
17. Çakan, A.: Separating QMA from QCIP with a classical oracle, or, the power of quantum proofs over classical interaction for quantum verifiers. arXiv:2610.03648 (2026). https://arxiv.org/abs/2610.03648
18. Çakan, A., Goyal, V., Shmueli, O.: Public-key quantum fire and key-fire from classical oracles. In: Advances in Cryptology—CRYPTO 2026, LNCS, vol. 16804, pp. 341–371. Springer (2026). https://doi.org/10.1007/978-3-032-35409-9_12
19. Casper, O., Nehoran, B., Sattath, O.: Conversable quantum fire in the standard model. arXiv:2609.36335 (2026). https://arxiv.org/abs/2609.36335
20. Farhi, E., Gosset, D., Hassidim, A., Lutomirski, A., Nagaj, D., Shor, P.: Quantum state restoration and single-copy tomography for ground states of Hamiltonians. Physical Review Letters **105**(19), 190503 (2010). https://doi.org/10.1103/PhysRevLett.105.190503
21. Fefferman, B., Kimmel, S.: Quantum vs. classical proofs and subset verification. In: 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018), LIPIcs, vol. 117, pp. 22:1–22:23. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2018). https://doi.org/10.4230/LIPIcs.MFCS.2018.22
22. Gavinsky, D.: Classical interaction cannot replace a quantum message. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008), pp. 95–102. ACM (2008). https://doi.org/10.1145/1374376.1374393
23. Grewal, S., Kretschmer, W.: Unentanglement and post-measurement branching in quantum interactive proofs. arXiv:2509.15319 (2025). https://arxiv.org/abs/2509.15319
24. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC 1996), pp. 212–219. ACM (1996). https://doi.org/10.1145/237814.237866

25. Irani, S., Natarajan, A., Nirkhe, C., Rao, S., Yuen, H.: Quantum search-to-decision reductions and the state synthesis problem. In: 37th Computational Complexity Conference (CCC 2022), LIPIcs, vol. 234, pp. 5:1–5:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2022). https://doi.org/10.4230/LIPIcs.CCC.2022.5
26. Jain, R., Ji, Z., Upadhyay, S., Watrous, J.: QIP = PSPACE. Journal of the ACM **58**(6), Article 30, 27 pages (2011). https://doi.org/10.1145/2049697.2049704
27. Jain, R., Kerenidis, I., Kuperberg, G., Santha, M., Sattath, O., Zhang, S.: On the power of a unique quantum witness. Theory of Computing **8**, 375–400 (2012). https://doi.org/10.4086/toc.2012.v008a017
28. Kitaev, A.Y., Shen, A.H., Vyalyi, M.N.: Classical and Quantum Computation. Graduate Studies in Mathematics, vol. 47. American Mathematical Society, Providence (2002). https://doi.org/10.1090/gsm/047
29. Kitaev, A., Watrous, J.: Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In: Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC 2000), pp. 608–617. ACM (2000). https://doi.org/10.1145/335305.335387
30. Kundu, S., Lalonde, O.: Near-optimal entanglement-communication tradeoffs for remote state preparation. arXiv:2602.09428 (2026). https://arxiv.org/abs/2602.09428
31. Li, X., Liu, Q., Pelecanos, A., Yamakawa, T.: Classical vs quantum advice and proofs under classically-accessible oracle. In: 15th Innovations in Theoretical Computer Science Conference (ITCS 2024), LIPIcs, vol. 287, pp. 72:1–72:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2024). https://doi.org/10.4230/LIPIcs.ITCS.2024.72
32. Liu, Q., Mutreja, A., Yuen, H.: QMA vs QCMA and pseudorandomness. In: Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025), pp. 1520–1531. ACM (2025). https://doi.org/10.1145/3717823.3718296
33. Mahadev, U.: Classical verification of quantum computations. SIAM Journal on Computing **51**(4), 1172–1229 (2022). https://doi.org/10.1137/20M1371828
34. Natarajan, A., Nirkhe, C.: A distribution testing oracle separation between QMA and QCMA. Quantum **8**, 1377 (2024). https://doi.org/10.22331/q-2024-06-17-1377
35. Nehoran, B., Zhandry, M.: A computational separation between quantum no-cloning and no-telegraphing. In: 15th Innovations in Theoretical Computer Science Conference (ITCS 2024), LIPIcs, vol. 287, pp. 82:1–82:23. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2024). https://doi.org/10.4230/LIPIcs.ITCS.2024.82
36. Shamir, A.: IP = PSPACE. Journal of the ACM **39**(4), 869–877 (1992). https://doi.org/10.1145/146585.146609
37. Zhandry, M.: Toward separating QMA from QCMA with a classical oracle. In: 16th Innovations in Theoretical Computer Science Conference (ITCS 2025), LIPIcs, vol. 325, pp. 95:1–95:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2025). https://doi.org/10.4230/LIPIcs.ITCS.2025.95