# Quantum Versus Classical Proofs and Advice 

Scott Aaronson* Greg Kuperberg ${ }^{\dagger}$


#### Abstract

This paper studies whether quantum proofs are more powerful than classical proofs, or in complexity terms, whether QMA = QCMA. We prove two results about this question. First, we give a "quantum oracle separation" between QMA and QCMA. More concretely, we show that any quantum algorithm needs $\Omega\left(\sqrt{\frac{2^{n}}{m+1}}\right)$ queries to find an $n$-qubit "marked state" $|\psi\rangle$, even if given an $m$-bit classical description of $|\psi\rangle$ together with a quantum black box that recognizes $|\psi\rangle$. We also prove a matching upper bound. Second, we show that, in the one previously-known case where quantum proofs seemed to help, classical proofs are basically just as powerful. In particular, Watrous gave a QMA protocol for verifying non-membership in finite groups. Under plausible group-theoretic assumptions, we give a QCMA protocol for the same problem. Even with no assumptions, our protocol makes only polynomially many queries to the group oracle. Both of our results apply equally to the problem of quantum versus classical advice-that is, of whether BQP/qpoly equals BQP/poly. We end with some conjectures about quantum versus classical oracles, and about the problem of achieving a classical oracle separation between QMA and QCMA.


## 1 Introduction

If someone hands you a quantum state, is that more "useful" than being handed a classical string with a comparable number of bits? In particular, are there truths that you can efficiently verify, and are there problems that you can efficiently solve, using the quantum state but not using the string? These are the questions that this paper addresses, and that it answers in several contexts.

Recall that QMA, or Quantum Merlin-Arthur, is the class of decision problems for which a "yes" answer can be verified in quantum polynomial time, with help from a polynomial-size quantum witness state $|\psi\rangle$. Many results are known about QMA: for example, it has natural complete problems [13], allows amplification of success probabilities [16], and is contained in PP [16].

Yet as Aharonov and Naveh [2] pointed out in 2002, the very definition of QMA raises a fundamental question. Namely: is it really essential that the witness be quantum, or does it suffice for the algorithm verifying the witness to be quantum? To address this question, Aharonov and Naveh defined the class QCMA, or "Quantum Classical MerlinArthur," to be the same as QMA except that now the witness is classical. ${ }^{1}$ We can then ask whether QMA = QCMA. Not surprisingly, the answer is that we don't know.

If we can't decide whether two complexity classes are equal, the usual next step is to construct a relativized world that separates them. Even though relativized worlds are known to be unrealistic in some cases, they would still provide some evidence that the classes are different. But in the case of QMA versus QCMA, even this limited goal has remained elusive.

Closely related to the question of quantum versus classical proofs is that of quantum versus classical advice. Compared to a proof, advice has the advantage that it can be trusted, but the disadvantage that it can't be tailored to a particular input. More formally, let BQP/qpoly be the class of problems solvable in quantum polynomial time, with help from a polynomial-size "quantum advice state" $\left|\psi_{n}\right\rangle$ that depends only on the input length $n$. Then the question is whether BQP/qpoly = BQP/poly, where BQP/poly is the class of problems solvable in quantum polynomial time with help from polynomial-size classical advice. Aaronson [1] showed that $\mathrm{BQP} / \mathrm{qpoly} \subseteq \mathrm{PP} /$ poly, which at least tells us that quantum advice is not "infinitely" more powerful than classical advice. But, like the QMA versus QCMA question, the BQP/qpoly versus BQP/poly question has remained open, with not even an oracle separation known.

[^0]
### 1.1 Our Results

This paper introduces new tools with which to attack QMA versus QCMA and related questions.
First, we achieve an oracle separation between QMA and QCMA, but only by broadening the definition of "oracle." In particular, we introduce the notion of a quantum oracle, which is just an infinite sequence of unitaries $U=\left\{U_{n}\right\}_{n \geq 1}$ that a quantum algorithm can apply in a black-box fashion. Just as a classical oracle models a subroutine to which an algorithm has black-box access, so a quantum oracle models a quantum subroutine, which can take quantum input and produce quantum output. We are able to give a quantum oracle that separates QMA from QCMA:

Theorem 1 There exists a quantum oracle $U$ such that $\mathrm{QMA}^{U} \neq \mathrm{QCMA}^{U}$.
Similarly, there exists a quantum oracle $V$ such that $\mathrm{BQP}^{V} / \mathrm{qpoly} \neq \mathrm{BQP}^{V} /$ poly. Indeed, all of our results about QMA versus QCMA immediately carry over to BQP/qpoly versus BQP/poly-leading us to conjecture that the two questions "stand or fall together," even though we do not know of an implication between them.

Theorem 1 implies that any proof of QMA = QCMA would require "quantumly nonrelativizing techniques"-that is, techniques that are sensitive to the presence of quantum oracles. Any classically nonrelativizing technique is also quantumly nonrelativizing, but we do not know whether the converse holds. ${ }^{2}$

Underlying Theorem 1 is the following lower bound. Suppose a unitary oracle $U_{n}$ acts on $n$ qubits, and suppose there exists a secret $n$-qubit "marked state" $\left|\psi_{n}\right\rangle$ such that $U_{n}\left|\psi_{n}\right\rangle=-\left|\psi_{n}\right\rangle$, but $U_{n}|\varphi\rangle=|\varphi\rangle$ whenever $\left\langle\varphi \mid \psi_{n}\right\rangle=$ 0 . Then even if a quantum algorithm is given $m$ bits of classical advice about $\left|\psi_{n}\right\rangle$, that algorithm still needs $\Omega\left(\sqrt{\frac{2^{n}}{m+1}}\right)$ queries to $U_{n}$ to find $\left|\psi_{n}\right\rangle$. Note that if $m=0$, then this reduces to the standard fact that Grover search in $2^{n}$ dimensions requires $\Omega\left(\sqrt{2^{n}}\right)$ queries. At the other extreme, if $m \approx 2^{n}$ then our bound gives nothing-not surprisingly, since the classical advice might contain explicit instructions for preparing $\left|\psi_{n}\right\rangle$. The key point is that, if $m$ is not exponentially large, then exponentially many queries to $U_{n}$ are needed. In Section 3.1, we show that our lower bound is tight when $m \geq 2 n$, by giving an algorithm that finds $\left|\psi_{n}\right\rangle$ using $O\left(\sqrt{2^{n} / m}\right)$ queries.

Both the upper and lower bounds hinge on geometric results about partitionings of the unit sphere. We can picture an $m$-bit advice string as partitioning the $2^{n}$-dimensional unit sphere into $2^{m}$ "advice regions." The problem is then to analyze how long Grover search takes, if the marked state $\left|\psi_{n}\right\rangle$ is restricted to one of those regions.

Having separated QMA from QCMA by a quantum oracle, we next revisit the question of whether these classes can be separated by a classical oracle. Right now, we know of only one candidate problem for such a separation in the literature: the Group Non-Membership (GNM) problem, which Watrous [22] placed in QMA even though Babai [3] showed that it is not in MA. In this problem, Arthur is given black-box access to a finite group $G$, together with a subgroup $H \leq G$ specified by its generators and an element $x \in G$. Arthur's goal is to verify that $x \notin H$, using a number of group operations polynomial in $\log |G|$. (Note that the group membership problem is in NP by an easy argument.) In Watrous's protocol, the quantum witness is simply an equal superposition $|H\rangle$ over the elements of $H$. Given such a witness, Arthur can check non-membership by comparing the states $|H\rangle$ and $|x H\rangle$, and can similarly check the veracity of $|H\rangle$ by comparing it to $|h H\rangle$, where $h$ is an almost-uniformly random element of $H$.

Evidently a classical proof of non-membership would have to be completely different. Nevertheless, in Section 4 we show the following:

## Theorem 2 GNM has polynomially-bounded QCMA query complexity.

Theorem 2 implies that it is pointless to try to prove a classical oracle separation between QMA and QCMA by lower-bounding the quantum query complexity of Group Non-Membership. If such a separation is possible, then a new approach will be needed.

The idea of the proof of Theorem 2 is that Merlin can "pull the group out of the black box." In other words, he can claim an embedding of a model group $\Gamma$ into $G$. This claim is entirely classical, but verifying it requires solving the Normal Hidden Subgroup Problem (NHSP) in $\Gamma$. This problem has low query complexity by a result of Ettinger, Høyer, and Knill [10], but is not known to be in BQP. In addition, analyzing the description of $\Gamma$ is not known to be computationally efficient. Nonetheless, in Section 4.1 we discuss evidence that NHSP is in BQP and that non-membership for $\Gamma$ is in NP. Based on this evidence, we conjecture the following:

[^1]
## Conjecture 3 GNM is in QCMA.

Given our results in Section 4, the question remains of whether there is some other way to prove a classical oracle separation between QMA and QCMA. In Section 5, we conjecture that the answer is yes:

Conjecture 4 There exists a classical oracle $A$ such that $\mathrm{QMA}^{A} \neq \mathrm{QCMA}^{A}$. Furthermore, this can be proven by exhibiting an oracle problem with polynomial QMA query complexity but exponential QCMA query complexity.

In other words, we conjecture that separating QMA from QCMA by a classical oracle is much easier than separating them in the unrelativized world.

The reason we believe Conjecture 4 is that it seems possible, for many purposes, to "encode" a quantum oracle into a classical one. In Section 5 we explain more concretely what we mean by that, and present some preliminary results. For example, we show that there exists a BQP algorithm that maps an oracle string $A$ to an $n$-qubit pure state $\left|\psi_{A}\right\rangle$, such that if $A$ is uniformly random, then $\left|\psi_{A}\right\rangle$ is (under a suitable metric) close to uniformly random under the Haar measure. On the negative side, we show that any quantum algorithm that applies an $N$-dimensional unitary $U_{A}$ after making a single quantum query to a classical oracle $A$, can apply at most $4^{N}$ distinct unitaries.

We end in Section 6 with some open problems.

## 2 Preliminaries

Throughout this paper, we refer to the set of $N$-dimensional pure states as $\mathbb{C P}^{N-1}$ (that is, complex projective space with $N-1$ dimensions). We use $\operatorname{Pr}$ to denote probability, and E to denote expectation.

We assume familiarity with standard complexity classes such as BQP and MA. For completeness, we now define QMA, QCMA, BQP/qpoly, and BQP/poly.

Definition 5 QMA is the class of languages $L \subseteq\{0,1\}^{n}$ for which there exists a polynomial-time quantum verifier $\mathcal{Q}$ and a polynomial $p$ such that, for all $x \in\{0,1\}^{n}$ :
(i) If $x \in L$ then there exists a $p(n)$-qubit quantum proof $|\varphi\rangle$ such that $\mathcal{Q}$ accepts with probability at least $2 / 3$ given $|x\rangle|\varphi\rangle$ as input.
(ii) If $x \notin L$ then $\mathcal{Q}$ accepts with probability at most $1 / 3$ given $|x\rangle|\varphi\rangle$ as input, for all purported proofs $|\varphi\rangle$.

The class QCMA is defined similarly, except that $|\varphi\rangle$ is replaced by a classical string $z \in\{0,1\}^{p(n)}$.
Definition 6 BQP/qpoly is the class of languages $L \subseteq\{0,1\}^{n}$ for which there exists a polynomial-time quantum algorithm $\mathcal{Q}$, together with a set of states $\left\{\left|\psi_{n}\right\rangle\right\}_{n \geq 1}$ (where $\left|\psi_{n}\right\rangle$ has size $p(n)$ for some polynomial $p$ ), such that for all $x \in\{0,1\}^{n}$ :
(i) If $x \in L$ then $\mathcal{Q}$ accepts with probability at least $2 / 3$ given $|x\rangle\left|\psi_{n}\right\rangle$ as input.
(ii) If $x \notin L$ then $\mathcal{Q}$ accepts with probability at most $1 / 3$ given $|x\rangle\left|\psi_{n}\right\rangle$ as input.

The class $B Q P /$ poly is defined similarly, except that $\left|\psi_{n}\right\rangle$ is replaced by a classical string $a_{n} \in\{0,1\}^{p(n)}$.
Let us now explain what we mean by a "quantum oracle." For us, a quantum oracle is simply an infinite sequence of unitary transformations, $U=\left\{U_{n}\right\}_{n>1}$. We assume that each $U_{n}$ acts on $p(n)$ qubits for some known polynomial $p$. We also assume that given an $n$-bit string as input, a quantum algorithm calls only $U_{n}$, not $U_{m}$ for any $m \neq n$. This assumption is only made for simplicity; our results would go through without it. ${ }^{3}$ When there is no danger of confusion, we will refer to $U_{n}$ simply as $U$.

We now describe the oracle access mechanism. Assume a quantum computer's state has the form

$$
|\Phi\rangle=\sum_{z, b} \alpha_{z, b}|z\rangle|b\rangle\left|\phi_{z, b}\right\rangle,
$$

[^2]where $|z\rangle$ is a workspace register, $|b\rangle$ is a control qubit, and $\left|\phi_{b, z}\right\rangle$ is a $p(n)$-qubit answer register. Then to "query $U_{n}$ " means to apply the $(p(n)+1)$-qubit unitary transformation that maps $|\Phi\rangle$ to
$$
\left|\Phi^{\prime}\right\rangle=\sum_{z}|z\rangle\left(\alpha_{z, 0}|0\rangle\left|\phi_{z, 0}\right\rangle+\alpha_{z, 1}|1\rangle U_{n}\left|\phi_{z, 1}\right\rangle\right) .
$$

Let $\mathcal{C}$ be a quantum complexity class, and let $U=\left\{U_{n}\right\}_{n \geq 1}$ be a quantum oracle. Then by $\mathcal{C}^{U}$, we will mean the class of problems solvable by a $\mathcal{C}$ machine that, given an input of length $n$, can query $U_{n}$ at unit cost as many times as it likes.

In defining the notion of quantum oracle, several choices present themselves that have no counterpart for classical oracles. Even though these choices will not matter for our results, it seems worthwhile to mention them, since they might arise in future work on the subject. First, we implicitly assumed that if we can apply $U$, then we can also apply controlled- $U$ (that is, $U$ conditioned on the control qubit $|b\rangle$ ). Should we make such an assumption? Second, should we assume that if we can apply $U$, then we can also apply $U^{-1}$ ?

Arguably the answer to both questions should be 'yes'—since given a quantum circuit for $U$, we could produce a quantum circuit for controlled- $U$ or $U^{-1}$ in a completely routine way, one that leaves the circuit's overall structure intact. ${ }^{4}$ Still, it would be interesting to know whether disallowing controlled- $U$ or $U^{-1}$ would enable us to prove more quantum oracle separations. (Note that if we disallow these operations, then the set of inequivalent quantum oracles becomes larger.)

Another question is whether we could prove more oracle separations by allowing nonunitary quantum oraclesthat is, oracles that map pure states to mixed states. In this case, as long as the unitary oracle $U$ is not required to come with $U^{-1}$, the answer seems to be no. For given any $n$-qubit quantum operation $\mathcal{E}$, we can construct a $2 n$-qubit unitary operation $U$, whose induced action on the first $n$ qubits is $\mathcal{E}$. This $U$ might potentially reveal information in the second $n$ qubits. However, we should be able to prevent that by composing $U$ with a unitary that "scrambles" the second $n$ qubits (so that they might as well be thrown away), without affecting the first $n$ qubits.

All quantum oracles considered in this paper will be unitary and self-inverse (that is, $U=U^{-1}$ ). Also, while our algorithm in Section 3.1 will need to apply controlled- $U$, that is only for the technical reason that we will define $U$ so that $U|\psi\rangle=-|\psi\rangle$ if $|\psi\rangle$ is the marked state, and $U|\varphi\rangle=|\varphi\rangle$ whenever $\langle\varphi \mid \psi\rangle=0$. If we stipulated instead that $U|\psi\rangle|b\rangle=|\psi\rangle|b \oplus 1\rangle$ and $U|\varphi\rangle|b\rangle=|\varphi\rangle|b\rangle$ whenever $\langle\varphi \mid \psi\rangle=0$, then $U$ alone would suffice.

## 3 Quantum Oracle Separations

The aim of this section is to prove Theorem 1: that there exists a quantum oracle $U$ such that $\mathrm{QMA}^{U} \neq \mathrm{QCMA}^{U}$. The same ideas will also yield a quantum oracle $V$ such that $\mathrm{BQP}^{V} / \mathrm{qpoly} \neq \mathrm{BQP}^{V} /$ poly.

To prove these oracle separations, we first need some lemmas about probability measures on quantum states. Let $\mu$ be the uniform probability measure over $N$-dimensional pure states (that is, over $\mathbb{C P}^{N-1}$ ). The following notion will play a key role in our argument.

Definition 7 For all $p \in[0,1]$, a probability measure $\sigma$ over $\mathbb{C P}^{N-1}$ is called $p$-uniform if $p \sigma \leq \mu$.
Intuitively, a $p$-uniform measure is what we end up with if we start with the uniform prior over all pure states $|\psi\rangle$, and then condition on $\log 1 / p$ bits of classical information about $|\psi\rangle$.

We are interested in the following question: among all $p$-uniform probability measures $\sigma$, which is the one that maximizes $\mathrm{E}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2}\right]$ ? We can think of $\mathbb{C P}^{N-1}$ as a container, which contains a fluid $\sigma$ that is gravitationally attracted to the state $|0\rangle$. Then intuitively, the answer is clear: the way to maximize $\mathrm{E}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2}\right]$ is to "fill the container from the bottom," subject to the density constraint $p \sigma \leq \mu$. In other words, the optimal $\sigma$ should be the uniform measure over the region $\mathcal{R}(p)$ given by $|\langle\psi \mid 0\rangle| \geq h(p)$, where $h(p)$ is chosen so that the volume of $\mathcal{R}(p)$ is a $p$ fraction of the total volume of $\mathbb{C P}^{N-1}$. The following lemma makes this intuition rigorous.
Lemma 8 Among all p-uniform probability measures $\sigma$ over $\mathbb{C P}^{N-1}$, the one that maximizes $\mathrm{E}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2}\right]$ is $\tau(p)$, the uniform measure over the region $\mathcal{R}(p)$ defined above.

[^3]Proof. Since $|\langle\psi \mid 0\rangle|^{2}$ is nonnegative, we can write

$$
\underset{|\psi\rangle \in \sigma}{\mathrm{E}}\left[|\langle\psi \mid 0\rangle|^{2}\right]=\int_{0}^{\infty} \operatorname{Pr}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2} \geq y\right] d y
$$

We claim that setting $\sigma:=\tau(p)$ maximizes the integrand for every value of $y$. Certainly, then, setting $\sigma:=\tau(p)$ maximizes the integral itself as well.

To prove the claim, we consider two cases. First, if $y \leq h(p)^{2}$, then

$$
\operatorname{Pr}_{|\psi\rangle \in \tau(p)}\left[|\langle\psi \mid 0\rangle|^{2} \geq y\right]=1
$$

which is certainly maximal. Second, if $y>h(p)^{2}$, then

$$
\operatorname{Pr}_{|\psi\rangle \in \tau(p)}\left[|\langle\psi \mid 0\rangle|^{2} \geq y\right]=\frac{1}{p} \cdot \operatorname{Pr}_{|\psi\rangle \in \mu}\left[|\langle\psi \mid 0\rangle|^{2} \geq y\right] .
$$

This is maximal as well, since

$$
\operatorname{Pr}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2} \geq y\right] \leq \frac{1}{p} \cdot \operatorname{Pr}_{|\psi\rangle \in \mu}\left[|\langle\psi \mid 0\rangle|^{2} \geq y\right] .
$$

for all $p$-uniform probability measures $\sigma$.
Lemma 8 completely describes the probability measure that maximizes $\operatorname{EX}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2}\right]$, except for one detail: the value of $h(p)$ (or equivalently, the radius of $\mathcal{R}(p)$ ). The next lemma completes the picture.

Lemma 9 For all p,

$$
h(p)=\sqrt{1-p^{1 /(N-1)}}=\Theta\left(\sqrt{\frac{\log 1 / p}{N}}\right)
$$

Proof. We will show that for all $h$,

$$
\operatorname{Pr}_{|\psi\rangle \in \mu}[|\langle\psi \mid 0\rangle| \geq h]=\left(1-h^{2}\right)^{N-1}
$$

where $\mu$ is the uniform probability measure over $\mathbb{C P}^{N-1}$. Setting $p:=\operatorname{Pr}_{|\psi\rangle \in \mu}[|\langle\psi \mid 0\rangle| \geq h]$ and solving for $h$ then yields the lemma.

Let $\vec{z}=\left(z_{0}, \ldots, z_{N-1}\right)$ be a complex vector; then let $\vec{r}=\left(r_{0}, \ldots, r_{N-1}\right)$ and $\vec{\theta}=\left(\theta_{0}, \ldots, \theta_{N-1}\right)$ be real vectors such that $z_{k}=r_{k} e^{i \theta_{k}}$ for each coordinate $k$. Also, let $\mathcal{D}$ be a Gaussian probability measure on $\mathbb{C}^{N}$, with density function

$$
P(\vec{z})=P(\vec{r})=\frac{1}{\pi^{N}} e^{-\|\vec{r}\|_{2}^{2}}
$$

Let $d \vec{r}$ be shorthand for $d r_{0} \cdots d r_{N-1}$. Then we can express the probability that $|\langle\psi \mid 0\rangle| \geq h$ as

$$
\begin{aligned}
& \operatorname{Pr}_{|\psi\rangle \in \mu}[|\langle\psi \mid 0\rangle| \geq h]=\underset{\vec{z} \in \mathcal{D}}{\operatorname{Pr}}\left[\left|z_{0}\right| \geq h\|\vec{z}\|_{2}\right] \\
& =\underset{\vec{r}, \vec{\theta}}{\operatorname{Pr}}\left[r_{0} \geq h\|\vec{r}\|_{2}\right] \\
& =\int_{\vec{r}, \vec{\theta}: r_{0} \geq h\|\vec{r}\|_{2}} P(\vec{r}) r_{0} \cdots r_{N-1} d \vec{r} d \vec{\theta} \\
& =(2 \pi)^{N} \int_{\vec{r}: r_{0} \geq h\|\vec{r}\|_{2}} \frac{1}{\pi^{N}} e^{-\|\vec{r}\|_{2}^{2}} r_{0} \cdots r_{N-1} d \vec{r} \\
& =\int_{r_{1}, \ldots, r_{N-1}=0}^{\infty}\left(\int_{r_{0}=h \sqrt{\frac{r_{1}^{2}+\cdots+r_{N-1}^{2}}{1-h^{2}}}}^{\infty} 2 e^{-r_{0}^{2}} r_{0} d r_{0}\right) 2^{N-1} e^{-r_{1}^{2}-\cdots-r_{N-1}^{2}} r_{1} d r_{1} \cdots r_{N-1} d r_{N-1} \\
& =\int_{r_{1}, \ldots, r_{N-1}=0}^{\infty} e^{-\left(r_{1}^{2}+\cdots+r_{N-1}^{2}\right) \cdot h^{2} /\left(1-h^{2}\right) 2^{N-1} e^{-r_{1}^{2}-\cdots-r_{N-1}^{2}} r_{1} d r_{1} \cdots r_{N-1} d r_{N-1}} \\
& =\int_{r_{1}, \ldots, r_{N-1}=0}^{\infty} 2^{N-1} e^{-\left(r_{1}^{2}+\cdots+r_{N-1}^{2}\right) /\left(1-h^{2}\right)} r_{1} d r_{1} \cdots r_{N-1} d r_{N-1} \\
& =\left(\int_{r=0}^{\infty} 2 e^{-r^{2} /\left(1-h^{2}\right)} r d r\right)^{N-1} \\
& =\left(1-h^{2}\right)^{N-1} \text {. }
\end{aligned}
$$

By combining Lemmas 8 and 9, we can now prove a key fact: that if $|\psi\rangle$ is drawn from a $p$-uniform probability measure, then for every mixed state $\rho$, the squared fidelity between $|\psi\rangle$ and $\rho$ has a small expectation.

Lemma 10 Let $\sigma$ be a p-uniform probability measure over $\mathbb{C P}^{N-1}$. Then for all $\rho$,

$$
\underset{|\psi\rangle \in \sigma}{\mathrm{E}}[\langle\psi| \rho|\psi\rangle]=O\left(\frac{1+\log 1 / p}{N}\right) .
$$

Proof. If $p \leq e^{-\Omega(N)}$ then the lemma is certainly true, so suppose $p \geq e^{-O(N)}$. Since the concluding inequality is linear in $\rho$, we can assume without loss of generality that $\rho$ is a pure state. Indeed, by symmetry we can assume that $\rho=|0\rangle\langle 0|$. So our aim is to upper-bound $\mathrm{E}_{|\psi\rangle \in \sigma}\left[|\langle\psi \mid 0\rangle|^{2}\right]$, where $\sigma$ is any $p$-uniform probability measure. By Lemma 8, we can assume without loss of generality that $\sigma=\tau(p)$ is the uniform measure over all $|\psi\rangle$ such that $|\langle\psi \mid 0\rangle| \geq h(p)$. Then letting

$$
\begin{aligned}
|\psi\rangle & =\alpha_{0}|0\rangle+\cdots+\alpha_{N-1}|N-1\rangle \\
r & =\sqrt{\left|\alpha_{1}\right|^{2}+\cdots+\left|\alpha_{N-1}\right|^{2}}
\end{aligned}
$$

we have

$$
\begin{aligned}
\underset{|\psi\rangle \in \tau(p)}{\mathrm{E}}\left[|\langle\psi \mid 0\rangle|^{2}\right] & =\underset{|\psi\rangle:\left|\alpha_{0}\right| \geq h(p)}{\mathrm{E}}\left[\left|\alpha_{0}\right|^{2}\right] \\
& =\underset{|\psi\rangle: r^{2} \leq 1-h(p)^{2}}{\mathrm{E}}\left[1-r^{2}\right] \\
& =\frac{\int_{0}^{\sqrt{1-h(p)^{2}}} r^{2 N-3}\left(1-r^{2}\right) d r}{\int_{0}^{\sqrt{1-h(p)^{2}}} r^{2 N-3} d r} \\
& =\frac{\left[\frac{r^{2 N-2}}{2 N-2}-\frac{r^{2 N}}{2 N}\right]_{0}^{\sqrt{1-h(p)^{2}}}}{\left[\frac{r^{2 N}}{2 N}\right]_{0}^{\sqrt{1-h(p)^{2}}}} \\
& =\frac{1-\left(1-\frac{1}{N}\right)\left(1-h(p)^{2}\right)}{\left(1-\frac{1}{N}\right)\left(1-h(p)^{2}\right)} \\
& =O\left(\frac{1}{N}+h(p)^{2}\right) \\
& =O\left(\frac{1+\log 1 / p}{N}\right)
\end{aligned}
$$

where the last line follows from Lemma 9.
We are finally ready to prove the main result of this section: that any quantum algorithm needs $\Omega\left(\sqrt{\frac{2^{n}}{m+1}}\right)$ queries to find an $n$-qubit marked state $|\psi\rangle$, even if given $m$ bits of classical advice about $|\psi\rangle$.

Theorem 11 Suppose we are given oracle access to an n-qubit unitary $U$, and want to decide which of the following holds:
(i) There exists an n-qubit "quantum marked state" $|\psi\rangle$ such that $U|\psi\rangle=-|\psi\rangle$, but $U|\phi\rangle=|\phi\rangle$ whenever $\langle\phi \mid \psi\rangle=0$; or
(ii) $U=I$ is the identity operator.

Then even if we have an $m$-bit classical witness $w$ in support of case (i), we still need $\Omega\left(\sqrt{\frac{2^{n}}{m+1}}\right)$ queries to verify the witness, with bounded probability of error.

Proof. If $m=\Omega\left(2^{n}\right)$ then the theorem is certainly true, so suppose $m=o\left(2^{n}\right)$. Let $A$ be a quantum algorithm that queries $U$. Also, let $U_{|\psi\rangle}$ be an $n$-qubit unitary such that $U_{|\psi\rangle}|\psi\rangle=-|\psi\rangle$, but $U_{|\psi\rangle}|\phi\rangle=|\phi\rangle$ whenever $\langle\phi \mid \psi\rangle=0$. Then $A$ 's goal is to accept if and only if $U=U_{|\psi\rangle}$ for some $|\psi\rangle$.

For each $n$-qubit pure state $|\psi\rangle$, let us fix a classical witness $w \in\{0,1\}^{m}$ that maximizes the probability that $A$ accepts, given $U_{|\psi\rangle}$ as oracle. Let $S(w)$ be the set of $|\psi\rangle$ 's associated with a given witness $w$. Since the $S(w)$ 's form a partition of $\mathbb{C P}^{2^{n}-1}$, clearly there exists a witness, call it $w^{*}$, such that

$$
\operatorname{Pr}_{|\psi\rangle \in \mu}\left[|\psi\rangle \in S\left(w^{*}\right)\right] \geq \frac{1}{2^{m}}
$$

Fix that $w^{*}$ (or in other words, hardwire $w^{*}$ into $A$ ). Then to prove the theorem, it suffices to establish the following claim: $A$ cannot distinguish the case $U=U_{|\psi\rangle}$ from the case $U=I$ by making $o\left(\sqrt{\frac{2^{n}}{m+1}}\right)$ queries to $U$, with high probability if $|\psi\rangle$ is chosen uniformly at random from $S\left(w^{*}\right)$.

To prove the claim, we use a generalization of the hybrid argument of Bennett et al. [6]. Suppose that $A$ makes $T$ queries to $U$. (Technically speaking, we should also allow queries to controlled- $U$, but this will make no difference in our analysis.) Then for all $0 \leq t \leq T$, let $\left|\Phi_{t}\right\rangle$ be the final state of $A$, assuming that $U=I$ for the first $t$ queries, and $U=U_{|\psi\rangle}$ for the remaining $T-t$ queries. Thus $\left|\Phi_{0}\right\rangle$ is the final state in case (i), while $\left|\Phi_{T}\right\rangle$ is the final state
in case (ii). We will argue that $\left|\Phi_{t}\right\rangle$ cannot be very far from $\left|\Phi_{t-1}\right\rangle$, with high probability over the choice of marked state $|\psi\rangle$. Intuitively, this is because the computations of $\left|\Phi_{t}\right\rangle$ and $\left|\Phi_{t-1}\right\rangle$ differ in only a single query, and with high probability that query cannot have much overlap with $|\psi\rangle$. We will then conclude, by the triangle inequality, that $\left|\Phi_{0}\right\rangle$ cannot be far from $\left|\Phi_{T}\right\rangle$ unless $T$ is large.

More formally, let $\rho_{t}$ be the marginal state of the query register just before the $t^{t h}$ query, assuming the "control case" $U=I$. Also, let $\rho_{t}=\sum p_{i}\left|\varphi_{i}\right\rangle\left\langle\varphi_{i}\right|$ be an arbitrary decomposition of $\rho_{t}$ into pure states. Then for every $i$, the component of $\left|\varphi_{i}\right\rangle$ orthogonal to $|\psi\rangle$ is unaffected by the $t^{t h}$ query. Therefore

$$
\begin{aligned}
\|\left|\Phi_{t}\right\rangle-\left|\Phi_{t-1}\right\rangle \|_{2} & \leq \sum_{i} p_{i} \cdot 2\left|\left\langle\varphi_{i} \mid \psi\right\rangle\right| \\
& =2 \sum_{i} p_{i} \sqrt{\left\langle\psi \mid \varphi_{i}\right\rangle\left\langle\varphi_{i} \mid \psi\right\rangle} \\
& \leq 2 \sqrt{\sum_{i} p_{i}\left\langle\psi \mid \varphi_{i}\right\rangle\left\langle\varphi_{i} \mid \psi\right\rangle} \\
& =2 \sqrt{\langle\psi| \rho_{t}|\psi\rangle},
\end{aligned}
$$

where the third line uses the Cauchy-Schwarz inequality (the average of the square root is at most the square root of the average). Now let $\sigma$ be the uniform probability measure over $S\left(w^{*}\right)$, and observe that $\sigma$ is $2^{-m}$-uniform. So by Lemma 10,

$$
\begin{aligned}
\underset{|\psi\rangle \in \sigma}{\mathrm{E}}\left[\|\left|\Phi_{t}\right\rangle-\left|\Phi_{t-1}\right\rangle \|_{2}\right] & \leq 2 \underset{|\psi\rangle \in \sigma}{\mathrm{E}}\left[\sqrt{\langle\psi| \rho_{t}|\psi\rangle}\right] \\
& \leq 2 \sqrt{\underset{|\psi\rangle \in \sigma}{\mathrm{E}}\left[\langle\psi| \rho_{t}|\psi\rangle\right]} \\
& \leq 2 \sqrt{\frac{1+\ln \left(1 / 2^{-m}\right)}{2^{n}}} \\
& =O\left(\sqrt{\frac{m+1}{2^{n}}}\right)
\end{aligned}
$$

where the second line again uses the Cauchy-Schwarz inequality. Finally,

$$
\begin{aligned}
\left.\underset{|\psi\rangle \in S\left(w^{*}\right)}{\mathrm{E}}\left[\|| | \Phi_{T}\right\rangle-\left|\Phi_{0}\right\rangle \|_{2}\right] & \leq \sum_{t=1}^{T} \underset{|\psi\rangle \in S\left(w^{*}\right)}{\mathrm{E}}\left[\|\left|\Phi_{t}\right\rangle-\left|\Phi_{t-1}\right\rangle \|_{2}\right] \\
& =O\left(T \sqrt{\frac{m+1}{2^{n}}}\right)
\end{aligned}
$$

by the triangle inequality. This implies that, for $\left|\Phi_{T}\right\rangle$ and $\left|\Phi_{0}\right\rangle$ to be distinguishable with $\Omega$ (1) bias, we must have $T=\Omega\left(\sqrt{\frac{2^{n}}{m+1}}\right)$.

Using Theorem 11, we can immediately show a quantum oracle separation between QMA and QCMA.
Proof of Theorem 1. Let $L$ be a unary language chosen uniformly at random. Then the oracle $U=\left\{U_{n}\right\}_{n \geq 1}$ is defined as follows: if $0^{n} \in L$, then $U_{n}$ negates a randomly chosen marked state $\left|\psi_{n}\right\rangle$; otherwise $U_{n}$ is the identity operation on $n$ qubits. Clearly $L \in \mathrm{QMA}^{U}$; the witness is just $\left|\psi_{n}\right\rangle$ itself. On the other hand, Theorem 11 implies that $L \notin \mathrm{QCMA}^{U}$ with probability 1 . We omit the standard diagonalization argument.

We can similarly show a quantum oracle separation between BQP/qpoly and BQP/poly.
Theorem 12 There exists a quantum oracle $U$ such that $\mathrm{BQP}^{U} / \mathrm{qpoly} \neq \mathrm{BQP}^{U} /$ poly.
Proof. In this case $U_{n}$ will act on $2 n$ qubits. Let $L$ be a binary language chosen uniformly at random, and let $L(x)=1$ if $x \in L$ and $L(x)=0$ otherwise. Also, for all $n$, let $\left|\psi_{n}\right\rangle$ be an $n$-qubit state chosen uniformly at random. Then $U_{n}$ acts as follows: for all $x \in\{0,1\}^{n}$,

$$
U_{n}\left(\left|\psi_{n}\right\rangle|x\rangle\right)=(-1)^{L(x)}\left|\psi_{n}\right\rangle|x\rangle
$$

but

$$
U_{n}(|\phi\rangle|x\rangle)=|\phi\rangle|x\rangle
$$

whenever $\left\langle\phi \mid \psi_{n}\right\rangle=0$. Clearly $L \in \mathrm{BQP}^{U}$ /qpoly; we just take $\left|\psi_{n}\right\rangle$ as the advice. On the other hand, Theorem 11 implies that $L \notin \mathrm{BQP}^{U}$ /poly with probability 1 .

### 3.1 Upper Bound

In this section we show that the lower bound of Theorem 11 is basically tight. In particular, let $U$ be an $n$-qubit quantum oracle, and suppose we are given an $m$-bit classical proof that $U$ is not the identity, but instead conceals a marked state $|\psi\rangle$ such that $U|\psi\rangle=-|\psi\rangle$. Then provided $2 n \leq m \leq 2^{n}$, a quantum algorithm can verify the proof by making $O\left(\sqrt{2^{n} / m}\right)$ oracle calls to $U$. This matches our lower bound when $m \geq 2 n$. ${ }^{5}$

Let $N=2^{n}$ be the dimension of $U$ 's Hilbert space. Then the idea of our algorithm is to use a "mesh" of states $\left|\phi_{1}\right\rangle, \ldots,\left|\phi_{M}\right\rangle \in \mathbb{C P}^{N-1}$, at least one of which has nontrivial overlap with every pure state in $\mathbb{C P}^{N-1}$. A classical proof can then help the algorithm by telling it the $\left|\phi_{i}\right\rangle$ that is closest to $|\psi\rangle$. More formally, define the $h$-ball about $|\phi\rangle$ to be the set of $|\varphi\rangle$ such that $|\langle\phi \mid \varphi\rangle| \geq h$. Then define an $h$-mesh for $\mathbb{C P}^{N-1}$ of size $M$ to be a set of states $\left|\phi_{1}\right\rangle, \ldots,\left|\phi_{M}\right\rangle$ such that every $|\psi\rangle \in \mathbb{C} \mathbb{P}^{N-1}$ is contained in the $h$-ball about $\left|\phi_{i}\right\rangle$ for some $i$. We will rely on the following theorem, which follows from Corollary 1.2 of Böröczky and Wintsche [7].

Theorem 13 ([7]) For all $0 \leq h<1$, there exists an $h$-mesh for $\mathbb{C P}^{N-1}$ of size

$$
O\left(\frac{N^{3 / 2} \log \left(1+N h^{2}\right)}{\left(1-h^{2}\right)^{N}}\right)
$$

Unfortunately, Böröczky and Wintsche do not provide an explicit construction of such an $h$-mesh; they only prove that it exists. ${ }^{6}$ We will return to this issue later.

We now prove the upper bound.
Theorem 14 Suppose we have an n-qubit quantum oracle $U$ such that either (i) $U=U_{|\psi\rangle}$ for some $|\psi\rangle$, or (ii) $U=I$ is the identity operator. Then given an $m$-bit classical witness $w$ in support of case ( $i$ ), there exists a quantum algorithm that verifies the witness using $O\left(\sqrt{2^{n} / m}+1\right)$ queries to $U$, provided $m \geq 2 n$.

Proof. By Theorem 13, there exists an $h$-mesh $\mathcal{S}$ for $\mathbb{C P}^{2^{n}-1}$ of cardinality

$$
|\mathcal{S}|=O\left(\frac{2^{3 n / 2} \log \left(1+2^{n} h^{2}\right)}{\left(1-h^{2}\right)^{2^{n}}}\right)
$$

Setting $|\mathcal{S}|=2^{m}$ gives us

$$
m \leq \frac{3 n}{2}+2^{n} \log \left(\frac{1}{1-h^{2}}\right)+O(1+\log n)
$$

Solving for $h$, we obtain

$$
h \geq \sqrt{\frac{m-3 n / 2-O(1+\log n)}{2^{n}}}
$$

which is $\Omega\left(\sqrt{m / 2^{n}}\right)$ provided $m \geq 2 n$. So there exists a collection of $M=2^{m}$ states, $\left|\phi_{1}\right\rangle, \ldots,\left|\phi_{M}\right\rangle \in \mathbb{C P}^{2^{n}-1}$, such that for every $|\psi\rangle$, there exists an $i$ such that $\left|\left\langle\phi_{i} \mid \psi\right\rangle\right| \geq h$ where $h=\Omega\left(\sqrt{m / 2^{n}}\right)$.

Given an oracle $U=U_{|\psi\rangle}$, the witness $w \in\{0,1\}^{m}$ will simply encode an index $i$ such that $\left|\left\langle\phi_{i} \mid \psi\right\rangle\right| \geq h$. If we prepare $\left|\phi_{i}\right\rangle$ and feed it to $U$, then the probability of finding the marked state $|\psi\rangle$ is $\left|\left\langle\phi_{i} \mid \psi\right\rangle\right|^{2} \geq h^{2}$. Furthermore,

[^4]if we do find $|\psi\rangle$, we will know we did (i.e. a control qubit will be $|1\rangle$ instead of $|0\rangle$ ). From these facts, it follows immediately from the amplitude amplification theorem of Grover [11] and Brassard et al. [8] that we can find $|\psi\rangle$ with probability $\Omega$ (1) using
$$
O\left(\sqrt{\frac{1}{h^{2}}}+1\right)=O\left(\sqrt{\frac{2^{n}}{m}}+1\right)
$$
oracle calls to $U$.
Of course, if we care about computational complexity as well as query complexity, then it is not enough for the $h$-mesh $\left|\phi_{1}\right\rangle, \ldots,\left|\phi_{M}\right\rangle$ to exist. We also need there to be an efficient algorithm that prepares $\left|\phi_{i}\right\rangle$ given $i$. One of us (Kuperberg) has found an explicit construction of $h$-meshes based on BCH codes, for which such an algorithm exists. Details will be provided in a subsequent paper.

## 4 Group Non-Membership

The Group Non-Membership (GNM) problem is defined as follows. We are given a finite group $G$, a subgroup $H \leq G$, and an element $x \in G$. The problem is to decide whether $x \notin H$.

But how are $G, H$, and $x$ specified? To abstract away the details of this question, we will use Babai and Szemerédi's model of black-box groups [4]. In this model, we know generators for $H$, and we know how to multiply and invert the elements of $G$, but we "do not know anything else." More formally, we are given access to a group oracle $\mathcal{O}$, which represents each element $x \in G$ by a randomly-chosen label $\ell(x) \in\{0,1\}^{n}$ for some $n \gg \log _{2}|G|$. We are also given the labels of generators $\left\langle h_{1}, \ldots, h_{l}\right\rangle$ for $H$. We are promised that every element has a unique label.

Suppose that our quantum computer's state has the form

$$
|\Phi\rangle=\sum_{x, y \in G, z} \alpha_{x, y, z}|\ell(x), \ell(y)\rangle|z\rangle
$$

where $\ell(x)$ and $\ell(y)$ are labels of group elements and $|z\rangle$ is a workspace register. Then the oracle $\mathcal{O}$ maps this state to

$$
\mathcal{O}|\Phi\rangle=\sum_{x, y \in G, z} \alpha_{x, y, z}\left|\ell(x), \ell\left(x y^{-1}\right)\right\rangle|z\rangle
$$

Note that if the first register does not contain valid labels of group elements, then $\mathcal{O}$ can behave arbitrarily. Thus, from now on we will ignore labels, and talk directly about the group elements they represent. Using $\mathcal{O}$, it is easy to see that we can perform group inversion (by putting the identity element $e$ in the $x$ register) and multiplication (by first inverting $y$, then putting $y^{-1}$ in the $y$ register), as well as any combination of these operations.

We will show that GNM has polynomially-bounded QCMA query complexity. In other words, if $x \notin H$, then Merlin can provide Arthur with a poly ( $n$ )-bit classical witness of that fact, which enables Arthur to verify it with high probability using poly $(n)$ quantum queries to the group oracle $\mathcal{O}$.

To prove this result, we first need to collect various facts from finite group theory. Call $g_{1}, \ldots, g_{k}$ an efficient generating set for a finite group $G$ if (i) $k=O(\log |G|)$, and (ii) every $x \in G$ is expressible as $g_{1}^{e_{1}} \cdots g_{k}^{e_{k}}$ where $e_{1}, \ldots, e_{k} \in\{0,1\}$. The following lemma follows immediately from a theorem of Erdös and Rényi [9], and can also be proven directly.

Lemma 15 Every finite group $G$ has an efficient generating set.
Given finite groups $\Gamma$ and $G$, we say that functions $f, g: \Gamma \rightarrow G$ are $\varepsilon$-close if

$$
\operatorname{Pr}_{x \in \Gamma}[f(x) \neq g(x)] \leq \varepsilon .
$$

Also, recall that $f: \Gamma \rightarrow G$ is a homomorphism if $f(x y)=f(x) f(y)$ for all $x, y \in \Gamma$. The following two propositions relate $\varepsilon$-closeness to homomorphisms.

Proposition 16 If two homomorphisms $f, g: \Gamma \rightarrow G$ are $(1 / 2-\varepsilon)$-close for any $\varepsilon>0$, then $f=g$.

Proof. Fix $x \in \Gamma$; then for all $y \in \Gamma$, we have $f(x)=f(y) f\left(y^{-1} x\right)$ and $g(x)=g(y) g\left(y^{-1} x\right)$. By the union bound,

$$
\begin{aligned}
\operatorname{Pr}_{y \in \Gamma}\left[f(y)=g(y) \wedge f\left(y^{-1} x\right)=g\left(y^{-1} x\right)\right] & \geq 1-\operatorname{Pr}_{y \in \Gamma}[f(y) \neq g(y)]-\operatorname{Pr}_{y \in \Gamma}\left[f\left(y^{-1} x\right) \neq g\left(y^{-1} x\right)\right] \\
& >0 .
\end{aligned}
$$

Hence there exists a $y$ such that $f(y)=g(y)$ and $f\left(y^{-1} x\right)=g\left(y^{-1} x\right)$. But this implies that $f(x)=g(x)$.
In particular, Proposition 16 implies that if a function $f$ is $1 / 5$-close to a homomorphism, then it is $1 / 5$-close to a unique homomorphism ( $1 / 5$ being an arbitrary constant less than $1 / 4$ ).

Proposition 17 (Ben-Or et al. [5]) Given finite groups $\Gamma$ and $G$, a function $f: \Gamma \rightarrow G$, and a real number $\varepsilon>0$, if

$$
\operatorname{Pr}_{x, y \in \Gamma}[f(x y) \neq f(x) f(y)] \leq \varepsilon
$$

then $f$ is $\varepsilon$-close to a homomorphism.
Together, Propositions 16 and 17 have the following easy corollary.
Corollary 18 Given finite groups $\Gamma$ and $G$ and a function $f: \Gamma \rightarrow G$, there exists a randomized algorithm that makes $O(1)$ oracle queries to $f$, accepts with probability 1 if $f$ is a homomorphism, and rejects with probability at least $2 / 3$ if $f$ is not $1 / 5$-close to a homomorphism. Also, if $f$ is $1 / 5$-close to some homomorphism $\underset{\sim}{f}$, then there exists a randomized algorithm that, given an input $x \in \Gamma$, makes $O(r)$ oracle queries to $f$, and outputs $\widetilde{f}(x)$ with probability at least $1-1 / 2^{r}$.

Proof. The first algorithm simply chooses $O(1)$ pairs $x, y \in \Gamma$ uniformly at random, accepts if $f(x y)=f(x) f(y)$ for all of them, and rejects otherwise. Let $k=O(r)$. Then the second algorithm chooses $z_{1}, \ldots, z_{k} \in \Gamma$ uniformly at random, and outputs the plurality answer among $f\left(z_{1}\right) f\left(z_{1}^{-1} x\right), \ldots, f\left(z_{k}\right) f\left(z_{k}^{-1} x\right)$ (breaking ties arbitrarily).

Interestingly, despite the simplicity of the next result, it is not known how to prove it without using the Classification of Finite Simple Groups.

Proof. Let $F_{\text {simple }}(N)$ be the number of simple groups of order $N$ up to isomorphism. Neumann [18] showed in 1969 that if $F_{\text {simple }}(N)=N^{O\left(\left(\log _{2} N\right)^{2}\right)}$, then $F(N)=N^{O\left(\left(\log _{2} N\right)^{2}\right)}$ as well. Since the Classification of Finite Simple Groups established that $F_{\text {simple }}(N) \leq 2$ (see Lubotzky [15] for example), the theorem follows.

Finally, recall that the Hidden Subgroup Problem (HSP) is defined as follows. We are given a finite group $G$, and oracle access to a function $f: G \rightarrow \mathbb{Z}$. We are promised that there exists a "hidden subgroup" $H \leq G$ such that $f(x)=f(y)$ if and only if $x$ and $y$ belong to the same left coset of $H$. The problem is then to output a set of generators for $H$. Whether HSP can be solved in quantum polynomial time, for various non-abelian groups $G$, is one of the most actively studied questions in quantum computing. However, if we only care about query complexity, then Ettinger, Høyer, and Knill [10] proved the following useful result.

Theorem 20 ([10]) For all finite groups $G$, there exists a quantum algorithm that solves HSP using only polylog $(|G|)$ quantum queries to $f$ (together with a possibly exponential amount of postprocessing).

We can now prove Theorem 2: that GNM has polynomially-bounded QCMA query complexity.
Proof of Theorem 2. Let $G$ be a group of order at most $2^{n}$, and let $\mathcal{O}$ be a group oracle that maps each element of $G$ to an $n$-bit label. Also, given (the labels of) group elements $x, h_{1}, \ldots, h_{m} \in G$, let $H$ be the subgroup of $G$ generated by $\left\langle h_{1}, \ldots, h_{m}\right\rangle$. Then the problem is to decide if $x \notin H$.

In our QCMA protocol for this problem, Merlin's witness will consist of the following:

- An explicit "model group" $\Gamma$, of order at most $2^{n}$.
- A list of elements $\gamma_{1}, \ldots, \gamma_{k} \in \Gamma$, where $k=O(\log |\Gamma|)$.
- A corresponding list $g_{1}, \ldots, g_{k} \in G$.
- Another list $z, \lambda_{1}, \ldots, \lambda_{m} \in \Gamma$.

By Theorem 19, there are at most $2^{\text {poly }(n)}$ groups of order $|\Gamma| \leq 2^{n}$ up to isomorphism. From this it follows that Merlin can specify the witness using only poly $(n)$ bits.

Now if Merlin is honest, then the witness will satisfy the following three properties:
(1) $\gamma_{1}, \ldots, \gamma_{k}$ is an efficient generating set for $\Gamma$.
(2) $z \notin \Lambda$, where $\Lambda$ is the subgroup of $\Gamma$ generated by $\left\langle\lambda_{1}, \ldots, \lambda_{m}\right\rangle$.
(3) There exists an embedding $\tilde{f}: \Gamma \rightarrow G$, such that (i) $\tilde{f}\left(\gamma_{i}\right)=g_{i}$ for all $i \in\{1, \ldots, k\}$, (ii) $\tilde{f}\left(\lambda_{j}\right)=h_{j}$ for all $j \in\{1, \ldots, m\}$, and (iii) $\widetilde{f}(z)=x$.

Suppose for the moment that (1)-(3) all hold. Then there exists an embedding $\tilde{f}: \Gamma \rightarrow G$, which maps the set $\left\langle\gamma_{1}, \ldots, \gamma_{k}\right\rangle$ in $\Gamma$ to the set $\left\langle g_{1}, \ldots, g_{k}\right\rangle$ in $G$. Furthermore, this embedding satisfies $\widetilde{f}(\Lambda)=H$ and $\widetilde{f}(z)=x$. Since $z \notin \Lambda$ by (2), it follows that $x \notin H$ as well, which is what Arthur wanted to check.

So it suffices to verify (1)-(3). In the remainder of the proof, we will explain how to do this using a possibly exponential amount of computation, but only poly $(n)$ quantum queries to the group oracle $\mathcal{O}$.

First, since properties (1) and (2) only involve the explicit group $\Gamma$, not the black-box group $G$, Arthur can verify these properties "free of cost." In other words, regardless of how much computation he needs, he never has to query the group oracle.

The nontrivial part is to verify (3). It will be convenient to split (3) into the following sub-claims:
(3a) There exists a homomorphism $\tilde{f}: \Gamma \rightarrow G$ such that $\tilde{f}\left(\gamma_{i}\right)=g_{i}$ for all $i \in\{1, \ldots, k\}$.
(3b) $\tilde{f}$ satisfies $\tilde{f}(z)=x$ and $\tilde{f}\left(\lambda_{j}\right)=h_{j}$ for all $j \in\{1, \ldots, m\}$.
(3c) $\tilde{f}$ is injective (i.e. is an embedding into $G$ ).
To verify (3a), first Arthur fixes a "canonical representation" of each element $\gamma \in \Gamma$. This representation has the form

$$
\gamma=\gamma_{1}^{e_{1}} \cdots \gamma_{k}^{e_{k}}
$$

where $\left\langle\gamma_{1}, \ldots, \gamma_{k}\right\rangle$ is the efficient generating set for $\Gamma$, and $e_{1}, \ldots, e_{k} \in\{0,1\}$ are bits depending on $\gamma$. Next he defines a function $f: \Gamma \rightarrow G$ by

$$
f(\gamma):=g_{1}^{e_{1}} \cdots g_{k}^{e_{k}}
$$

for all $\gamma \in \Gamma$. By using the canonical representation of $\gamma$, Arthur can evaluate $f(\gamma)$ using at most $k-1$ queries to the group oracle $\mathcal{O}$. Finally Arthur appeals to Corollary 18. If $f$ is not $1 / 5$-close to a homomorphism, then by using $O(1)$ queries to $f$, with high probability Arthur can detect that $f$ is not a homomorphism. In that case Merlin has been caught cheating, so Arthur rejects. On the other hand, if $f$ is $1 / 5$-close to some homomorphism $\widetilde{f}$, then by using $O(\log |\Gamma|)$ queries to $f$, with high probability Arthur can "correct" $f$ to $\widetilde{f}$. In that case it remains only to check that $\widetilde{f}\left(\gamma_{i}\right)=g_{i}$ for all $i \in\{1, \ldots, k\}$.

Once Arthur has an efficient procedure for computing $\tilde{f}$-that is, a procedure that involves only poly $(n)$ queries to $\mathcal{O}$-he can then verify property (3b) directly.

To verify (3c), Arthur runs the algorithm of Ettinger, Høyer, and Knill [10] for the Hidden Subgroup Problem. Notice that, since $\tilde{f}: \Gamma \rightarrow G$ is a homomorphism, there must be a "hidden subgroup" $K \leq \Gamma$ - namely the kernel of $\widetilde{f}$-such that $\tilde{f}$ is constant on cosets of $K$ and distinct on distinct cosets. Furthermore, $\tilde{f}$ is injective if and only if $K$ is trivial. But deciding whether $K$ is trivial is just an instance of HSP, and can therefore be solved using poly ( $n$ ) quantum queries by Theorem 20.

### 4.1 Computational Complexity

Theorem 2 showed that one can always verify group non-membership using a polynomial-size classical witness, together with polynomially many quantum queries to the group oracle $\mathcal{O}$. Unfortunately, while the query complexity is polynomial, the computational complexity might be exponential. However, as mentioned in Section 1.1, we conjecture that this shortcoming of Theorem 2 can be removed, and that GNM is in QCMA for any group oracle $\mathcal{O}$.

In our QCMA protocol, the main computational problem that needs to be solved is not the general HSP, but rather the Normal Hidden Subgroup Problem (NHSP)—that is, HSP where the hidden subgroup is normal. This is because the kernel of a homomorphism is always a normal subgroup. Hallgren, Russell, and Ta-Shma [12] showed that NHSP is in BQP for any explicit group $\Gamma$, provided the quantum Fourier transform over $\Gamma$ can be implemented efficiently. Furthermore, Moore, Rockmore, and Russell [17] showed that many classes of finite groups $G$ have an explicit model $\Gamma \cong G$ for which this assumption holds.

However, even if it can be shown that NHSP is in BQP, there are two remaining obstacles to showing that GNM is in QCMA. First, we need to be able to verify group non-membership in the explicit model group $\Gamma$-possibly with the help of additional classical information from Merlin. And second, we need an efficient algorithm to compute the function $\tilde{f}: \Gamma \rightarrow G$ for every $\gamma \in \Gamma$, even though $\widetilde{f}$ is explicitly defined only on the generators $\gamma_{1}, \ldots, \gamma_{k}$.

More precisely, we need that for every finite group $G$, there should exist an explicit model group $\Gamma \cong G$, together with a list of generators $\gamma_{1}, \ldots, \gamma_{k} \in \Gamma$ with $k=O($ polylog $|G|)$, such that
(i) NHSP over $\Gamma$ is in BQP,
(ii) GNM is over $\Gamma$ in QCMA, and
(iii) Every $\gamma \in \Gamma$ can be efficiently decomposed into a product of $\gamma_{1}, \ldots, \gamma_{k}$.

These steps have already been completed for several classes of groups. For example, if $G$ is abelian, then there exists a model $\Gamma=\mathbb{Z} / r_{1} \times \cdots \times \mathbb{Z} / r_{k}$ for which NHSP is in BQP by the work of Shor [20] and Kitaev [14]; GNM is in P by linear algebra; and the classification of finite abelian groups yields an efficient decomposition. If $G$ is isomorphic to the symmetric group $S_{n}$, then for the model $\Gamma=S_{n}$, we have that NHSP is trivial (since the only normal subgroup is $A_{n}$ ); GNM is in P by the work of Sims [21]; and $S_{n}$ is efficiently generated by transpositions. Indeed, Babai [3] has conjectured that every finite group $G$ has an explicit model group $\Gamma$ for which GNM is in NP $\cap$ coNP. We conjecture that all three steps can be completed-first for finite simple groups, using their classification, and then for arbitrary groups using Jordan-Holder composition series.

## 5 Mimicking Random Quantum Oracles

We have seen, on the one hand, that there exists a quantum oracle separating QMA from QCMA; and on the other hand, that separating these classes by a classical oracle seems much more difficult. Together, these results raise a general question: how much "stronger" are quantum oracles than classical ones? In particular, are there complexity classes $\mathcal{C}$ and $\mathcal{D}$ that can be separated by quantum oracles, but such that separating them by classical oracles is almost as hard as separating them in the unrelativized world? Whatever the answer, we conjecture that QMA and QCMA are not examples of such classes. The reason is that it seems possible, using only classical oracles, to approximate quantum oracles similar to ones that would separate QMA from QCMA.

To illustrate, let $\sigma$ be the uniform probability measure over $2^{n} \times 2^{n}$ unitary diagonal matrices. (In other words, each diagonal entry of $D \in \sigma$ is a random complex number with norm 1.) Also, let $H^{\otimes n}$ be a tensor product of $n$ Hadamard matrices. Then let $\varsigma_{k}$ be the probability measure over $2^{n} \times 2^{n}$ unitary matrices

$$
U=D_{k} H^{\otimes n} D_{k-1} H^{\otimes n} \cdots H^{\otimes n} D_{1} H^{\otimes n}
$$

induced by drawing each $D_{i}$ independently from $\sigma$. In other words, $U \in \varsigma_{k}$ is obtained by first applying a Hadamard gate to each qubit, then a random $2^{n} \times 2^{n}$ diagonal matrix, then Hadamard gates again, then another random diagonal matrix, and so on $k$ times.

Note that we can efficiently apply such a $U$ —at least to polynomially many bits of precision—if given a classical random oracle $A$. To do so, we simply implement the random diagonal matrix $D_{i}$ as

$$
\sum_{x \in\{0,1\}^{n}} \alpha_{x}|x\rangle \rightarrow \sum_{x \in\{0,1\}^{n}} \omega^{A(i, x)} \alpha_{x}|x\rangle,
$$

where $A(i, x)$ is a uniformly random $n$-bit integer indexed by $i$ and $x$, and $\omega=e^{2 \pi i / 2^{n}}$.
Now let $\mu$ be the uniform probability measure over $2^{n} \times 2^{n}$ unitary matrices. If $k \ll 2^{n}$, then $\varsigma_{k}$ is not close to $\mu$ in variation distance, since the former has only $\Theta\left(k 2^{n}\right)$ degrees of freedom while the latter has $\Theta\left(k 4^{n}\right) .{ }^{7}$ On the other hand, we conjecture that a $U$ drawn from $\varsigma_{k}$ will "look random" to any polynomial-time algorithm, and that this property can be used to prove a classical oracle separation between QMA and QCMA.

Let us explain what we mean in more detail. Suppose we are given access to an $n$-qubit unitary oracle $U$, and want to decide whether
(i) $U$ was drawn uniformly at random (that is, from $\mu$ ), or
(ii) $U$ was drawn uniformly at random conditioned on there existing $n / 2$-qubit pure states $|\psi\rangle$ and $|\varphi\rangle$ such that $U\left(|0\rangle^{\otimes n / 2}|\psi\rangle\right) \approx|0\rangle^{\otimes n / 2}|\varphi\rangle$.

In case (i), the states $|\psi\rangle$ and $|\varphi\rangle$ will exist only with negligible probability. ${ }^{8}$ It follows that the above problem is in QMA ${ }^{U}$ —since if case (ii) holds, then a succinct quantum proof of that fact is just $|\psi\rangle$ itself. We now state three conjectures about this problem, in increasing order of difficulty.

Conjecture 21 The above problem is not in QCMA ${ }^{U}$. In other words, if case (ii) holds, there is no succinct classical proof of that fact that can be verified with high probability using poly $(n)$ quantum queries to $U$.

Presumably Conjecture 21 can be proved using ideas similar to those in Section 3. If so, then the next step is to replace the uniform measure $\mu$ by the "pseudorandom" measure $\varsigma_{k}$.

Conjecture 22 Suppose that instead of being drawn from $\mu$, the unitary $U$ is drawn from $\varsigma_{k}$ for some $k=\Omega(n)$. Then the probability that there exist $n / 2-q u b i t$ states $|\psi\rangle$ and $|\varphi\rangle$ such that $U\left(|0\rangle^{\otimes n / 2}|\psi\rangle\right) \approx|0\rangle^{\otimes n / 2}|\varphi\rangle$ is still negligibly small.

Now suppose we want to decide whether
(i') $U$ was drawn from $\varsigma_{k}$, or
(ii') $U$ was drawn from $\varsigma_{k}$ conditioned on there existing $n / 2$-qubit states $|\psi\rangle$ and $|\varphi\rangle$ such that $U\left(|0\rangle^{\otimes n / 2}|\psi\rangle\right) \approx$ $|0\rangle^{\otimes n / 2}|\varphi\rangle$.

Also, let $A$ be a classical oracle that encodes the diagonal matrices $D_{1}, \ldots, D_{k}$ such that

$$
U=D_{k} H^{\otimes n} D_{k-1} H^{\otimes n} \cdots H^{\otimes n} D_{1} H^{\otimes n}
$$

If Conjecture 22 is true, then case (ii') can be verified in QMA $^{A}$. So to obtain a classical oracle separation between QMA and QCMA, the one remaining step would be to prove the following.

Conjecture 23 Case (ii') cannot be verified in QCMA ${ }^{A}$.

### 5.1 From Random Oracles to Random Unitaries

The previous discussion immediately suggests even simpler questions about the ability of classical oracles to mimic quantum ones. In particular, could a BQP machine use a classical random oracle to prepare a uniformly random $n$-qubit pure state? Also, could it use such an oracle to apply a random $n$-qubit unitary?

In this section we answer the first question in the affirmative, and present partial results about the second question. We first need a notion that we call the " $\varepsilon$-smoothing" of a probability measure.

[^5]Definition 24 Let $\sigma$ be a probability measure over $|\psi\rangle \in \mathbb{C P}^{2^{n}-1}$. Then the $\varepsilon$-smoothing of $\sigma$, or $\mathcal{S}_{\varepsilon}(\sigma)$, is the probability measure obtained by first drawing a state $|\psi\rangle$ from $\sigma$, and then drawing a state $|\varphi\rangle$ uniformly at random subject to $\langle\varphi \mid \psi\rangle \geq 1-\varepsilon$.

Let $\mu$ be the uniform measure over $\mathbb{C P} \mathbb{P}^{2^{n}-1}$. Also, let $Q$ be a quantum algorithm that queries a classical oracle $A$. Suppose that, given $0^{n}$ as input, $Q^{A}$ outputs the pure state $\left|\psi_{A}\right\rangle \in \mathbb{C P}^{2^{n}-1}$. Then we say that $Q$ "approximates the uniform measure within $\varepsilon$ " if, as we range over uniform random $A \subseteq\{0,1\}^{n}$, the induced probability measure $\sigma$ over $\left|\psi_{A}\right\rangle$ satisfies $\left\|\mathcal{S}_{\varepsilon}(\sigma)-\mu\right\| \leq \varepsilon$.

Theorem 25 For all polynomials p, there exists a quantum oracle algorithm $Q$ that runs in expected polynomial time, and that approximates the uniform measure within $2^{-p(n)}$.

Proof Sketch. The algorithm $Q$ is as follows: first prepare a uniform superposition over $n$-bit strings. Then, using the classical random oracle $A$ as a source of random bits, map this state to

$$
|\Psi\rangle=\frac{1}{2^{n / 2}} \sum_{x \in\{0,1\}^{n}}|x\rangle\left(\sqrt{1-\left|\alpha_{x}\right|^{2}}|0\rangle+\alpha_{x}|1\rangle\right)
$$

where each $\alpha_{x}$ is essentially a Gaussian random variable. More precisely, let $q(n)=(n+p(n))^{2}$. Then each $\alpha_{x}$ is drawn independently from a complex Gaussian distribution with mean 0 and variance $1 / q(n)$, with the two technicalities that (1) $\alpha_{x}$ is rounded to $q(n)$ bits of precision, and (2) the cutoff $\left|\alpha_{x}\right| \leq 1$ is imposed. (By a tail bound, with overwhelming probability we will have $\left|\alpha_{x}\right| \leq 1$ for all $x$ anyway.)

Next measure the second register of $|\Psi\rangle$ in the standard basis. The outcome $|1\rangle$ will be observed with probability $\Omega(1 / q(n))$. Furthermore, conditioned on $|1\rangle$ being observed, one can check that the distribution $\sigma$ over the reduced state of the first register satisfies $\left\|\mathcal{S}_{2-p(n)}(\sigma)-\mu\right\| \leq 2^{-p(n)}$. (We omit the calculation.) Hence it suffices to repeat the algorithm $O(q(n))$ times.

Theorem 25 shows that, by using a classical random oracle $A$, we can efficiently prepare a uniformly random $n$-qubit state $\left|\psi_{A}\right\rangle$. But what if we want to use a random oracle to apply a uniformly random $n$-qubit unitary $U_{A}$ ? It is clear that we can do this if we have exponential time: given an oracle $A$, we simply query an exponentially long prefix $A^{*}$ of $A$, and then treat $A^{*}$ as an explicit description of a quantum circuit for $U_{A}$. But what if we can make only polynomially many quantum queries to $A$ ? We do not know whether that suffices for applying a random unitary; indeed, we do not even have a conjecture about this.

What we can show is that a single quantum query to $A$ does not suffice for applying a random unitary. In particular, suppose every entry of an $n$-qubit unitary matrix $U_{A}$ is a degree- 1 polynomial in the bits of $A$ (as it must be, if $U_{A}$ is the result of a single quantum query). Then $U_{A}$ can assume at most $4^{2^{n}}$ distinct values as we range over the possible $A$ 's, as opposed to the $\Omega\left(c^{2^{2 n}}\right)$ that would be needed to approximate every $n$-qubit unitary. To prove this statement, we first need a lemma about matrices satisfying a certain algebraic relation.

Lemma 26 Let $E_{1}, \ldots, E_{M}$ be nonzero $N \times N$ matrices over $\mathbb{C}$, and suppose that $E_{i} E_{j}^{\dagger}+E_{j} E_{i}^{\dagger}=0$ for all $i \neq j$. Then $M \leq 2 N$.

Proof. Suppose by contradiction that $M>2 N$. Let $e_{i}^{(k)}$ be vector in $\mathbb{C}^{N}$ corresponding to the $k^{t h}$ row of $E_{i}$. Then the condition $E_{i} E_{j}^{\dagger}+E_{j} E_{i}^{\dagger}=0$ implies that

$$
e_{i}^{(k)} \cdot e_{j}^{(l)}+e_{j}^{(k)} \cdot e_{i}^{(l)}=0
$$

for all $i \neq j$ and $k, l$, where $\cdot$ denotes the complex inner product. Now for all $i$, let $k(i)$ be the minimum $k$ such that $e_{i}^{(k)} \neq 0$, and consider the vectors $e_{1}^{(k(1))}, \ldots, e_{M}^{(k(M))} \in \mathbb{C}^{N}$. Certainly these vectors are not all orthogonal-indeed, since $M>2 N$, there must exist $i \neq j$ such that $\operatorname{Re}\left(e_{i}^{(k(i))} \cdot e_{j}^{(k(j))}\right) \neq 0$. There are now two cases: if $k(i)=k(j)$, then

$$
e_{i}^{(k(i))} \cdot e_{j}^{(k(i))}+e_{j}^{(k(i))} \cdot e_{i}^{(k(i))} \neq 0
$$

and we are done. On the other hand, if $k(i) \neq k(j)$, then

$$
e_{j}^{(k(i))} \cdot e_{i}^{(k(j))}=-e_{i}^{(k(i))} \cdot e_{j}^{(k(j))}
$$

is nonzero. Hence $e_{j}^{(k(i))}$ and $e_{i}^{(k(j))}$ must themselves be nonzero. But if $k(i)>k(j)$, then this contradicts the minimality of $k(i)$, while if $k(i)<k(j)$ then it contradicts the minimality of $k(j)$.

We can now prove the main result.
Theorem 27 Let $U(X)$ be an $N \times N$ matrix, every entry of which is a degree-1 complex polynomial in variables $X=\left(x_{1}, \ldots, x_{k}\right)$. Suppose $U(X)$ is unitary for all $X \in\{0,1\}^{k}$. Then $U(X)$ can assume at most $4^{N}$ distinct values as we range over $X \in\{0,1\}^{k}$.

Proof. By suitable rotation, we can assume without loss of generality that $U\left(0^{k}\right)$ is the $N \times N$ identity $I$. Let $X_{i}$ be the $k$-bit string with a ' 1 ' only in the $i^{t h}$ position, and let $E_{i}:=U\left(X_{i}\right)-I$. Then for all $i$,

$$
\begin{aligned}
E_{i} E_{i}^{\dagger} & =\left(U\left(X_{i}\right)-I\right)\left(U\left(X_{i}\right)^{\dagger}-I^{\dagger}\right) \\
& =I-U\left(X_{i}\right)-U\left(X_{i}\right)^{\dagger}+I \\
& =-E_{i}-E_{i}^{\dagger}
\end{aligned}
$$

Next, for all $i \neq j$, let $X_{i j}$ be the $k$-bit string with ' 1 's only in the $i^{t h}$ and $j^{t h}$ positions. Since $U(X)$ is an affine function of $X$, we have $U\left(X_{i j}\right)=I+E_{i}+E_{j}$. Therefore

$$
\begin{aligned}
0 & =U\left(X_{i j}\right) U\left(X_{i j}\right)^{\dagger}-I \\
& =\left(I+E_{i}+E_{j}\right)\left(I^{\dagger}+E_{i}^{\dagger}+E_{j}^{\dagger}\right)-I \\
& =\left(E_{i} E_{i}^{\dagger}+E_{j} E_{j}^{\dagger}\right)+\left(E_{i} E_{j}^{\dagger}+E_{j} E_{i}^{\dagger}\right)+\left(E_{i}+E_{i}^{\dagger}\right)+\left(E_{j}+E_{j}^{\dagger}\right) \\
& =E_{i} E_{j}^{\dagger}+E_{j} E_{i}^{\dagger}
\end{aligned}
$$

Here the first line uses unitarity, and the fourth line uses the fact that $E_{i}+E_{i}^{\dagger}=-E_{i} E_{i}^{\dagger}$ and $E_{j}+E_{j}^{\dagger}=-E_{j} E_{j}^{\dagger}$. Lemma 26 now implies that there can be at most $2 N$ nonzero $E_{i}$ 's. Hence $U(X)$ can depend nontrivially on at most $2 N$ bits of $X$, and can assume at most $2^{2 N}$ values.

## 6 Open Problems

The most obvious problems left open by this paper are, first, to prove a classical oracle separation between QMA and QCMA, and second, to prove that the Group Non-Membership problem is in QCMA. We end by listing four other problems.

- The class QMA (2) is defined similarly to QMA, except that now there are two quantum provers who are guaranteed to share no entanglement. Can we give a quantum oracle relative to which QMA $(2) \neq$ QMA?
- Can we give a quantum oracle relative to which BQP/qpoly $\not \subset \mathrm{QMA} /$ poly? This would show that Aaronson's containment BQP/qpoly $\subseteq P P /$ poly [1] is in some sense close to optimal.
- Can we use the ideas of Section 5 to give a classical oracle relative to which $\mathrm{BQP} \not \subset \mathrm{PH}$ ? What about a classical oracle relative to which $N P \subseteq B Q P$ but $\mathrm{PH} \not \subset \mathrm{BQP} ?^{9}$
- Is there a polynomial-time quantum oracle algorithm $Q$, such that for every $n$-qubit unitary transformation $U$, there exists a classical oracle $A$ such that $Q^{A}$ approximately implements $U$ ? Alternatively, can we prove that any such algorithm would require more than poly $(n)$ queries to $A ?^{10}$


## 7 Acknowledgments

We thank Dorit Aharonov, Robert Beals, Robert Guralnick, Bill Kantor, and Cris Moore for helpful correspondence.

[^6]
## References

[1] S. Aaronson. Limitations of quantum advice and one-way communication. Theory of Computing, 1:1-28, 2005. quant-ph/0402095.
[2] D. Aharonov and T. Naveh. Quantum NP - a survey. quant-ph/0210077, 2002.
[3] L. Babai. Bounded round interactive proofs in finite groups. SIAM J. Discrete Math, 5(1):88-111, 1992.
[4] L. Babai and E. Szemerédi. On the complexity of matrix group problems I. In Proc. IEEE FOCS, pages 229-240, 1984.
[5] M. Ben-Or, D. Coppersmith, M. Luby, and R. Rubinfeld. Non-abelian homomorphism testing, and distributions close to their self-convolutions. In Proceedings of RANDOM, pages 273-285. Springer-Verlag, 2004. ECCC TR04-052.
[6] C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. SIAM J. Comput., 26(5):1510-1523, 1997. quant-ph/9701001.
[7] K. Böröczky Jr. and G. Wintsche. Covering the sphere by equal spherical balls. In Discrete and Computational Geometry: The Goodman-Pollack Festschrift, pages 237-253. 2003.
[8] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. In S. J. Lomonaco and H. E. Brandt, editors, Quantum Computation and Information, Contemporary Mathematics Series. AMS, 2002. quant-ph/0005055.
[9] P. Erdös and A. Rényi. Probabilistic methods in group theory. J. Analyse Mathématique, 14:127-138, 1965.
[10] M. Ettinger, P. Høyer, and E. Knill. The quantum query complexity of the hidden subgroup problem is polynomial. Inform. Proc. Lett., 91(1):43-48, 2004. quant-ph/0401083.
[11] L. K. Grover. A framework for fast quantum mechanical algorithms. In Proc. ACM STOC, pages 53-62, 1998. quant-ph/9711043.
[12] S. Hallgren, A. Russell, and A. Ta-Shma. The hidden subgroup problem and quantum computation using group representations. SIAM J. Comput., 32(4):916-934, 2003. Conference version in STOC'2000, p. 627-635.
[13] J. Kempe, A. Kitaev, and O. Regev. The complexity of the Local Hamiltonian problem. SIAM J. Comput., 35(5):1070-1097, 2006. quant-ph/0406180.
[14] A. Kitaev. Quantum measurements and the abelian stabilizer problem. ECCC TR96-003, quant-ph/9511026, 1996.
[15] A. Lubotzky. Enumerating boundedly generated finite groups. J. Algebra, 238:194-199, 2001.
[16] C. Marriott and J. Watrous. Quantum Arthur-Merlin games. Computational Complexity, 14(2):122-152, 2005.
[17] C. Moore, D. N. Rockmore, and A. Russell. Generic quantum Fourier transforms. In Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA), pages 778-787, 2004. quant-ph/0304064.
[18] P. M. Neumann. An enumeration theorem for finite groups. Quart. J. Math. Ser., 2(20):395-401, 1969.
[19] A. Shamir. IP=PSPACE. J. ACM, 39(4):869-877, 1992.
[20] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM J. Comput., 26(5):1484-1509, 1997. Earlier version in IEEE FOCS 1994. quant-ph/9508027.
[21] C. Sims. Computational methods in the study of permutation groups. In Computational Problems in Abstract Algebra, pages 169-183. Pergamon Press, 1970.
[22] J. Watrous. Succinct quantum proofs for properties of finite groups. In Proc. IEEE FOCS, pages 537-546, 2000. cs.CC/0009002.


[^0]:    *University of Waterloo. Email: scott@scottaaronson.com.
    ${ }^{\dagger}$ UC Davis. Email: greg@math.ucdavis.edu.
    ${ }^{1}$ Some say that this class would more accurately be called CMQA, for "Classical Merlin Quantum Arthur." But QCMA has apparently stuck.

[^1]:    ${ }^{2}$ Of course, we assume here that we are comparing two quantum complexity classes, for which the notion of a quantum oracle makes sense! Otherwise, one could argue that there exists a quantum oracle relative to which (say) BQP $\not \subset$ PSPACE, since a PSPACE machine cannot access a quantum oracle.

[^2]:    ${ }^{3}$ If one made the analogous assumption in classical complexity-that given an input of length $n$, an algorithm can query the oracle only on strings of length $n$-one could simplify a great many oracle results without any loss of conceptual content.

[^3]:    ${ }^{4}$ One might object that the arithmetization at the heart of the IP $=$ PSPACE theorem [19] also leaves a circuit's "overall structure" intact. But inverting a gate or conditioning it on a control qubit seems less drastic to us than enlarging its base field.

[^4]:    ${ }^{5}$ When $m \ll 2 n$, the best upper bound we know is the trivial $O\left(\sqrt{2^{n}}\right)$. However, we conjecture that $O\left(\sqrt{2^{n} / m}\right)$ is achievable in this case as well.
    ${ }^{6}$ Note that we cannot just start from an explicit construction of a sphere-packing, and then double the radius of the spheres to get a covering. We could do that if we wanted a covering of $\mathbb{C P}^{N-1}$ by small balls. But in our case, $h$ is close to zero, which means that the balls already have close to the maximal radius.

[^5]:    ${ }^{7}$ Admittedly, it is still conceivable that the finite-precision version of $\varsigma_{k}$ is close in variation distance to the finite-precision version of $\mu$. However, a more sophisticated argument that counts distinguishable unitaries rules out that possibility as well.
    ${ }^{8}$ Indeed, the reason we did not ask for $(n-1)$-qubit states $|\psi\rangle$ and $|\varphi\rangle$ such that $U(|0\rangle|\psi\rangle) \approx|0\rangle|\varphi\rangle$ is that such states will exist generically. Asking for $(n-2)$-qubit states $|\psi\rangle$ and $|\varphi\rangle$ such that $U(|00\rangle|\psi\rangle) \approx|00\rangle|\varphi\rangle$ might suffice, but we wish to stay on the safe side.

[^6]:    ${ }^{9}$ Note that a simple relativizing argument shows that if $N P \subseteq B P P$ then $P H \subseteq B P P$.
    ${ }^{10}$ We do not even know whether a single query suffices. Note that Theorem 27 does not apply here, since we have dropped the requirement that $Q^{A}$ must implement some $n$-qubit unitary (as opposed to a more general superoperator) for every oracle $A$.

