Quantum Information and Computation, Vol. 9, No. 1 (2009) 0081–0089


# ON PERFECT COMPLETENESS FOR QMA

SCOTT AARONSON[a]
*Department of Electrical Engineering and Computer Science, MIT*
*Cambridge, MA 02139-4307, USA*
aaronson@csail.mit.edu

Received June 4, 2006
Revised Auguest 12, 2008

Whether the class QMA (Quantum Merlin Arthur) is equal to $\mathsf{QMA}_1$, or QMA with one-sided error, has been an open problem for years. This note helps to explain why the problem is difficult, by using ideas from real analysis to give a "quantum oracle" relative to which $\mathsf{QMA} \neq \mathsf{QMA}_1$. As a byproduct, we find that there are facts about quantum complexity classes that are classically relativizing but not quantumly relativizing, among them such "trivial" containments as $\mathsf{BQP} \subseteq \mathsf{ZQEXP}$.

*Keywords*:

*Communicated by*: R Cleve & J Watrous

## 1 Introduction

The complexity class MA (Merlin-Arthur) was introduced by Babai [1] in 1985. Intuitively, MA is a probabilistic version of NP; it contains all problems for which an omniscient wizard Merlin can convince a probabilistic polynomial-time verifier Arthur of a "yes" answer, by a one-round protocol in which Merlin sends Arthur a purported proof $z$, and then Arthur checks $z$. In the usual definition, if the answer to the problem is "yes" then there should exist a string $z$ that makes Arthur accept with probability at least 2/3 (this property is called *completeness*), while if the answer is "no" then no $z$ should make Arthur accept with probability more than 1/3 (this property is called *soundness*).

One of the first questions people asked about MA was whether it can be made to have *perfect completeness* (also called *one-sided error*): that is, whether the 2/3 in the above definition can be replaced by 1. In other words, can we assume without loss of generality that *Arthur never rejects a valid proof?* This question was answered in the affirmative by Zachos and Fürer [2], using a technique introduced earlier by Lautemann [3] to show $\mathsf{BPP} \subseteq \Sigma_2^{\mathsf{P}}$ (for a different proof see Goldreich and Zuckerman [4]).

A decade ago, Kitaev [5] and Watrous [6] introduced a quantum analogue of MA, called QMA (Quantum Merlin Arthur). Loosely speaking, QMA is the same as MA, except that the verifier Arthur is a polynomial-time *quantum* algorithm, and the proof sent by Merlin is a quantum state $|\psi\rangle$ with polynomially many qubits. We know a reasonable amount about QMA (see Aharonov and Naveh [7] for a survey). Like MA, for example, QMA allows

[a]Supported by MIT and by the Keck Foundation (through xQIT)

exponential amplification of completeness and soundness [8], is contained in PP [8], and has natural complete promise problems [5].

However, the question of whether QMA can be made to have perfect completeness has resisted attack. At first a mere nuisance, this question has increasingly cropped up in quantum complexity theory. For example, two years ago Bravyi [9] defined a quantum analogue of the $k$-SAT problem, and showed it complete for the complexity class $\mathsf{QMA}_1$, meaning QMA with one-sided error. But showing quantum $k$-SAT is QMA-complete would require further showing that $\mathsf{QMA}_1 = \mathsf{QMA}$, or equivalently, that QMA protocols can be made to have perfect completeness. What makes the situation even stranger is that, if we allow multiple rounds of interaction between the prover and verifier (yielding the class QIP), then quantum interactive proof systems *can* be made to have perfect completeness [10].

In this note we help explain this puzzling state of affairs, by giving a *quantum oracle* $\mathcal{U}$ relative to which $\mathsf{QMA}_1^{\mathcal{U}} \neq \mathsf{QMA}^{\mathcal{U}}$. A quantum oracle, as defined by Aaronson and Kuperberg [11], is simply a unitary transformation on quantum states that can be applied in black-box fashion. Our result implies that there is no "black-box method" to convert QMA protocols into $\mathsf{QMA}_1$ protocols, in the same sense that there are black-box methods to convert MA protocols into $\mathsf{MA}_1$ protocols, and to convert QMA protocols into QMA protocols with exponentially small error. If a proof of $\mathsf{QMA}_1 = \mathsf{QMA}$ exists, it will instead have to use "quantumly nonrelativizing techniques": techniques that are sensitive to the presence of quantum oracles.

Somewhat surprisingly, our separation proof has almost nothing to do with complexity theory, and instead hinges on real analysis. Our oracle will act on just a single qubit, and will rotate that qubit by an angle $\theta$ that is either 0 or far from zero. We will show that any $\mathsf{QMA}_1$ protocol to convince a time-bounded verifier that $\theta \neq 0$, using any finite-sized quantum proof, would lead to a matrix $E(\theta)$ that depends analytically on $\theta$, yet whose maximum eigenvalue has the "piecewise" behavior shown in Figure 1. We will then use results from real analysis to show that this behavior cannot occur.

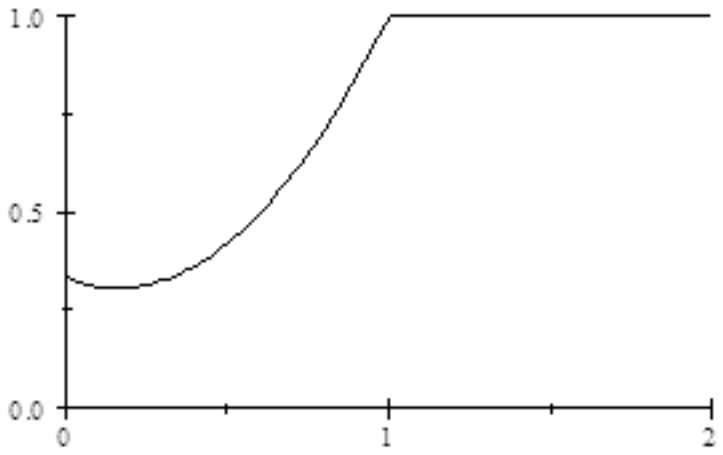


Figure 1: As we vary $\theta$, the largest eigenvalue $a(\theta)$ of the matrix $E(\theta)$ must start out small, but then "plateau" at $a(\theta) = 1$. We will show that this contradicts the analyticity of $E(\theta)$.

Since our argument does not depend on the running time of the $\mathsf{QMA}_1$ machine (so long as it is finite), the same argument will yield quantum oracles $\mathcal{U}$ such that $\mathsf{BQP}^{\mathcal{U}} \not\subset \mathsf{QMA}_1\mathsf{EXP}^{\mathcal{U}}$ and even $\mathsf{BQP}^{\mathcal{U}} \not\subset \mathsf{QMA}_1\mathsf{EEEXP}^{\mathcal{U}}$. This, in turn, has a somewhat surprising implication: that quantum oracles can invalidate even complexity class containments that hold for "trivial" reasons in the unrelativized world. In particular, we will argue that there are extremely simple proof techniques—including the representation of quantum amplitudes by explicit sequences of bits—that are classically relativizing but not quantumly relativizing (at least when applied to one-sided-error complexity classes). Unfortunately, knowing this does not by itself seem to help in finding a proof that $\mathsf{QMA}_1 = \mathsf{QMA}$.

Some argue that any quantum complexity class involving perfect completeness is "inherently unphysical," and we do not wish to dispute this. Indeed, our results could even be taken as further evidence for this point of view. On the other hand, classes like $\mathsf{QIP}$ and $\mathsf{QMA}$ could also be seen as "unphysical" (since there are no Merlins), yet few quantum computing researchers would deny that their study has led to major insights. On a related topic, let us stress that our result does not depend on restricting the set of gates available to the $\mathsf{QMA}_1$ machine: it works assuming *any* countable set of gates. The key issue, then, is not any limitation of the gate basis, but simply the underlying requirement of perfect completeness.

The rest of the paper is organized as follows. Section 2 reviews some preliminaries from complexity theory and real analysis, Section 3 proves the main result, Section 4 discusses the implications for quantum oracles, and Section 5 concludes with some extensions and open problems.

## 2 Preliminaries

In what follows, we assume familiarity with standard complexity classes such as $\mathsf{QMA}$ and $\mathsf{MA}$. See the Complexity Zoo[b] for definitions. For completeness, we now define the class $\mathsf{QMA}_1$, or $\mathsf{QMA}$ with one-sided error.

**Definition 1** *A language $L \subseteq \{0,1\}^*$ is in $\mathsf{QMA}_1$ if there exists a uniform polynomial-size quantum circuit family $\{C_n\}_{n\geq 1}$, and a polynomial $p$, such that for all inputs $x \in \{0,1\}^n$:*

- *(Perfect Completeness) If $x \in L$, then there exists a $p(n)$-qubit quantum witness $|\varphi\rangle$ such that $C_n(x, |\varphi\rangle)$ accepts with certainty.*
- *(Constant Soundness) If $x \notin L$, then $C_n(x, |\varphi\rangle)$ accepts with probability at most $1/2$ for all $|\varphi\rangle$.*

One can similarly define $\mathsf{MA}_1$ as the class of languages $L$ for which there exists a randomized polynomial-time algorithm $A$ such that for all inputs $x \in \{0,1\}^n$, if $x \in L$ then there exists a witness $w \in \{0,1\}^{p(n)}$ such that $A(x,w)$ accepts with certainty, while if $x \notin L$ then $A(x,w)$ accepts with probability at most $1/2$ for all $w$. As mentioned before, the result of Zachos and Fürer [2] implies that $\mathsf{MA}_1 = \mathsf{MA}$, whereas we do not know whether $\mathsf{QMA}_1 = \mathsf{QMA}$.

Note that, because of the perfect completeness condition, the definition of $\mathsf{QMA}_1$ might depend on the particular basis of gates used to generate $C_n$. Indeed, the natural way to show that $\mathsf{QMA}_1$ does *not* depend of the basis of gates would presumably be to show that

[b] http://www.complexityzoo.com

$\mathsf{QMA}_1 = \mathsf{QMA}$, the very task for which we are pointing out an obstacle! For our purposes, though, we can take *any* countable set of 1- and 2-qubit gates as the gate basis of the $\mathsf{QMA}_1$ machine, regardless of how many bits are needed to describe those gates. For example, we could take the set of all 1- and 2-qubit gates that are computably describable. Our separation results will still go through, and such an assumption can only make our results stronger. (Furthermore, the only reason the gate basis needs to be countable is so that a diagonalization argument will go through. If the quantum oracle $U$ could be chosen *subsequent* to the choice of $\mathsf{QMA}_1$ machine, then we could even handle 1- and 2-qubit gates with arbitrary complex-valued transition probabilities.)

Following Aaronson and Kuperberg [11], we define a quantum oracle $\mathcal{U}$ to be simply a collection of unitary operations $\{U_n\}_{n \geq 1}$, where each $U_n$ acts on some number of qubits $q(n)$ (in this note $q(n)$ will always be 1). Let $\mathsf{C}$ be a quantum complexity class. Then by $\mathsf{C}^{\mathcal{U}}$, we mean the class of problems solvable by a $\mathsf{C}$ machine that can, at any time step, apply any $U_n \in \mathcal{U}$ to any subset of its qubits at unit cost. While this is admittedly an informal definition, for any $\mathsf{C}$ of interest to us it is easy to give a reasonable formalization. While there *are* ambiguities in defining $\mathsf{C}^{\mathcal{U}}$, none of those ambiguities will turn out to matter for us. For example, we can assume (if we like) that a $\mathsf{C}^{\mathcal{U}}$ machine is also able to apply $U_n^{-1}$ and controlled-$U_n$, possibly with different values of $n$ in different branches of a superposition. None of these decisions will affect our results.

We now turn to reviewing some facts from real analysis. Recall that a function $f : R \rightarrow R$ is called *real analytic* if for every $x_0 \in R$, the Taylor series about $x_0$ is convergent and equal to $f(x)$ for all $x$ close enough to $x_0$. Every real analytic function is smooth, but the converse does not hold.

We will need the following theorem of Alekseevsky et al. (Theorem 5.1 in [12]):

**Theorem 1 ([12])** *Let*

$$p(\theta)(x) = b_0(\theta) + b_1(\theta)\, x + b_2(\theta)\, x^2 + \cdots + b_N(\theta)\, x^N$$

*be a real polynomial in $x$ with all real roots, parameterized by $\theta \in R$. Suppose the coefficients $b_0(\theta), \ldots, b_N(\theta)$ are all real analytic functions of $\theta$. Then there exist real analytic functions $\lambda_1(\theta), \ldots, \lambda_N(\theta)$ such that $\{\lambda_1(\theta), \ldots, \lambda_N(\theta)\}$ is the set of roots of $p(\theta)(x)$ for all $\theta \in R$.*

We also need the following basic fact:

Let $f : R \rightarrow R$ be a real analytic function. If there exists an open set $(x, y) \subset R$ on which $f$ is constant, then $f$ is constant everywhere.

Note that Proposition 2 is false with smooth functions in place of real analytic ones.[c] This is why we need analyticity for our result.

## 3 Result

We first need a more-or-less standard fact (proved for completeness), which recasts the problem of finding an optimal $\mathsf{QMA}$ witness as a principal eigenvalue problem.

**Lemma 1** *Let $V$ be a quantum verifier that takes as input a $Q$-qubit quantum witness $|\varphi\rangle$, and that makes $T$ queries to a quantum oracle described by a unitary matrix $U$. Also, let $a(U)$ be the acceptance probability of $V^U$ maximized over all possible $|\varphi\rangle$. Then there exists a $2^Q \times 2^Q$ complex-valued matrix $E(U)$ such that*

[c] The standard counterexample is $f(x) = 0$ for $x \leq 0$ and $f(x) = e^{-1/x^2}$ for $x > 0$.

*(i) Every entry of $E(U)$ is a polynomial in the entries of $U$, of degree at most $2T$.*

*(ii) $E(U)$ is Hermitian for all $U$.*

*(iii) $a(U)$ equals the largest eigenvalue of $E(U)$, for all $U$.*

**Proof.** L .et $a(U, |\varphi\rangle)$ be the acceptance probability of $V^U$ on input $|\varphi\rangle$. Then clearly there exist vectors $\{|v_i\rangle\}_{i=1}^{2^Q}$ (not necessarily normalized, and depending on $U$) such that

$$a(U, |\varphi\rangle) = \sum_i |\langle v_i|\varphi\rangle|^2 .$$

Furthermore, by an observation of Beals et al. [13], every entry of every $|v_i\rangle$ must be a polynomial in the entries of $U$, of degree at most $T$. (This is because initially the entries are degree-0 polynomials, and every query to the oracle can increase the degree by at most 1.) So if we set $E := \sum_i |v_i\rangle \langle v_i|$, then $E$ is a $2^Q \times 2^Q$ Hermitian matrix, every entry of which is a polynomial of degree at most $2T$. Furthermore $a(U, |\varphi\rangle) = \langle\varphi|E|\varphi\rangle$, which implies that

$$a(U) = \max_{|\varphi\rangle} a(U, |\varphi\rangle) = \max_{|\varphi\rangle} \langle\varphi|E|\varphi\rangle$$

which is just the largest eigenvalue of $E$.

We now prove the main result.

**Theorem 2** *There exists a quantum oracle $\mathcal{U}$ such that $\mathsf{QMA}_1^{\mathcal{U}} \neq \mathsf{QMA}^{\mathcal{U}}$.*

**Proof.** L .et $\theta$ be a real number, and let $U = U(\theta)$ be the one-qubit unitary transformation

$$\begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}.$$

Given oracle access to $U$, we consider the problem of deciding whether $\theta = 0$ (the NO case) or $1 \leq \theta \leq 2$ (the YES case), promised that one of these holds. Of course this problem is easily solved by a quantum computer, with bounded error probability, using $O(1)$ queries to $U$. On the other hand, we will show that this problem does not admit a perfect-completeness $\mathsf{QMA}$ protocol, with any finite number of queries to $U$ and any finite-sized quantum proof.

To see this, let $V$ be a verifier, let $T$ be the number of queries that $V$ makes to $U$, and let $Q$ be the number of qubits in $V$'s quantum witness. Also, let $a(\theta)$ be the acceptance probability of $V$ assuming $U = U(\theta)$, maximized over all $Q$-qubit quantum witnesses $|\varphi\rangle$. Then by Lemma 1, there exists a $2^Q \times 2^Q$ complex-valued matrix $E(\theta)$ such that

(i) Every entry of $E(\theta)$ is a polynomial in $\cos\theta$ and $\sin\theta$, of degree at most $2T$.

(ii) $E(\theta)$ is Hermitian for all $\theta \in R$.

(iii) $a(\theta)$ equals the largest eigenvalue of $E(\theta)$, for all $\theta \in R$.

Let $N = 2^Q$, and let $\lambda_1(\theta), \ldots, \lambda_N(\theta)$ be the eigenvalues of $E(\theta)$. Then the $\lambda_i(\theta)$'s are roots of a degree-$N$ characteristic polynomial parameterized by $\theta$:

$$p(\theta)(x) = b_0(\theta) + b_1(\theta)\, x + b_2(\theta)\, x^2 + \cdots + b_N(\theta)\, x^N .$$

Each coefficient $b_j(\theta)$ is a polynomial in the entries of $E(\theta)$ of degree at most $N$, and hence, by (i), a polynomial in $\cos\theta$ and $\sin\theta$ of degree at most $2TN$. By (ii), the $\lambda_i(\theta)$'s are all real, and therefore the $b_j(\theta)$'s must be real as well for all $\theta$. Combining these facts, we find that each $b_j(\theta)$ is a real analytic function of $\theta$ (for note that $\cos\theta$ and $\sin\theta$ are real analytic functions, and real analytic functions are closed under composition). By Theorem 1, then, we can take $\lambda_1(\theta), \ldots, \lambda_N(\theta)$ to be real analytic functions as well.

By (iii), the acceptance probability $a(\theta)$ of $V$ (maximized over all witnesses) is equal to $\max_i \lambda_i(\theta)$. If $V$ is a valid $\mathsf{QMA}_1$ verifier, then we must have $a(0) \leq 1/2$, but $a(\theta) = 1$ for all real $1 \leq \theta \leq 2$. Since $N$ is finite and the $\lambda_i(\theta)$'s are continuous, this implies that there exists an $i \in [N]$ such that $\lambda_i(0) \leq 1/2$, but $\lambda_i(\theta) = 1$ for all $\theta$ in some open interval $(x, y) \subset [1, 2]$. But this contradicts the analyticity of $\lambda_i$ by Proposition 2. Hence there must be a choice of $\theta$ such that $V$ does not solve the problem correctly given $U = U(\theta)$ as oracle.

We now simply diagonalize over all $n$ to achieve the desired oracle separation. More formally, let $\mathcal{U}$ be a collection of quantum oracles $U_1, U_2, \ldots$, such that $U_n = U(\theta_n)$ rotates by the angle $\theta_n \in \{0\} \cup [1, 2]$. Also, let $L$ be a unary language such that $0^n \in L$ if and only if $\theta_n \neq 0$. Then clearly $L \in \mathsf{BQP}^{\mathcal{U}}$, and hence $L \in \mathsf{QMA}^{\mathcal{U}}$, for all choices of $\mathcal{U}$. On the other hand, we claim that $\mathcal{U}$ can be chosen so that $L \notin \mathsf{QMA}_1^{\mathcal{U}}$. To see this, let $M_1, M_2, \ldots$ be an enumeration of $\mathsf{QMA}_1$ machines. Then for each $i$, we simply choose an $n_i$ so large that $U_{n_i}$ cannot have been queried by machines $M_1, \ldots, M_{i-1}$, and then set $\theta_{n_i}$ so that $M_i$ fails on input $0^{n_i}$. (In other words, either $\theta_{n_i} = 0$ and there exists a witness $|\varphi\rangle$ causing $M_i$ to accept with probability greater than $1/2$, or else $\theta_{n_i} \in [1, 2]$ and no witness causes $M_i$ to accept with probability 1.) This is clearly possible by the argument above.

Notice that the proof of Theorem 2 breaks down if either $T$ (the number of queries to the unitary $U$) or $Q$ (the size of the witness) is infinite. This is not an accident. If $T$ is infinite, then a quantum algorithm can determine $\theta$ exactly, with no need for a witness. If $Q$ is infinite, then the witness $|\varphi\rangle$ can describe $\theta$ to infinite precision, and verifying the description (with perfect completeness) requires just a single query to $U$.

## 4 Discussion

Perhaps the strangest aspect of Theorem 2 is its lack of dependence on the polynomial running time of the $\mathsf{QMA}_1$ machine. For example, the same argument gives a quantum oracle $\mathcal{U}$ such that $\mathsf{BQP}^{\mathcal{U}} \not\subset \mathsf{QMA}_1\mathsf{EXP}^{\mathcal{U}}$, where $\mathsf{QMA}_1\mathsf{EXP}$ is the exponential-time version of $\mathsf{QMA}_1$, and even $\mathsf{BQP}^{\mathcal{U}} \not\subset \mathsf{QMA}_1\mathsf{EEEXP}^{\mathcal{U}}$. Indeed, just by using Proposition 2 about real analytic functions, without Lemma 1 or the theorem of Alekseevsky et al. [12], one can construct a quantum oracle $\mathcal{U}$ such that (for example) $\mathsf{BQP}^{\mathcal{U}} \not\subset \mathsf{ZQEXP}^{\mathcal{U}}$, where $\mathsf{ZQEXP}$ is Zero-Error Quantum Exponential-Time.[d]

What makes this strange is that we know, by trivial relativizing arguments, that $\mathsf{BQP} \subseteq \mathsf{EXP} \subseteq \mathsf{ZQEXP}$. Reflecting on the apparent contradiction, one might suspect that the quan-

[d] If we just want to separate $\mathsf{BQP}$ from $\mathsf{ZQP}$ (Zero-Error Quantum Polynomial-Time), this can be done with an ordinary classical oracle. Indeed we can easily construct an oracle $A$ such that $\mathsf{BPP}^A \not\subset \mathsf{ZQP}^A$, by considering a problem where the answer is YES if $A(y) = 1$ for most $y \in \{0,1\}^n$, or NO if $A(y) = 0$ for most $y \in \{0,1\}^n$. Such a problem is trivially in $\mathsf{BPP}^A$, but can be shown not to be in $\mathsf{ZQP}^A$ using the polynomial method of Beals et al. [13]. It would be nice if the same trick gave us a classical oracle $A$ such that $\mathsf{BPP}^A \not\subset \mathsf{QMA}_1^A$, but of course it does not, since the result of Zachos and Fürer [2] (which is relativizing) implies that $\mathsf{BPP}^A \subseteq \mathsf{MA}^A = \mathsf{MA}_1^A \subseteq \mathsf{QMA}_1^A$ for all $A$.

tum oracle separations are "cheating" somehow. But this is not the case; the correct resolution is simply that results like $\mathsf{BQP} \subseteq \mathsf{ZQEXP}$, while classically relativizing, must be quantumly non-relativizing! But how could that be?

If we carefully write out a proof that $\mathsf{BQP} \subseteq \mathsf{ZQEXP}$, we see what the problem is. Since $\mathsf{ZQEXP}$ is a zero-error class, the "obvious" proof will have to proceed not by direct simulation of the $\mathsf{BQP}$ machine, but by representing the amplitudes of the $\mathsf{BQP}$ machine in some explicit way. (In other words, by mimicking the proofs of containments such as $\mathsf{BQP} \subseteq \mathsf{EXP}$ or $\mathsf{BQP} \subseteq \mathsf{PSPACE}$ [14].) But the technique of explicitly representing amplitudes, simple though it seems, is already quantumly non-relativizing: it can break down if there is a quantum oracle $\mathcal{U}$, some property of which must be decided without error!

Some readers might conclude from this that quantum oracles are illegitimate; others, that the whole problem comes from the introduction of one-sided-error quantum complexity classes like $\mathsf{QMA}_1$. Our own view is that questions of "complexity-theoretic legitimacy" need to be decided on a case-by-case basis. In the present case, the real substance of our result is that any proof of $\mathsf{QMA}_1 = \mathsf{QMA}$ will need to involve explicit representation of amplitudes (or something similar), rather than just black-box composition of quantum circuits.

It remains a major challenge to find a quantumly non-relativizing technique that both (i) goes beyond the known classically non-relativizing techniques such as arithmetization, and (ii) fails to relativize even with two-sided-error complexity classes.

## 5 Extensions and Open Problems

By analogy to our quantum oracle separating $\mathsf{QMA}_1$ from $\mathsf{QMA}$, one might ask whether it is possible to construct a "randomized oracle" $R$ separating $\mathsf{MA}_1$ from $\mathsf{MA}$.[e] This would show that the proof of $\mathsf{MA}_1 = \mathsf{MA}$ due to Zachos and Fürer [2] must have been "randomly non-relativizing." Indeed such a randomized oracle separation is possible: simply have $R$ either output 0 whenever it is queried (the NO case), or else output 0 or 1 with equal probability (the YES case). It is obvious that these two cases can be distinguished by a $\mathsf{BPP}^R$ machine, using $O(1)$ queries to $R$. On the other hand, because of the perfect completeness requirement, the two cases *cannot* be distinguished by an $\mathsf{MA}_1^R$ machine: having a witness in support of the YES case clearly makes no difference.

However, this classical counterpart of our result really *does* feel like cheating! With the randomized oracle, perfect completeness is unachievable for trivial information-theoretic reasons, even assuming an infinitely long $\mathsf{MA}$ witness. With the quantum oracle, by contrast, perfect completeness *would* be achievable, if there were only some way to specify $\theta$ to infinite precision using the quantum witness $|\varphi\rangle$. This is of course what Theorem 2 rules out.

The above discussion immediately suggests another question. In constructing the quantum oracle $\mathcal{U}$, can we ensure that the angles $\theta_n$ are all rational numbers (or belong to some other dense countable set)? Indeed, the proof of Theorem 2 can easily be modified to achieve this. This is because of the following extension of Proposition 2:

[e] It is also interesting to see why a classical version of our argument does *not* yield an ordinary classical oracle $A$ such that $\mathsf{MA}_1^A \neq \mathsf{MA}^A$, thereby contradicting the result of Zachos and Fürer [2] (which is relativizing). The answer turns out to involve the fact that in the classical case, Merlin can take advantage of the individual oracle bits, rather than just the total amplitude for a '1' outcome. To put it another way: in the classical case, there is no such thing as an oracle that is both continuous and deterministic.

Given a real analytic function $f : R \to R$, if there exists an open set $(x, y) \subset R$ such that $f(z) = 1$ for all rational points $z \in (x, y)$, then $f(z) = 1$ identically.

However, there is an interesting difference between the real and rational cases. In the case where the $\theta_n$'s are real, it is possible to construct a *single* quantum oracle $\mathcal{U}$ such that $\mathsf{BQP}^{\mathcal{U}} \not\subset \mathsf{QMA}_1\mathsf{TIME}(f(n))^{\mathcal{U}}$ for all functions $f$. For example, choosing each $\theta_n$ to be 0 with probability 1/2, or uniformly distributed in $[1, 2]$ with probability 1/2, will yield such a $\mathcal{U}$ with probability 1, by an argument due to Bennett and Gill [15]. In the rational case, such a strong separation is also achievable, but only by choosing the numerator and denominator of each rational number $\theta_n$ to grow faster than any computable function of $n$. If we sidestep the issue of computability, say by giving the function $f(n)$ to the $\mathsf{QMA}_1\mathsf{TIME}(f(n))$ machine as advice, then it is not hard to show the following:

Given any quantum oracle $\mathcal{U} = \{U_n\}_{n \geq 1}$ with rational angles $\{\theta_n\}_{n \geq 1}$, there exists a function $f$ such that $\mathsf{BQP}^{\mathcal{U}}$ is simulable by a zero-error quantum algorithm that makes $f(n)$ queries to $\mathcal{U}$.

We end with two open problems. First, suppose the rotation angle $\theta_n$ cannot assume a continuum of values, but only a large finite set of values $S_n$. Is it then the case that either $T$ must scale like $|S_n|^{\Omega(1)}$ or $Q$ must scale like $\Omega(\log |S_n|)$? What is the optimal tradeoff between $T$ and $Q$? Are quantum witnesses (in this setting) ever more powerful than classical witnesses of comparable size?

Second, can we prove a *classical* oracle separation between $\mathsf{QMA}_1$ and $\mathsf{QMA}$?

## Acknowledgments

I thank Andy Drucker and the anonymous reviewers for helpful comments, and Greg Kuperberg and Dave Xiao for discussions of related problems several years ago.

## References

1. L. Babai. Trading group theory for randomness. In *Proc. ACM STOC*, pages 421–429, 1985.
2. S. Zachos and M. Fürer. Probabilistic quantifiers vs. distrustful adversaries. In *Proc. Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 443–455. Springer-Verlag, 1987.
3. C. Lautemann. BPP and the polynomial hierarchy. *Inform. Proc. Lett.*, 17:215–217, 1983.
4. O. Goldreich and D. Zuckerman. Another proof that $BPP \subseteq PH$ (and more). ECCC TR97-045, 1997.
5. A. Kitaev, A. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 2002.
6. J. Watrous. Succinct quantum proofs for properties of finite groups. In *Proc. IEEE FOCS*, pages 537–546, 2000. cs.CC/0009002.
7. D. Aharonov and T. Naveh. Quantum NP - a survey. quant-ph/0210077, 2002.
8. C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005.
9. S. Bravyi. Efficient algorithm for a quantum analogue of 2-SAT. quant-ph/0602108, 2006.
10. A. Kitaev and J. Watrous. Parallelization, amplification, and exponential-time simulation of quantum interactive proof systems. In *Proc. ACM STOC*, pages 608–617, 2000.
11. S. Aaronson and G. Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. Previous version in Proceedings of CCC 2007. quant-ph/0604056.

12. D. Alekseevsky, A. Kriegl, M. Losik, and P. W. Michor. Choosing roots of polynomials smoothly. *Israel Journal of Mathematics*, 105:203–233, 1998.
13. R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. Earlier version in IEEE FOCS 1998, pp. 352-361. quant-ph/9802049.
14. E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997. First appeared in ACM STOC 1993.
15. C. H. Bennett and J. Gill. Relative to a random oracle A, $P^A \neq NP^A \neq coNP^A$ with probability 1. *SIAM J. Comput.*, 10(1):96–113, 1981.