

# Limits to black-box amplification in QMA

Scott Aaronson<sup>1</sup>, Phillip Harris<sup>2</sup>, and Frek Witteveen<sup>3</sup>

<sup>1</sup>University of Texas at Austin. Supported by the Simons Foundation

<sup>2</sup>University of Bonn

<sup>3</sup>QuSoft & CWI, Amsterdam

## Abstract

We study the limitations of black-box amplification in the quantum complexity class QMA. Amplification is known to boost any inverse-polynomial gap between completeness and soundness to exponentially small error, and a recent result (Jeffery and Witteveen, 2025) shows that completeness can in fact be amplified to be doubly exponentially close to 1. We prove that this is optimal for black-box procedures: we provide a quantum oracle relative to which no QMA verification procedure using polynomial resources can achieve completeness closer to 1 than doubly exponential, or a soundness which is super-exponentially small. This is proven by making the oracle separation from (Aaronson, 2009) between QMA and QMA<sub>1</sub> quantitative, using techniques from complex approximation theory.

## 1 Introduction

The complexity class QMA (Quantum Merlin–Arthur) is a quantum analogue of NP, where given a problem instance, a prover Merlin sends a quantum witness to a polynomial-time quantum verifier Arthur, in order to convince Arthur whether the problem is a yes-instance or a no-instance. Arthur then performs a polynomial-size quantum computation on the witness, after which he accepts or rejects the witness. Quantum complexity classes typically allow some probability of error. There are two parameters quantifying the allowed error: the *completeness*  $c$ , the probability with which Arthur accepts a valid witness in the yes-case, and the *soundness*  $s$ , the maximum acceptance probability in the no-case. It is well known that the precise choice of these parameters does not matter as long as there is a non-negligible gap: by amplification techniques [KSV02, MW05], one can boost a polynomially small gap  $c - s = 1/\text{poly}$  to completeness  $c = 1 - 2^{-\text{poly}}$  and soundness  $s = 2^{-\text{poly}}$  exponentially close to 1 and 0 respectively. This motivates the canonical definition, which takes  $c = \frac{2}{3}$  and  $s = \frac{1}{3}$ .

A long-standing open question is whether the completeness can in fact be made *perfect*. We denote by QMA<sub>1</sub> the variant of QMA where we take  $c = 1$  (and  $s = \frac{1}{3}$ , or any other constant). The question whether QMA equals QMA<sub>1</sub> is then: can every QMA protocol be modified so that Arthur always accepts a valid witness in the yes-case, while still rejecting no-instances with bounded probability? There are a number of closely related complexity classes which allow for perfect completeness. This is the case for the variant with only classical randomness MA, the variant with a classical proof and quantum verifier QCMA [JKNN12], the variant with multiple rounds QIP [MW05], and the variant with an exponentially small soundness-completeness gap PreciseQMA [FL18].

For QMA, this question has resisted resolution, and one explanation for this is given by Aaronson [Aar09], who gave a barrier to proving QMA = QMA<sub>1</sub> using black-box techniques. Here ‘black-box’ refers to an amplification procedure that does not use any properties of the verification circuit, other than those that define it. Concretely, [Aar09] proved that there exists a quantum oracle relative to which QMA  $\neq$  QMA<sub>1</sub>. Any proof showing QMA = QMA<sub>1</sub> would have to break down in the presence of a quantum oracle, i.e. it would have to be quantumly nonrelativizing.




Figure 1: The verifier would like to decide whether  $|\theta| \leq 3\pi/4$  (a yes-instance) or  $\theta = \pi$  (a no-instance), given black-box access to the oracle defined in Eq. (1). Here we give a sketch of what the maximal acceptance probability, optimized over all choices of witness, should look like, when we have completeness  $c = 1 - \delta$  close to 1, and soundness  $s = 1/3$ .

All known amplification techniques for QMA are of a black-box nature (so they are quantumly relativizing).

Recently, it was shown in Ref [JW25] that one can reach completeness *doubly exponentially* close to 1

$$c = 1 - 2^{-2^{\text{poly}}},$$

in QMA, improving over the previous best completeness which was exponentially close to 1. The techniques in [JW25] are of a black-box nature. This raises the question: can one get even closer to perfect completeness with black-box amplification procedures?

In this work we show that the doubly-exponential bound is in fact *optimal* in the black-box setting. We prove that there exists a quantum oracle relative to which no QMA protocol using polynomial resources can achieve completeness closer to 1 than doubly exponential.

The key idea in the oracle separation of [Aar09] is that for a simple choice of quantum oracle, the maximal acceptance probability of the verifier is an analytic function in one variable. Perfect completeness would require this function to be constant on an interval, but the analyticity then implies that the maximal acceptance probability is *always* 1, contradicting soundness. We use the same quantum oracle, but give a slightly different analysis and use a standard result bounding the growth of trigonometric polynomials, to make this separation quantitative. The key idea of the proof is as follows. The quantum oracle is given by

$$U(\theta) = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}, \quad (1)$$

for  $\theta \in [-\pi, \pi)$ . We then consider the problem where Arthur has to decide whether  $|\theta| \leq \pi - s$  (yes-instance) or  $\theta = \pi$  (no-instance), with black-box access to  $U(\theta)$ .

Suppose we are given some verifier for this task. The idea is to consider the accepting measurement operator  $P_{\text{acc}}(\theta)$  and the rejecting measurement operator  $P_{\text{rej}}(\theta) = I - P_{\text{acc}}(\theta)$ . The maximal acceptance probability  $p(\theta)$ , over all possible witnesses, equals the largest eigenvalue of  $P_{\text{acc}}(\theta)$ . If the verifier has completeness close to 1, and soundness  $\frac{1}{3}$ , the maximal acceptance probability, as a function of the angle  $\theta$ , must look something like shown in Fig. 1. We now study the following function:

$$p(\theta) = \det[P_{\text{rej}}(\theta)]. \quad (2)$$

If the verifier has completeness close to 1,  $P_{\text{rej}}(\theta)$  has an eigenvalue close to 0, and hence  $p(\theta)$  must be small on the accepting values of  $\theta$ . On the other hand, we can use soundness to bound how small  $p(\pi)$  can be. We can then use the fact that  $p(\theta)$  is a trigonometric polynomial, and a standard result on the growth of polynomials, to bound how close to 1 the completeness can be. As a side comment, we note that this also simplifies the argument of [Aar09]: if one has perfect completeness,  $p(\theta) = 0$  on an interval, but then  $p(\theta)$ , being a trigonometric polynomial and hence analytic, has to be identically zero for all  $\theta$ , implying that the soundness condition can not be satisfied.

We also investigate the analogous question for the soundness parameter in the definition of QMA. It is known that QMA with perfect soundness equals the complexity class NQP [KMY03], which is hard for the polynomial hierarchy [FGHP99], and thus likely different from QMA. The soundness can be amplified to be exponentially small with standard techniques. Can one do even better? Can one get, analogous to the situation for completeness, to doubly-exponentially small soundness? We show, using the same oracle but with the role of the yes- and no-instances reversed, and using

$$p(\theta) = \text{tr}[P_{\text{acc}}(\theta)] \tag{3}$$

that there is no black-box way to do so. This completely settles what completeness and soundness parameters are achievable with black-box amplification techniques in QMA.

A caveat is that the proof breaks down if the witness register is allowed to be infinite dimensional, as also observed in [Aar09]. This is not just a quirk of the proof technique: in [JW25] it is proven that (with black-box amplification techniques) one in fact can prove perfect completeness for QMA, provided that one allows computation on an infinite-dimensional Hilbert space with an appropriate gate set. This choice of gate set is such that the infinite-dimensional Hilbert space does not increase the computational power of BQP or QMA.

**Use of AI:** In an earlier version of this paper, we used a different function in the completeness case, namely the *rational* trigonometric function  $r(\theta) = \text{tr}[P_{\text{rej}}(\theta)^{-1}]$ , which was suggested to us by GPT-5-Thinking. To analyze this  $r$ , we used a bound on the growth of complex rational functions [Gon69], which in turn required the no-instances to be an interval instead of a single point. See [Aar25] for a discussion and comments. After the initial posting, one of us (PH, subsequently added as coauthor) suggested to replace  $\text{tr}[P_{\text{rej}}(\theta)^{-1}]$  by the simpler  $\det[P_{\text{rej}}(\theta)]$ .

## 2 Preliminaries

We first recall the definition of QMA.

**Definition 1.** For parameters  $c, s \in [0, 1]$  with  $c > s$ , the class  $\text{QMA}_{c,s}$  consists of promise problems  $L = (L_{\text{yes}}, L_{\text{no}})$  for which there exists a uniform family of polynomial-time quantum circuits  $\{V_x\}_{x \in \{0,1\}^n}$ , called verifiers, with the following properties:

- (Completeness) If  $x \in L_{\text{yes}}$ , then there exists a quantum witness state  $|\psi\rangle$  such that

$$\Pr[V_x \text{ accepts } |\psi\rangle] \geq c.$$

- (Soundness) If  $x \in L_{\text{no}}$ , then for all quantum states  $|\psi\rangle$ ,

$$\Pr[V_x \text{ accepts } |\psi\rangle] \leq s.$$

The class QMA is defined as  $\text{QMA}_{2/3,1/3}$ , since any inverse-polynomial gap between  $c$  and  $s$  can be amplified to these parameters by standard techniques [MW05]. The subclass  $\text{QMA}_1$  is defined as  $\text{QMA}_{1,s}$  for some constant  $s < 1$ , i.e., *QMA with perfect completeness*.

We will also consider QMA with access to a quantum oracle  $\mathcal{U}$ . Here, there is a collection of possible unitaries, and the verifier gets access to a unitary  $U$  from this set. The verifier has black-box access to  $U$ , its inverse, and to their controlled versions. We then denote by  $\text{QMA}^{\mathcal{U}}$  the class of problems that can be solved as in [Theorem 1](#) where Arthur additionally has access to a polynomial number of calls to  $\mathcal{U}$ .

## 2.1 Trigonometric polynomials

A *trigonometric polynomial* of degree  $d$  is a polynomial of degree  $d$  in  $\cos(\theta)$  and  $\sin(\theta)$ .

We use a standard bound on how fast such polynomials can grow: if  $p(\theta)$  is uniformly bounded on some interval, how large can it be outside the interval? The solution is similar to the case of standard polynomials, where the extremal polynomials are Chebyshev polynomials.

We use the following bound; see [Theorem 5.1.2](#) in [\[BE12\]](#):

**Theorem 2.** *Let  $p(\theta)$  be a trigonometric polynomial of degree  $d$ . Then for  $u \in (0, \pi/4]$ ,*

$$\max_{\theta} |p(\theta)| \leq \exp(8du) \max_{|\theta| \leq \pi - u} |p(\theta)| \quad (4)$$

That is, if  $p(\theta)$  is uniformly bounded on a subset of angles, we can bound the maximal value it can take.

## 3 Limits to black-box QMA amplification

In this section we prove tight black-box lower bounds for amplifying the completeness and soundness parameters of QMA. The following result shows that there is an oracle with respect to which QMA has completeness  $c(n)$  which is at most doubly-exponentially close to 1, and soundness which is at most exponentially small. Both bounds follow from the same oracle construction (as in [\[Aar09\]](#)) together with the polynomial growth bound in [Theorem 2](#).

**Theorem 3** (Limits to black-box amplification). *There exists a quantum oracle  $\mathcal{U}$  such that the following hold for any QMA verification procedure using  $\text{poly}(n)$  resources:*

1. (**Completeness barrier**). *For any black-box QMA amplification procedure that achieves completeness  $c = 1 - \delta$  and soundness  $s = 1/3$ , we have*

$$\delta \geq 2^{-2^{\text{poly}(n)}}.$$

*More precisely, relative to  $\mathcal{U}$ ,*

$$\text{QMA}^{\mathcal{U}} \neq \text{QMA}_{c(n), 1/3}^{\mathcal{U}} \quad \text{for } c(n) = 1 - o\left(2^{-2^{\text{poly}(n)}}\right).$$

2. (**Soundness barrier**). *For any black-box QMA amplification procedure that achieves completeness  $c = 2/3$  and soundness  $s = \delta$ , we have*

$$\delta \geq 2^{-\text{poly}(n)}.$$

*In particular, relative to  $\mathcal{U}$ ,*

$$\text{QMA}^{\mathcal{U}} \neq \text{QMA}_{2/3, s(n)}^{\mathcal{U}} \quad \text{for } s(n) = o\left(2^{-\text{poly}(n)}\right).$$

*Proof.* We use the same quantum oracle as in [\[Aar09\]](#). For  $\theta \in (-\pi, \pi]$ , define

$$U(\theta) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}.$$

The promise problem is to decide whether the oracle parameter satisfies  $|\theta| \leq \pi - u$  for some arbitrary constant  $u \in (0, \pi/4]$ , or  $\theta = \pi$ . Which of these two cases corresponds to yes-instances and no-instances is reversed for the proof of the completeness and soundness barrier, but in either case it is obvious that Arthur can do so with constant soundness and completeness (and even with trivial witness), so the problem is in QMA. We will now show, however, that there are limits to how close the completeness and soundness can get to 1 and 0 respectively using polynomial resources.

Consider any QMA verifier  $V$  that, on input  $x \in \{0, 1\}^n$ , uses

- $t = t(n) = \text{poly}(n)$  calls to the oracle,
- a witness register of  $m = m(n) = \text{poly}(n)$  qubits; let  $q = 2^m$  denote its dimension.

Let  $P_{\text{acc}}(\theta)$  denote the POVM element corresponding to acceptance and let  $P_{\text{rej}}(\theta) = I - P_{\text{acc}}(\theta)$ . Write the eigenvalues of  $P_{\text{acc}}(\theta)$  as  $1 \geq \lambda_1(\theta) \geq \dots \geq \lambda_q(\theta) \geq 0$ . For each fixed  $\theta$ , the optimal acceptance probability, over all choices of witness, equals  $\lambda_1(\theta)$ .

Every matrix element of  $U(\theta)$  and  $U(\theta)^\dagger$  (and hence also of their controlled versions) is affine in  $\cos(\theta)$  and  $\sin(\theta)$ . Hence, by expanding the verifier circuit and projecting on the accept outcome, each entry of  $P_{\text{acc}}(\theta)$  and  $P_{\text{rej}}(\theta)$  is a trigonometric polynomial of degree  $2t$  in  $\theta$ .

**Completeness barrier** We let  $|\theta| \leq \pi - u$  be the yes-instances, and  $\theta = \pi$  be the no-instance. Assume the verifier achieves completeness  $1 - \delta$  and soundness  $1/3$ :

$$\lambda_1(\theta) \geq 1 - \delta \text{ for yes-instances,} \quad \lambda_1(\theta) \leq \frac{1}{3} \text{ for a no-instance.}$$

We consider the following function:

$$p(\theta) = \det[P_{\text{rej}}(\theta)] = \prod_{i=1}^q (1 - \lambda_i(\theta)),$$

which is a trigonometric polynomial of degree  $2tq$ .

We now bound  $p$  for yes- and no-instances. For yes-instances,  $1 - \lambda_1(\theta) \leq \delta$ , so  $p(\theta) \leq \delta$ . For the no-instance, all  $\lambda_i(\theta) \leq 1/3$  and hence  $p(\pi) \geq (2/3)^q$ .

By [Theorem 2](#),

$$(2/3)^q \leq \max_{\theta} |p(\theta)| \leq \exp(16utq) \max_{|\theta| \leq \pi - u} |p(\theta)| \leq \exp(16utq)\delta.$$

In other words,

$$\delta \geq (2/3)^q \exp(-16utq) \tag{5}$$

and since  $q = 2^{\text{poly}(n)}$ , and  $t = \text{poly}(n)$ , this proves [1](#).

**Soundness barrier** We let  $|\theta| \leq \pi - u$  be the no-instances, and  $\theta = \pi$  be the yes-instance. Assume the verifier achieves completeness  $2/3$  and soundness  $\delta$ , so

$$\begin{aligned} \lambda_1(\theta) &\leq \delta \text{ for no-instances} \\ \lambda_1(\theta) &\geq \frac{2}{3} \text{ for a yes-instance.} \end{aligned}$$

We now consider

$$p(\theta) = \text{tr}[P_{\text{acc}}(\theta)] = \sum_{i=1}^q \lambda_i(\theta).$$

This is a trigonometric polynomial of degree  $2t$  in  $\theta$ . We can bound it as

$$\begin{aligned} p(\theta) &\leq q\lambda_1(\theta) \leq q\delta \text{ for no-instances,} \\ p(\theta) &\geq \lambda_1(\theta) \geq \frac{2}{3} \text{ for a yes-instance.} \end{aligned}$$

By [Theorem 2](#),

$$\frac{2}{3} \leq \max_{\theta} |p(\theta)| \leq \exp(16ut) \max_{|\theta| \leq \pi - u} |p(\theta)| \leq \exp(16ut)q\delta.$$

In other words,

$$\delta \geq \frac{2}{3q} \exp(-16ut) \tag{6}$$

and since  $q = 2^{\text{poly}(n)}$ , and  $t = \text{poly}(n)$ , this proves [2](#). □

Note that in the proof of [Theorem 3](#), the crucial difference between the completeness and the soundness cases is that the degree of the function  $p(\theta)$  does not depend on the witness dimension  $q$  for the soundness.

Since black-box amplification procedures can also *achieve* completeness  $c = 1 - 2^{-2^{\text{poly}(n)}}$  and soundness  $s = 2^{-\text{poly}(n)}$ , this gives the optimal completeness and soundness parameters for QMA which can be achieved through black-box reductions.

## 4 Discussion

While previously it was known how to achieve completeness and soundness exponentially close to 1 and 0 respectively, this work, together with [\[JW25\]](#), shows that with black-box reductions one can achieve completeness doubly exponentially close to 1, and soundness exponentially small, and that this is optimal. The asymmetry between soundness and completeness is natural from the definition: completeness only requires a single eigenvalue of the accepting measurement operator to be very close to 1, while soundness is a condition on all possible witness states and requires *all* eigenvalues of the accepting measurement operator to be close to 0.

Let us call attention to something conceptually strange about our result. Recall that  $\text{QMA} \subseteq \text{PP} = \text{PostBQP}$ , with the containment believed to be strict. Nevertheless, we saw in this paper that “the approximation theory object corresponding to QMA,” namely the largest eigenvalue of an  $\exp(n) \times \exp(n)$  Hermitian matrix of low-degree polynomials, is in some respects *more* powerful than “the approximation theory object corresponding to PP,” namely a low-degree rational function. In particular, the former allows amplification to doubly exponentially small completeness error, while the latter does not.

How is this possible? In our view, the resolution is simply that, to prove  $\text{QMA} \subseteq \text{PP}$ , it suffices to calculate only a crude estimate of the largest eigenvalue in question—something that *is* possible using low-degree polynomials. Indeed, if we ask for more precise information about the largest eigenvalue, we get the class `PreciseQMA`, which equals `PSPACE` [\[FL18\]](#) and hence is presumably more powerful than  $\text{PostBQP} = \text{PP}$ .

Of course, the central open question we leave is whether  $\text{QMA} = \text{QMA}_1$  for some, or all, finite-dimensional gate sets.

## Acknowledgments

We thank Stacey Jeffery for helpful discussions. FW was supported by the European Union (ERC, ASC-Q, 101040624). PH was supported by Deutsche Forschungsgemeinschaft (DFG) Grant No. SFB-TRR 358/1 2023 — 491392403.

## References

- [Aar09] Scott Aaronson. On perfect completeness for QMA. *Quantum Information and Computation*, 9:81–89, 2009. arXiv: [0806.0450](https://arxiv.org/abs/0806.0450).
- [Aar25] Scott Aaronson. Shtetl Optimized: The QMA Singularity. <https://scottaaronson.blog/?p=9183>, 2025.
- [BE12] Peter Borwein and Tamás Erdélyi. *Polynomials and polynomial inequalities*, volume 161. Springer Science & Business Media, 2012.
- [FGHP99] Stephen Fenner, Frederic Green, Steven Homer, and Randall Pruim. Determining acceptance possibility for a quantum computation is hard for the polynomial hierarchy. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 455(1991):3953–3966, 1999. arXiv: [quant-ph/9812056](https://arxiv.org/abs/quant-ph/9812056).
- [FL18] Bill Fefferman and Cedric Yen-Yu Lin. A complete characterization of unitary quantum space. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, pages 4–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018. arXiv: [1604.01384](https://arxiv.org/abs/1604.01384).
- [Gon69] AA Gončar. Zolotarev problems connected with rational functions. *Mathematics of the USSR-Sbornik*, 7(4):623, 1969.
- [JKNN12] Stephen P. Jordan, Hirotada Kobayashi, Daniel Nagaj, and Harumichi Nishimura. Achieving perfect completeness in classical-witness quantum Merlin-Arthur proof systems. *Quantum Information and Computation*, 12(5-6):461–471, 2012. arXiv: [1111.5306](https://arxiv.org/abs/1111.5306).
- [JW25] Stacey Jeffery and Frek Witteveen.  $\text{QMA} = \text{QMA}_1$  with an infinite counter. arXiv: [2506.15551](https://arxiv.org/abs/2506.15551), 2025.
- [KMY03] Hirotada Kobayashi, Keiji Matsumoto, and Tomoyuki Yamakami. Quantum Merlin-Arthur proof systems: Are multiple Merlins more helpful to Arthur? In *International Symposium on Algorithms and Computation*, pages 189–198. Springer, 2003. arXiv: [quant-ph/0306051](https://arxiv.org/abs/quant-ph/0306051).
- [KSV02] Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi. *Classical and quantum computation*. Number 47 in Graduate Studies in Mathematics. American Mathematical Society, 2002.
- [MW05] Chris Marriott and John Watrous. Quantum Arthur–Merlin games. *computational complexity*, 14(2):122–152, 2005. arXiv: [cs/0506068](https://arxiv.org/abs/cs/0506068).