QMA Lower Bounds for Batch Verification via Approximate Degree

Mark Bun
Boston University, MA, USA

Mandar Juvekar
Boston University, MA, USA

Samuel King
Georgetown University, Washington, DC, USA

Abstract

We study batch verification in QMA query and communication complexity, where the goal is to understand how the resources needed to verify m copies of a Boolean function f depend on m. We give a general technique for proving lower bounds on the witness-query tradeoff needed to batch verify a function f in terms of its approximate degree. Applying this technique to an explicit family of DNF formulas f, we show that attempting to save even a constant factor on the witness length of the baseline approach to batch verifying f necessitates a large polynomial increase in the query cost. We also obtain new lower bounds on the QMA query complexity of read-once CNF formulas and on the surjectivity and k-element distinctness functions. Our lower bounds also lift to give communication analogs of these results.

2012 ACM Subject Classification Theory of computation → Quantum query complexity

Keywords and phrases QMA, approximate degree, query complexity, communication complexity, quantum computation

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.69

Category RANDOM

Related Version *Full Version*: `https://arxiv.org/abs/2607.08888` [8]

Funding *Mark Bun*: supported in part by NSF award CNS-2046425 and a Sloan Research Fellowship. *Mandar Juvekar*: supported in part by NSF award CNS-2046425 and an IBM PhD Fellowship. *Samuel King*: supported in part by NSF CAREER award CCF-2338730.

Acknowledgements We would like to thank Nadezhda Voronova for her insightful comments on a draft of this paper, Sasha Golovnev for invaluable discussions during this work, and Robin Kothari and Justin Thaler for stimulating our interest in these problems. We would also like to thank Robin Kothari for spotting an error in a previous manuscript.

1 Introduction

A powerful but untrusted server wants to convince a computationally constrained client that it has correctly performed a large list of computations that were delegated to it. It can do so by proving to the client the veracity of m statements $x_1, \ldots, x_m$. As communication is costly, the proof must be as short as possible. If the statements are instances of an NP (or more generally for this work, QMA) language, a baseline strategy is to send as a proof witnesses $w_1, \ldots, w_m$ for all of the statements – the client can then verify each statement individually. Both the communication and the client's computational cost in this strategy are m times the cost for a single statement. For what statements is it possible to do better? What tradeoffs are achievable between communication and computation?

This *batch verification* problem has generated exciting research in complexity theory and cryptography. Highly communication-efficient interactive proofs are known for batch verifying UP statements (i.e., NP statements with unique witnesses) [34, 35, 36], while a connection between batch verification and statistical witness indistinguishability [6] suggests extending such a result to all of NP may be impossible. Zero-knowledge batch verification protocols are known for problems admitting non-interactive statistical zero-knowledge proofs [25, 26, 31, 27]. Under cryptographic assumptions, computationally sound noninteractive batch arguments are known for all of NP [23, 16, 15, 18, 42, 22]; recent work has even constructed succinct interactive (classical) batch arguments for QMA [20]. In cryptography, the study of batch verification of NP statements has led to new constructions of succinct non-interactive arguments (SNARGs) [15, 24] and non-interactive zero-knowledge arguments (NIZKs) [14] for more expressive classes and from weaker assumptions. Finally, [32] recently initiated a study of quantum witness indistinguishable proofs and showed that they are implied by quantum batch proofs.

In this work, we study the batch verification problem for quantum Merlin-Arthur (QMA) proofs in the query and communication models. These models abstract computation to settings where techniques are available for proving strong lower bounds, yet which shed both conceptual and technical light on the corresponding Turing machine model. In the QMA *query* or *black-box* model, hereafter denoted by $\mathrm{QMA}^{\mathrm{dt}}$, a quantum verifier (the resource-bounded mortal "Arthur") wishes to evaluate a known function $f : \{-1,1\}^n \to \{-1,1\}$ on an input $x \in \{-1,1\}^n$.[1] He receives a witness consisting of w qubits from a prover (the all-powerful wizard "Merlin"), and is allowed to make q queries to x in quantum superposition. The protocol is correct if (1) whenever $f(x) = -1$ (i.e., x is a "true" input), there exists a witness that causes the verifier to accept with probability at least $2/3$, and (2) whenever $f(x) = 1$ (i.e., x is a "false" input), every witness causes the verifier to reject with probability at least $2/3$. The cost of such a protocol is the sum $w + q$ of the witness length and the number of verifier queries, though we are also interested in understanding the tradeoff between these parameters. The QMA query complexity of a Boolean function f, denoted $\mathsf{QMA}^{\mathsf{dt}}(f)$, is the least cost of a protocol computing f.

The batch verification problem for QMA query complexity can thus be stated as follows: under what conditions on f does the function $(\mathsf{AND}_m \circ f^m)(x_1, \ldots, x_m) := f(x_1) \wedge \cdots \wedge f(x_m)$ have low $\mathrm{QMA}^{\mathrm{dt}}$ complexity? But what should "low" mean here? In the quantum setting, Grover's search algorithm gives a way to "batch compute" all functions with bounded-error algorithms: when f has a bounded-error quantum query protocol (i.e., one that does not make use of a witness) with query cost q, Grover search implies that $\mathsf{AND}_m \circ f^m$ has a bounded-error quantum query protocol with cost only $O(\sqrt{m} \cdot q)$ [21]. This methodology extends to $\mathrm{QMA}^{\mathrm{dt}}$ protocols. If f admits a $\mathrm{QMA}^{\mathrm{dt}}$ protocol with witness length w_f and query complexity q, one can compute $\mathsf{AND}_m \circ f^m$ by taking as witness all m individual witnesses (for a total witness length of $w = m \cdot w_f$) and using Grover search to verify all m copies using $O(\sqrt{m} \cdot q)$ queries. Our batch verification problem thus becomes: for what functions f can one beat this baseline protocol for $\mathsf{AND}_m \circ f^m$ that uses Grover search to verify all m individual witnesses?

Our first main result establishes a strong limitation on batching even when f is highly structured and admits an efficient NP query protocol.

[1] The notation $\mathrm{QMA}^{\mathrm{dt}}$ alludes to yet another view of query algorithms as generalizations of decision trees. Here and throughout the paper we identify -1 with "true" and 1 with "false."

▶ **Theorem 1** (Informal). *For every $\delta > 0$ there is an explicitly given polynomial-size, constant-width DNF $f : \{-1,1\}^n \to \{-1,1\}$ such that every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$, $m = n^{O(1)}$, with witness length $w \le cm \log n$ for a suitable constant $c \in (0,1)$ must make at least $\widetilde{\Omega}(n^{1-\delta}\sqrt{m/w})$ queries.*

One can interpret this result as follows. Since the function f is computed by a polynomial-size, constant-width DNF, it admits an NP query protocol with witness length $O(\log n)$ and $O(1)$ queries. The baseline strategy for verifying $\mathsf{AND}_m \circ f^m$ thus entails a proof with witness length $O(m \log n)$ and uses $O(\sqrt{m})$ queries. Theorem 1 shows that if one seeks to beat this baseline using a protocol with witness length just a constant factor smaller than this $O(m \log n)$, then the requisite query complexity jumps to $\Omega(n^{1-\delta})$, which is $\Omega(m)$ whenever $m = O(n^{1-\delta})$.

We prove Theorem 1 by proving a more general connection between the $\mathrm{QMA}^{\mathrm{dt}}$ complexity of batch verifying f with the *approximate degree* of f. The ε-approximate degree of f, denoted $\mathrm{adeg}_\varepsilon(f)$, is the minimum degree of a real-valued polynomial that approximates f point-wise to additive error ε. Approximate degree is an extensively studied measure, and has many connections to randomized and quantum query and communication complexity (see, e.g., [13]). In particular, Beals et al. [5] showed that the acceptance probability of every q-query quantum algorithm can be written as a degree-$2q$ polynomial in its input, and hence the ε-error quantum query complexity of a function f is always at least half of its ε-approximate degree.

We show that lower bounds on the ε-approximate degree of f for various settings of ε yield witness-query tradeoffs for $\mathrm{QMA}^{\mathrm{dt}}$ batch verifying f.

▶ **Theorem 2** (Informal). *Let $f : \{-1,1\}^n \to \{-1,1\}$ be a function.*

1. *If $\mathrm{adeg}_{1-1/m}(f) > d$ then every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $w = O(m)$ requires $\Omega(d\sqrt{m/w})$ queries.*
2. *If $\mathrm{adeg}_{1/3}(f) > d$ then every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length at most m requires $\Omega(d)$ queries.*

As we discuss below, previous work derived $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds from more stringent variants of approximate degree lower bounds. However, to our knowledge, ours is the first to draw a direct connection between $\mathrm{QMA}^{\mathrm{dt}}$ complexity and the standard notion of approximate degree.

As an illustrative example, let f be the complement of the element distinctness function on n bits (defined in Section 2.1), which is known to have approximate degree $\tilde{\Omega}(n^{2/3})$ [3, 4]. At the same time, f admits a $\mathrm{QMA}^{\mathrm{dt}}$ protocol with witness length $O(\log n)$ and query complexity $O(\sqrt{\log n})$: the witness consists of the location of a collision, and the verifier uses Grover search to check the collision. So the naïve batching strategy achieves witness length $O(m \log n)$ and query complexity $O(\sqrt{m \log n})$. Item 2 of Theorem 2 shows that any protocol improving on this witness length by even an $O(\log n)$ factor must have its query complexity jump to $\widetilde{\Omega}(n^{2/3})$, which is significantly larger for any $m \ll n^{4/3}$. Item 1 of the theorem gives more fine-grained tradeoffs when one can obtain lower bounds on the approximate degree of f to *large* error. Similar tradeoffs apply to other functions that admit efficient $\mathrm{QMA}^{\mathrm{dt}}$ protocols such as DNFs and the complement of the surjectivity function.

Beyond the intrinsic interest of the batch verification problem, Theorem 2 also serves as a valuable technical tool for understanding the $\mathrm{QMA}^{\mathrm{dt}}$ complexity of specific functions that have been central to the study of quantum query complexity. Sherstov and Thaler [41] used $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds on functions of the form $\mathsf{AND}_m \circ f^m$ to give the sharpest known lower bounds for the influential element distinctness and permutation testing problems. For

technical reasons elaborated on below, their method falls short of applying to closely related functions, such as the two-level AND-OR tree, the k-element distinctness function, and the surjectivity function. Our framework for proving batch verification lower bounds bypasses the limitations from prior work, allowing us to prove the following lower bounds for these functions.

▶ **Theorem 3.** *Let* SURJ_n *denote the surjectivity function on* n *bits and* $k\mathsf{ED}_n$ *the* k*-element distinctness function on* n *bits.*

1. $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_{n^{1/3}} \circ \mathsf{OR}_{n^{2/3}}^{n^{1/3}}) = \Omega(n^{1/3})$.
2. $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{SURJ}_n) = \widetilde{\Omega}(n^{3/7})$.
3. $\mathsf{QMA}^{\mathsf{dt}}(k\mathsf{ED}_n) = \widetilde{\Omega}(n^{\frac{3k-1}{7k-1}})$.

The best lower bounds that followed straightforwardly from prior work were $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_{n^{1/3}} \circ \mathsf{OR}_{n^{2/3}}^{n^{1/3}}) = \Omega(n^{1/6})$, $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{SURJ}_n) = \widetilde{\Omega}(n^{1/4})$ and $\mathsf{QMA}^{\mathsf{dt}}(k\mathsf{ED}_n) \geq \mathsf{QMA}^{\mathsf{dt}}(\mathsf{ED}_n) = \widetilde{\Omega}(n^{2/5})$.

1.1 Technical Overview

Our main results follow from a general result (Corollaries 12 and 13, which formalize Theorem 2) showing that $\mathsf{AND}_m \circ f^m$ has high $\mathrm{QMA}^{\mathrm{dt}}$ complexity whenever f has high *approximate degree*.

Suppose there is a QMA protocol for a function F using witness length w and q queries. The in-place amplification technique of Marriott and Watrous [30] gives rise to a protocol for F with improved error 2^{-8w} using witness length w and $O(qw)$ queries. Replacing the witness in this protocol with the maximally mixed state yields a witness-free quantum query protocol for F, and hence an approximating polynomial p with the following properties:

$$F(x) = -1 \implies -2^{8w} \leq p(x) \leq -1,$$
$$F(x) = 1 \implies 1 - 2^{-8w} \leq p(x) \leq 1.$$

That is, p very tightly approximates F on its false inputs, but (roughly) subject to having the right sign behavior, can become exponentially large on true inputs. Thus, in order to prove a lower bound on the QMA query complexity of F, it suffices to rule out low-degree approximating polynomials p of this form. Up to this point, this lower bound method, which is inspired by the complexity class containment $\mathsf{QMA} \subseteq \mathsf{SBQP}$, is well-established for proving QMA query complexity lower bounds [1, 2, 41, 17]. We refer to such an approximating polynomial p as an "SBQP approximation" to F.

Prior work proved degree lower bounds for SBQP approximations by relating it to other better-studied notions of polynomial approximation, including ε-approximate degree in the "large error" regime where $\varepsilon \approx 1 - 2^{-w}$, threshold degree, and vanishing error one-sided approximate degree. We depart from prior work by studying SBQP approximations more directly. We do this using the method of dual polynomials, whereby one proves approximability lower bounds by constructing dual solutions to a linear program capturing the approximation problem. We introduce a notion of "relaxed SBQP duals" to capture feasible solutions in this dual space that are relatively clean to construct and reason about. A relaxed SBQP dual for a function F is a real-valued function Φ that is ε-correlated with F for some $\varepsilon \geq 2^{-8w}$, orthogonal to all low degree polynomials, and which has "almost one-sided error" in the sense that the total weight it places on $\{x : F(x) = -1 \wedge \Phi(x) > 0\}$ is at most $2^{-8w}\varepsilon$. The proof of our general theorem proceeds by showing that whenever f has high approximate degree, we can construct an appropriate relaxed SBQP dual for the function $\mathsf{AND}_m \circ f^m$. The connection outlined above, in turn, implies our $\mathrm{QMA}^{\mathrm{dt}}$ tradeoffs.

To construct this relaxed SBQP dual, we appeal to the fact that the high approximate degrees of AND_m and f are themselves witnessed by dual objects ψ and φ, respectively. Building on a construction of Sherstov [38], we show how to combine ψ and φ to produce a new dual witness to the high SBQP degree of the composed function $F = \mathsf{AND}_m \circ f^m$. Note that Sherstov's original construction was tailored to showing only that F has high approximate degree. Since we need a stronger SBQP degree lower bound, our analysis needs to first take advantage of the fact that ψ actually witnesses the high *one-sided* approximate degree of AND_m. That is, it rules out one-sided approximations which tightly approximate AND_m on false inputs, but may be arbitrarily unbounded on true inputs. This manifests in the dual witness ψ as the property that it places zero weight on the set $\{x : \mathsf{AND}_m(x) = -1 \wedge \psi(x) > 0\}$. Second, our analysis needs to keep tighter track of how errors on the "true" side of F accumulate to show that it has the required almost one-sided error.

With Theorem 2 established, Theorem 1 follows by taking the f to be an explicit, constant-width, polynomial-size DNF with approximate degree $\Omega(n^{1-\delta})$ as constructed by Sherstov [40]. We obtain the results in Theorem 3 by building on the ideas of Sherstov and Thaler [41]. For every parameter m, the functions $k\mathsf{ED}_n$ and SURJ_n contain, as subfunctions, $\mathsf{AND}_m \circ (k\mathsf{ED}_{n/m})^m$ and $\mathsf{AND}_m \circ \mathsf{SURJ}^m_{n/m}$, respectively. We can thus invoke our general theorem on these subfunctions for appropriate choices of m that balance the final witness length and query costs in our lower bound.

The reason why Sherstov and Thaler could use this insight to prove lower bounds for element distinctness and permutation testing, but not for k-element distinctness or surjectivity, owes to a gap between the type of polynomial approximation lower bounds they could prove on functions of the form $\mathsf{AND}_m \circ f^m$ and the type they identified as yielding SBQP approximation lower bounds. Specifically, they showed a lower bound on the vanishing-error approximate degree $\widetilde{\deg}_\varepsilon(\mathsf{AND}_m \circ f^m)$, but needed to invoke a stronger lower bound on the *one-sided* vanishing-error approximate degree $\widetilde{\deg}^-_\varepsilon(\mathsf{AND}_m \circ f^m)$ to rule out SBQP approximations. For some functions with special structure, including AND, element distinctness, and permutation testing, ordinary approximate degree and one-sided approximate degree fortuitously coincide, and hence vanishing-error approximate degree lower bounds imply $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds. But such a coincidence is not (yet) known for k-distinctness or surjectivity, and false for other simple functions such as the two-level AND-OR tree. By working directly with SBQP approximations via our relaxed SBQP duals, we can prove $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds for arbitrary composed functions $\mathsf{AND}_m \circ f^m$, regardless of whether ordinary and one-sided approximate degree agree.

1.2 Lifting to QMA Communication Lower Bounds

In the QMA *communication* model, denoted by $\mathrm{QMA}^{\mathrm{cc}}$, Alice holds an input $x \in \{-1,1\}^n$, Bob holds an input $y \in \{-1,1\}^n$, and they wish to compute a joint function $f : \{-1,1\}^n \times \{-1,1\}^n \to \{-1,1\}$ of their inputs. Bob receives a witness of w qubits from the prover and the parties have unlimited access to shared entanglement. They wish to minimize the amount of communication between them (including the witness) required to compute $f(x,y)$. The success criterion, cost of a protocol, and $\mathrm{QMA}^{\mathrm{cc}}$ complexity of a function f are defined analogously to the query model.

Sherstov's pattern matrix method [37] shows how to generically "lift" an approximate degree lower bound for a function f to a communication lower bound for a related function $G := f \circ g^n$, where g is a certain constant-sized two-party gadget. By adapting the pattern matrix method to show how to lift SBQP degree lower bounds to QMA communication lower

bounds, we establish QMA lower bounds for communication analogs of the batch verification problems we study. Our analysis builds closely on prior generalizations of the pattern matrix method lifting one-sided approximate degree to (multiparty) MA communication [19, 40] and to QMA communication [28], but the formulation of the statement that is most useful to us appears to be new. We defer these results to the full version of this paper [8].

1.3 Open Questions

- Is the witness-query tradeoff given by Theorem 2 tight? That is, can one find a function f such that for every $w = O(m)$, $\mathsf{AND}_m \circ f^m$ admits a $\mathrm{QMA}^{\mathrm{dt}}$ protocol with witness length w and query complexity $O(\mathrm{adeg}_{1-1/m}(f) \cdot \sqrt{m/w})$?
- Our $\mathrm{QMA}^{\mathrm{dt}}$ batch verification lower bounds depend on approximate degree lower bounds for the base function f. Recall that approximate degree is a lower bound on the quantum query complexity of f, but can be much smaller. Can one obtain similar batch verification lower bounds starting from arbitrary quantum query lower bounds on f?
- How tight are the $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds in Theorem 3? Can our ideas be used to obtain improved $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds for element distinctness and permutation testing?

2 Preliminaries

We let $\mathbb{N} = \{1, 2, \ldots\}$ denote the set of natural numbers and $\mathbb{R}$ denote the set of real numbers. We write $[n]$ for the set $\{1, 2, \ldots, n\}$.

We identify the Boolean values "true" and "false" with -1 and 1 respectively. The functions $\mathsf{OR}_n : \{-1,1\}^n \to \{-1,1\}$ and $\mathsf{AND}_n : \{-1,1\}^n \to \{-1,1\}$ are the usual functions $\mathsf{OR}_n(x) = \bigvee_{i=1}^n x_i$ and $\mathsf{AND}_n(x) = \bigwedge_{i=1}^n x_i$. The Hamming weight of a string $x \in \{-1,1\}^n$, denoted $|x|$, is the number of indices i such that $x_i = -1$. A function f is said to be computable by a *disjunctive normal form (DNF)* formula of size s and width w if there exist $S_1, S_2, \ldots, S_s \subseteq [n]$ with cardinality w such that $f(x) = \bigvee_{i \in [s]} \bigwedge_{j \in S_i} x_j$. The conjunctions $\bigwedge_{j \in S_i} x_j$ are called *terms*.

We use the following standard real-valued functions. We use $\exp_k$ to denote the base-k exponential, i.e., $\exp_k(x) = k^x$. The base-2 logarithm is denoted by $\log(\cdot)$. For any predicate P, $\mathbb{1}[P(x)]$ equals 1 if $P(x)$ is true and 0 otherwise. The notation $|x|$ represents the absolute value, Hamming weight, or cardinality of x depending on if x is a real number, binary string, or set respectively.

For real valued functions f and g with the same finite support X, we use the usual notation for the ℓ_1 norm and inner product.

$$\|f\|_1 = \sum_{x \in X} |f(x)|$$
$$\langle f, g \rangle = \sum_{x \in X} f(x) g(x)$$

We often refer to the inner product $\langle f, g \rangle$ as the *correlation* between f and g. The function f is *balanced* if $\sum_{x \in X} f(x) = 0$.

We use superscripts to denote the direct product of a function with itself: for every $f : X \to Y$ and $m \in \mathbb{N}$, $f^m : X^m \to Y^m$ is the function $f^m(x_1, \ldots, x_m) = (f(x_1), \ldots, f(x_m))$. The symbol $\circ$ is used for function composition, so that $(f \circ g)(x) = f(g(x))$. Consequently, for every $g : X \to Y$ and $f : Y^m \to Z$, $f \circ g^m$ denotes the "block composition" of f and g: $(f \circ g^m)(x_1, \ldots, x_m) = f(g(x_1), \ldots, g(x_m))$.

For any set S, we use 2^S to denote the set of all subsets of S. For every $T \subseteq [n]$, we denote by $\chi_T : \{-1,1\}^n \to \{-1,1\}$ the Fourier character $\chi_T(x) = \prod_{i \in T} x_i$. We recall that for every function $f : \{-1,1\}^n \to \mathbb{R}$ there exists a unique function $\widehat{f} : 2^{[n]} \to \mathbb{R}$ called the *Fourier transform* of f such that for all x,

$$f(x) = \sum_{T \subseteq [n]} \widehat{f}(T)\chi_T(x).$$

Furthermore, this unique function is given by

$$\widehat{f}(T) = 2^{-n} \sum_{x \in \{-1,1\}^n} f(x)\chi_T(x).$$

2.1 Polynomial approximation

We say that a polynomial p *ε-approximates* a function $f : \{-1,1\}^n \to \{-1,1\}$ if for all $x \in \{-1,1\}^n$, $|f(x) - p(x)| \leq \varepsilon$. The *$\varepsilon$-approximate degree*, denoted $\mathrm{adeg}_\varepsilon(f)$, of f is the minimum degree of a polynomial that ε-approximates f. A polynomial p *ε-one-sided approximates* f if the following conditions hold.

1. For every $x \in f^{-1}(-1)$, we have $p(x) \in (-\infty, -1 + \varepsilon]$.
2. For every $x \in f^{-1}(1)$, we have $p(x) \in [1 - \varepsilon, 1 + \varepsilon]$.

The *ε-one-sided approximate degree* of f, denoted $\mathrm{adeg}^-_\varepsilon(f)$, is the minimum degree of a polynomial that ε-one-sided approximates f. One can also define a complementary notion of one-sided approximation where the approximating polynomial must tightly approximate true inputs and is allowed to be arbitrarily positive on false inputs, though we do not consider that notion or its associated one-sided approximate degree measure in this paper. For both one-sided and regular approximate degree, when ε is not specified, we assume that $\varepsilon = 1/3$.

We recall the standard dual formulation of approximate degree.

▶ **Theorem 4** ([37, 11]). *Let $f : \{-1,1\}^n \to \{-1,1\}$ be a Boolean function. Let $d \geq 1$ be an integer and $\varepsilon > 0$ be a real number. Then $\mathrm{adeg}_\varepsilon(f) > d$ if and only if there exists a function $\psi : \{-1,1\}^n \to \mathbb{R}$ such that*

1. *$\langle f, \psi \rangle \geq \varepsilon$,*
2. *$\|\psi\|_1 = 1$, and*
3. *for every $T \subseteq [n]$ with $|T| \leq d$, $\langle \psi, \chi_T \rangle = 0$.*

Furthermore, $\mathrm{adeg}^-_\varepsilon(f) > d$ if and only if there exists a function ψ that satisfies the conditions above and has the additional property that for all x such that $f(x) = -1$, $\psi(x) \leq 0$.

The function ψ is called a *dual polynomial* for the lower bound $\mathrm{adeg}_\varepsilon(f) > d$ or $\mathrm{adeg}^-_\varepsilon(f) > d$ (depending on whether it has the additional property). Condition 3 is often referred to as ψ having "pure high degree at least d."

Sherstov showed that there exist functions computable by polynomial-size, constant-width DNF formulas that have near-maximal approximate degree.

▶ **Theorem 5** ([40]). *For all constants $\delta > 0$ and $c \geq 1$, for every $n \in \mathbb{N}$, there is an explicit function $f_n : \{-1,1\}^n \to \{-1,1\}$ computable by a polynomial-size, constant-width DNF such that $\mathrm{adeg}_{1-1/n^c}(f_n) = \Omega(n^{1-\delta})$.*

"Explicit" here means that there is a uniform algorithm that, given n, prints out a description of f_n in time polynomial in n.

The *surjectivity* function $\mathsf{SURJ}_{N,R} : \{-1,1\}^{N \log R} \to \{-1,1\}$ interprets its input as a list of N numbers in $[R]$ and evaluates to true if and only if every element of $[R]$ appears in the list at least once. As a default we let $n = N \log R$ and $R = N/2$. We write SURJ_n for $\mathsf{SURJ}_{N,R}$ with the default setting of parameters. Bun, Kothari, and Thaler [9] showed that the approximate degree of SURJ_n is $\widetilde{\Omega}(n^{3/4})$.

The *element distinctness* problem $\mathsf{ED}_{N,R} : \{-1,1\}^{N \log R} \to \{-1,1\}$ interprets its input as a list of N numbers in $[R]$ and evaluates to true if and only if no element appears more than once. The canonical setting of parameters is $N = R$, as the function is trivially false when $N > R$ and it was shown by Ambainis [4] that the ε-approximate degree of $\mathsf{ED}_{N,R}$ is the same for all $N \leq R$. Aaronson and Shi [3] showed that $\mathrm{adeg}(\mathsf{ED}_n) = \widetilde{\Theta}(n^{2/3})$ when R is large compared to N. The aforementioned result of Ambainis extended their bounds to our setting of parameters.

The *k-element distinctness* problem $k\mathsf{ED}_{N,R} : \{-1,1\}^{N \log R} \to \{-1,1\}$ is a generalization of element distinctness that evaluates to true if and only if no range element appears k or more times. We use the same canonical setting of parameters as for element distinctness: $k\mathsf{ED}_n = k\mathsf{ED}_{N,N}$ where $n = N \log N$. Mande, Thaler, and Zhu [29] showed that $\mathrm{adeg}(k\mathsf{ED}_n) = \widetilde{\Omega}(n^{3/4 - 1/(4k)})$.

Finally, for our lower bound on the $\mathrm{QMA}^{\mathrm{dt}}$ complexity of the two-level AND-OR tree we will need the lower bound $\mathrm{adeg}(\mathsf{AND}_m \circ \mathsf{OR}_n) = \Omega(\sqrt{mn})$, proved by Bun and Thaler [10] and Sherstov [39].

2.2 Quantum Merlin-Arthur protocols

Let $f : \{-1,1\}^n \to \{-1,1\}$ be a function. A *QMA query protocol* for f with witness length w, query complexity q, completeness error ε_1, and soundness error ε_2, is a quantum oracle algorithm A such that

1. For all $x \in f^{-1}(-1)$ there exists an w-qubit "witness" $|\psi\rangle$ such that $\Pr_A[A^x(|\psi\rangle) = -1] \geq 1 - \varepsilon_1$;
2. For all $x \in f^{-1}(1)$, for every w-qubit state $|\psi\rangle$, $\Pr_A[A^x(|\psi\rangle) = -1] \leq \varepsilon_2$; and
3. For all x and w, $A^x(w)$ makes at most q queries to the input x.

The algorithm A is often referred to as a *verifier*. If a protocol has completeness (resp. soundness) error 0, we say that it has perfect correctness (resp. soundness). When the completeness and/or soundness errors are not specified, we set them to the default value 1/3. The QMA query complexity of f, denoted $\mathsf{QMA}^{\mathsf{dt}}(f)$, is the minimum $w + q$ such that there exists a QMA query protocol for f with witness length w and query complexity q.

QMA communication complexity is defined analogously. A QMA communication protocol for a two-party function $f : X \times Y \to \{-1,1\}$ with witness length w, communication cost c, completeness error ε_1, and soundness error ε_2 is a two-party quantum communication protocol (A, B) with access to shared entanglement such that for all inputs (x, y) and w-qubit states $|\psi\rangle$, the execution $[A(x), B(y, |\psi\rangle)]$ involves at most c qubits of communication and satisfies completeness and soundness conditions analogous to the query case. For a more formal definition of this setting we refer the reader to [33]. The QMA communication complexity of f, denoted $\mathsf{QMA}^{\mathsf{cc}}(f)$ is the minimum $w + c$ such that there exists a QMA communication protocol for f with witness length w and communication cost c.

In models with classical witnesses (such as MA or QCMA), one can reduce completeness and soundness errors by repeatedly executing the verifier and outputting the majority output. This does not work for protocols where the witness is a quantum state, since executing the verifier might alter the witness state. Nevertheless, Marriott and Watrous showed that one

can, in fact, reduce the completeness and soundness errors for QMA protocols. The proof of this theorem can be adapted to achieve the same error reduction in the QMA communication setting.

▶ **Theorem 6** ([30]). *Let $f : \{-1,1\}^n \to \{-1,1\}$ be a function. Suppose there exists a $q(n)$-query quantum algorithm Q, numbers $a(n), b(n) \in [0,1]$, and $w(n), p(n) \leq \mathsf{poly}(n)$ such that*

1. *for all $x \in f^{-1}(-1)$ there exists a $w(n)$-qubit state $|\psi\rangle$ such that $\Pr_Q[Q^x(|\psi\rangle)$ accepts$] \geq a(n)$,*
2. *for all $x \in f^{-1}(1)$, for all $w(n)$-qubit states $|\psi\rangle$, $\Pr_Q[Q^x(|\psi\rangle)$ accepts$] \leq b(n)$, and*
3. *$a(n) - b(n) \geq 1/p(n)$.*

Then for every $r(n) \leq \mathsf{poly}(n)$ there exists a $\mathrm{QMA}^{\mathrm{dt}}$ protocol for f with completeness and soundness error $2^{-r(n)}$, witness length $w(n)$, and query complexity $O(q(n) \cdot r(n) \cdot p(n)^2)$.

2.3 Linear algebra

Our results on communication complexity require some standard concepts from matrix analysis which we recall here. Unless otherwise specified, all matrices we consider are over the field of real numbers. We use the notation $[A_{x,y}]_{x\in X, y\in Y}$ to refer to a $|X| \times |Y|$ matrix A whose rows and columns are indexed by elements of X and Y respectively and whose entries are $A_{x,y}$. For any two matrices $[A_{x,y}]_{x\in X,y\in Y}$ and $[B_{x,y}]_{x\in X,y\in Y}$, the inner product of A and B is defined as

$$\langle A, B\rangle = \sum_{x\in X, y\in Y} A_{x,y}B_{x,y}.$$

The singular values of an $m \times n$ matrix A are denoted by $\sigma_1(A) \geq \sigma_2(A) \geq \cdots \geq \sigma_{\min\{m,n\}}(A) \geq 0$. We recall four standard notions of matrix norm.

$$\|A\|_1 = \sum_{x,y} |A_{x,y}| \qquad \text{(The } L_1 \text{ norm)}$$

$$\|A\| = \max_i \sigma_i(A) \qquad \text{(The spectral norm)}$$

$$\|A\|_\Sigma = \sum_i \sigma_i(A) \qquad \text{(The trace norm)}$$

$$\|A\|_F = \sqrt{\sum_i \sigma_i(A)^2} \qquad \text{(The Frobenius norm)}$$

We will use the following well-known identities. For all compatible real matrices A and B,

$$\|AB\|_\Sigma \leq \|A\|_F \|B\|_F, \text{ and}$$
$$\langle A, B\rangle \leq \|A\|_\Sigma \|B\|.$$

3 SBQP Lower Bounds from High Approximate Degree

In this section we show that if a function f has high large-error approximate degree, then $\mathsf{AND}_m \circ f^m$ has high SBQP cost. We start by defining a notion of SBQP cost in terms of polynomial approximations and deriving a corresponding dual notion for it.

▶ **Definition 7.** *A polynomial (M,ε)-SBQP-approximates $f : \{-1,1\}^n \to \{-1,1\}$ if for all $x \in \{-1,1\}^n$,*

1. *if $f(x) = -1$ then $-M - 1 - \varepsilon \leq p(x) \leq -1 + \varepsilon$, and*
2. *if $f(x) = 1$ then $1 - \varepsilon \leq p(x) \leq 1 + \varepsilon$.*

▶ **Lemma 8.** *For all $f : \{-1,1\}^n \to \{-1,1\}$ and $M^*, \varepsilon, d \geq 0$, if there exists a function $\varphi : \{-1,1\}^n \to \mathbb{R}$ with the following properties, then there does not exist a polynomial of degree at most d that (M^*, ε)-SBQP-approximates f.*

1. $\langle \varphi, f \rangle \geq 2\varepsilon$,
2. $\|\varphi\|_1 \leq 1$,
3. *for all $S \subseteq [n]$ such that $|S| \leq d$, $\langle \varphi, \chi_S \rangle = 0$,*
4. $\sum_{x : f(x) = -1, \varphi(x) > 0} \varphi(x) \leq \varepsilon / M^*$.

Proof. Let $f : \{-1,1\}^n \to \{-1,1\}$ be arbitrary. Fix any $\varepsilon \geq 0$. The smallest M such that there exists a degree-d (M, ε)-SBQP-approximation to f is given by the following linear program.

$$\begin{aligned} \min_{\substack{M \\ p \text{ of deg } < d}} \quad & M & \\ \text{s.t.} \quad & p(x) \geq f(x) - \varepsilon & \forall x \in f^{-1}(1) \\ & p(x) + M \geq f(x) - \varepsilon & \forall x \in f^{-1}(-1) \\ & -p(x) \geq -f(x) - \varepsilon & \forall x \in \{-1,1\}^n \\ & M \geq 0 & \end{aligned}$$

The dual of this linear program is the following.

$$\begin{aligned} \max_{\psi^+, \psi^- : \{-1,1\}^n \to \mathbb{R}} \quad & \sum_{x \in \{-1,1\}^n} (\psi^+(x) - \psi^-(x)) f(x) - \varepsilon \sum_{x \in \{-1,1\}^n} (\psi^+(x) + \psi^-(x)) & \\ \text{s.t.} \quad & \left\langle \psi^+ - \psi^-, \chi_S \right\rangle = 0 & \forall S \subseteq [n], |S| \leq d \\ & \sum_{x : f(x) = -1} \psi^+(x) \leq 1 & \\ & \psi^-, \psi^+ \geq 0 & \end{aligned}$$

By a standard transformation (which appears, e.g., in the derivation of the dual formulation of approximate degree), any optimal solution to the dual program can be turned into one where ψ^+ and ψ^- have disjoint support. Thus, we can set $\psi = \psi^+ - \psi^-$ to get the following equivalent dual linear program.

$$\begin{aligned} \max_{\psi : \{-1,1\}^n \to \mathbb{R}} \quad & \langle \psi, f \rangle - \varepsilon \|\psi\|_1 & \\ \text{s.t.} \quad & \langle \psi, \chi_S \rangle = 0 & \forall S \subseteq [n], |S| \leq d \\ & \sum_{x : f(x) = -1, \psi(x) > 0} \psi(x) \leq 1 & \end{aligned}$$

By weak duality, to show that f does not admit an (M^*, ε)-SBQP-approximation it is sufficient to show that the dual program has objective at least M^*. Suppose we have a φ as in the theorem statement. Consider $\psi = M^* \varphi / \varepsilon$. This ψ satisfies the first constraint of the dual since φ has pure high degree at least d. It satisfies the second constraint since $\sum_{x : f(x) = -1, \psi(x) > 0} \psi(x) = \frac{M^*}{\varepsilon} \sum_{x : f(x) = -1, \varphi(x) > 0} \varphi(x) \leq 1$. Finally, the objective attained by ψ is $\frac{M^*}{\varepsilon} \langle \varphi, f \rangle - \varepsilon \cdot \frac{M^*}{\varepsilon} \|\varphi\|_1 \geq 2M^* - M^* = M^*$. ◀

Functions that satisfy the conditions in Lemma 8 are central to our proofs. We call a φ that has those properties with respect to a function f and parameters M, ε, and d a *relaxed (M, ε, d)-SBQP dual to f*. The word "relaxed" refers to the fact that while the existence of

φ is sufficient for a lower bound on the degree of an (M, ε)-SBQP approximation, it might not be necessary. This is because the proof of Lemma 8 loses tightness when it replaces ψ, which is obtained via LP duality, by the weaker (but easier to reason about) object φ. Just as SBQP approximations are natural restrictions of one-sided approximations (allowing large but bounded error on one side), property 4 of the dual is a natural relaxation of the one-sidedness property of adeg^- duals described in Theorem 4. We use relaxed SBQP duals to prove that if a function f has high approximate degree, then all SBQP approximations to $\mathsf{AND}_m \circ f^m$ must have high degree. To do so, we use the following well-known fact about the vanishing-error one-sided approximate degree of the AND function.

▶ **Theorem 9** ([7]). *There exist universal constants δ_{AND} and δ'_{AND} such that for every $\gamma \in [2^{-m}, 1/3]$, $\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)} < \mathrm{adeg}^-_\gamma(\mathsf{AND}_m) \leq \delta'_{\mathsf{AND}}\sqrt{m\log(1/\gamma)}$.*

Strictly speaking, [7] only implies $\mathrm{adeg}_\gamma(\mathsf{AND}_m) = \Theta(\sqrt{m\log(1/\gamma)})$, but it is easy to see that one-sidedness is without loss of generality for the AND function. (This observation goes back to Gavinsky and Sherstov [19].) Suppose $\mathrm{adeg}_\gamma(\mathsf{AND}_m) > d \geq 1$, and let ψ be a dual witness to that lower bound guaranteed by Theorem 4. Since ψ has pure high degree at least 1, we get that $0 = \langle \psi, 1\rangle = \psi(-\mathbf{1}_m) + \sum_{z \neq -\mathbf{1}_m} \psi(z)$. That ψ has correlation at least γ with AND_m combined with the fact that $-\mathbf{1}_m$ is the only -1 instance of AND_m gives us $\gamma \leq \langle \psi, \mathsf{AND}_m\rangle = -\psi(-\mathbf{1}_m) + \sum_{z \neq -\mathbf{1}_m} \psi(z)$. Combining the two equations yields $\psi(-\mathbf{1}_m) \leq -\gamma/2 < 0$, which means ψ is a witness for $\mathrm{adeg}^-_\gamma(\mathsf{AND}_m) > d$. The other ingredient in our construction is the following polynomial p_k.

▶ **Lemma 10** ([38]). *Let k be an even number in $[m-1]$, and define $p_k : [-1,1]^m \to \mathbb{R}$ as the unique degree-k multilinear polynomial such that for all $z \in \{-1,1\}^m$,*

$$p_k(z) = \prod_{i=1}^{k}(|z| - i).$$

Then $p_k(z) \geq 0$ for all $z \in [-1,1]^m$ and for every $\eta \in [0,1)$,

$$\mathop{\mathbb{E}}_{z\sim\Pi(\eta)}[p_k(z)] \leq p_k(\mathbf{1}_m)(1-\eta)^m\left\{1 + \frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1}\right\},$$

where $\Pi(\eta)$ is the distribution on $\{-1,1\}^m$ where each coordinate is -1 with probability η and 1 with probability $1-\eta$.

We now prove our main connection between approximate degree and SBQP degree. Our proof works by way of constructing a relaxed SBQP dual for $\mathsf{AND}_m \circ f^m$ from (1) a dual witness for the vanishing-error one-sided approximate degree of AND, (2) a dual for the assumed approximate degree lower bound on f, and (3) the polynomials p_k. The construction itself is a renormalized special case of the dual constructed in [38, Theorem 6.1]. Consequently, the correlation, ℓ_1 norm, and pure high degree analyses for the dual follow from the corresponding analyses in that work. Additionally, to get an SBQP dual, we need to show that the dual construction does not place too much mass on "true" inputs that it misclassifies. We do so by taking advantage of the fact that the dual for AND is one-sided and keeping tighter track of how errors on misclassified "true" inputs accumulate.

▶ **Theorem 11.** *Let n be a function of m, and let $\{f : \{-1,1\}^{n(m)} \to \{-1,1\}\}_{m\in\mathbb{N}}$ be a family of functions. Let $\gamma \in [2^{-m}, 1/3]$, $\varepsilon > 0$, $M > 0$, and $k \in \mathbb{N}$ be such that k is even and*

$$\varepsilon \leq \frac{\gamma}{2} - \frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} \text{ and} \tag{1}$$

$$\frac{\varepsilon}{M} \geq \frac{\eta^{k+1}}{2(1-\eta)^m}\binom{m}{k+1}. \tag{2}$$

If $\mathrm{adeg}_{1-\eta}(f) > d \geq 1$ then for sufficiently large m, every (M,ε)-SBQP approximation to $F = \mathsf{AND}_m \circ f^m$ has degree $\Omega(d \cdot (\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)} - k))$.

Proof. Fix n, $\eta \in (0,1]$, $f : \{-1,1\}^n \to \{-1,1\}$, $d \geq 1$, and suppose $\mathrm{adeg}_{1-\eta}(f) > d$. Suppose γ, ε, M, and k are such that equations 1 and 2 hold. By Lemma 8, it suffices to construct a relaxed $(M, \varepsilon, \Omega(d \cdot (\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)} - k)))$-SBQP dual Φ to F.

By Theorems 4 and 9 there is a dual witness $\psi : \{-1,1\}^m \to \mathbb{R}$ for $\mathrm{adeg}^-_\gamma(\mathsf{AND}_m) \geq \delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)}$. Similarly, there is a dual $\varphi : \{-1,1\}^n \to \mathbb{R}$ witnessing $\mathrm{adeg}_{1-\eta}(f)$. Let $\widetilde{\mu}$ be the distribution on $\{-1,1\}^n$ with probability mass function $\widetilde{\mu}(y) = |\varphi(y)|$, and let μ be a distribution consisting of m independent samples from $\widetilde{\mu}$. For every $z \in \{-1,1\}^m$ define μ_z to be μ conditioned on $(\ldots, \operatorname{sgn}\varphi(x_i), \ldots) = z$.

Note that $\mathbb{E}_{y\sim\widetilde{\mu}}[\operatorname{sgn}\varphi(y)] = \sum_{y\in\{-1,1\}^n} |\varphi(y)| \operatorname{sgn}\varphi(y) = \langle\varphi, \chi_\emptyset\rangle$. Since φ has pure high degree $d \geq 1$, $\langle\varphi,\chi_\emptyset\rangle = 0$ and so $\Pr_{y\sim\widetilde{\mu}}[\varphi(y) > 0] = \Pr_{y\sim\widetilde{\mu}}[\varphi(y) < 0]$. In other words, if $y \sim \widetilde{\mu}$, then $\operatorname{sgn}\varphi(y)$ is a uniformly random bit. Consequently, if $x \sim \mu$, then $(\ldots, \operatorname{sgn}\varphi(x_i), \ldots)$ is uniformly distributed over $\{-1,1\}^m$.

Since φ is a dual for $\mathrm{adeg}_{1-\eta}(f) > d$, $\langle\varphi, f\rangle \geq 1-\eta$. Thus $1-\eta \leq \sum_{y\in\{-1,1\}^n} f(y)\varphi(y) = \mathbb{E}_{y\sim\widetilde{\mu}}[f(y)\operatorname{sgn}\varphi(y)]$. Notice that for all y, $\frac{1-f(y)\operatorname{sgn}\varphi(y)}{2}$ is equal to 1 if $\operatorname{sgn}\varphi(y) \neq f(y)$ and 0 otherwise. Thus

$$\Pr_{y\sim\mu}[\operatorname{sgn}\varphi(y) \neq f(y)] = \mathop{\mathbb{E}}_{y\sim\widetilde{\mu}}\left[\frac{1-f(y)\operatorname{sgn}\varphi(y)}{2}\right] \leq \frac{\eta}{2}.$$

Now define

$$\eta^+ = \Pr_{y\sim\mu}[f(y) \neq \operatorname{sgn}\varphi(y) \mid \varphi(y) > 0],$$
$$\eta^- = \Pr_{y\sim\mu}[f(y) \neq \operatorname{sgn}\varphi(y) \mid \varphi(y) < 0],$$

and $\alpha : \{-1,1\}^n \to [-1,1]$ as

$$\alpha(y) = \begin{cases} \frac{1-2\eta+\eta^-}{1-\eta^-} & \text{if } f(y) = \operatorname{sgn}\varphi(y) = -1 \\ \frac{1-2\eta+\eta^+}{1-\eta^+} & \text{if } f(y) = \operatorname{sgn}\varphi(y) = 1 \\ -1 & \text{otherwise.} \end{cases}$$

The range of α is $[-1,1]$ because $\max\{\eta^+,\eta^-\} \leq \eta$, which holds since $1-\eta \leq 2\Pr_{y\sim\widetilde{\mu}}[\operatorname{sgn}\varphi(y) \neq f(y)] = \frac{1}{2}(1-2\eta^+) + \frac{1}{2}(1-2\eta^-)$. For $x \in (\{-1,1\}^n)^m$, we write $\alpha^m(x)$ to mean $(\ldots, \alpha(x_i), \ldots)$. Note that for $z \in \{-1,1\}^m$ and $i \in [m]$, $\mathbb{E}_{x\sim\mu_z}[\alpha(x_i)] = 1-2\eta$.

We are now ready to present the dual. Let $p_k : [-1,1]^m \to [0,\infty)$ be the degree-k multilinear polynomial given by Lemma 10. Our relaxed SBQP dual $\Phi : (\{-1,1\}^n)^m \to \mathbb{R}$ is defined as

$$\Phi(x_1,\ldots,x_m) = \frac{2^m}{p_k(\ldots,1-2\eta,\ldots)} \cdot \psi(\ldots,\operatorname{sgn}\varphi(x_i),\ldots) \cdot p_k(\alpha^m(x)) \prod_{i=1}^m |\varphi(x_i)|.$$

As we now describe, the required correlation, ℓ_1 norm, and pure high degree properties follow from the corresponding claims in [38]. First, [38, Claim 6.3] tells us that

$$\langle F, \Phi \rangle \geq \gamma - \frac{2\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} \geq 2\varepsilon,$$

where the last inequality applies Eq. 1. That $\|\Phi\|_1 = 1$ follows from [38, Claim 6.2]. Finally, [38, Eq. 6.7] and the fact that ψ has pure high degree at least $\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)}$ implies that Φ has pure high degree $\Omega(d \cdot (\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)} - k))$.

It remains to show that Φ does not place too much mass on -1 instances that it "misclassifies." Notice that the only term in the definition of Φ that is not non-negative is $\psi(\ldots, \operatorname{sgn}\varphi(x_i), \ldots)$. For convenience, let $\beta = p_k(\ldots, 1-2\eta, \ldots)$. We can write the error on -1 instances as

$$\begin{aligned}
&\sum_{x:F(x)=-1,\Phi(x)>0} \Phi(x) = \sum_{x\in(\{-1,1\}^n)^m} \Phi(x)\cdot \mathbb{1}[F(x)=-1 \wedge \Phi(x)>0] \\
&= \frac{2^m}{\beta} \sum_{x\in(\{-1,1\}^n)^m} \psi(\ldots,\operatorname{sgn}\varphi(x_i),\ldots)\cdot p_k(\alpha^m(x_i))\cdot \mathbb{1}[F(x)=-1\wedge\Phi(x)>0]\prod_{i=1}^m |\varphi(x_i)| \\
&= \frac{2^m}{\beta} \mathop{\mathbb{E}}_{x\sim\mu}\left[\psi(\ldots,\operatorname{sgn}\varphi(x_i),\ldots)\cdot p_k(\alpha^m(x))\cdot \mathbb{1}[F(x)=-1\wedge\Phi(x)>0]\right] \\
&= \frac{1}{\beta} \sum_{z\in\{-1,1\}^m} \mathop{\mathbb{E}}_{x\sim\mu_z}\left[\psi(z)\cdot p_k(\alpha^m(x))\cdot \mathbb{1}[F(x)=-1\wedge\psi(z)>0]\right] \\
&= \frac{1}{\beta} \sum_{z:\psi(z)>0} \psi(z) \mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot \mathbb{1}[F(x)=-1]\right]. \qquad (3)
\end{aligned}$$

The second-to-last equality holds since $\Pr_{x\sim\mu}[(\ldots,\operatorname{sgn}\varphi(x_i),\ldots)=z] = 2^{-m}$. Fix z. By multilinearity of p_k, $\mathbb{E}_{x\sim\mu_z}[p_k(\alpha^m(x))] = \beta$. Thus,

$$\begin{aligned}
&\left|\mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\mathbb{1}[F(x)=-1]\right] - \beta\cdot\mathbb{1}[\mathsf{AND}_m(z)=-1]\right| \\
&\qquad = \left|\mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\mathbb{1}[F(x)=-1] - p_k(\alpha^m(x))\cdot\mathbb{1}[\mathsf{AND}_m(z)=-1]\right]\right| \\
&\qquad \leq \mathop{\mathbb{E}}_{x\sim\mu_z}\left[\,|p_k(\alpha^m(x))|\cdot\left|\mathbb{1}[F(x)=-1] - \mathbb{1}[\mathsf{AND}_m(z)=-1]\right|\right].
\end{aligned}$$

Recall that $F(x) = \mathsf{AND}_m(\ldots, f(x_i), \ldots)$. So $|\mathbb{1}[F(x)=-1] - \mathbb{1}[\mathsf{AND}_m(z)=-1]|$ is 1 when exactly one of $(\ldots, f(x_i), \ldots)$ and z is equal to $-\mathbf{1}_m$ and 0 otherwise. By the non-negativity of p_k on inputs in $[-1,1]^m$, it follows that

$$\begin{aligned}
&\left|\mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\mathbb{1}[F(x)=-1]\right] - \beta\cdot\mathbb{1}[\mathsf{AND}_m(z)=-1]\right| \\
&\qquad \leq \mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\mathbb{1}[(\ldots,f(x_i),\ldots)\neq z]\right] \\
&\qquad = \mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\big(1-\mathbb{1}[(\ldots,f(x_i),\ldots)=z]\big)\right] \\
&\qquad = \beta - \mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\mathbb{1}[(\ldots,f(x_i),\ldots)=z]\right].
\end{aligned}$$

Now [38, Eq. 6.8] shows that

$$\mathop{\mathbb{E}}_{x\sim\mu_z}\left[p_k(\alpha^m(x))\cdot\mathbb{1}[(\ldots,f(x_i),\ldots)=z]\right] \geq (1-\eta)^m p_k(\mathbf{1}_m),$$

which means

$$\left| \mathop{\mathbb{E}}_{x\sim\mu_z} \left[p_k(\alpha^m(x)) \cdot \mathbb{1}[F(x) = -1] \right] - \beta \cdot \mathbb{1}[\mathsf{AND}_m(z) = -1] \right| \leq \beta - (1-\eta)^m p_k(\mathbf{1}_m).$$

Combining this bound with Eq. 3, we get

$$\begin{aligned} \sum_{x:F(x)=-1,\Phi(x)>0} \Phi(x) &\leq \frac{1}{\beta} \sum_{z:\psi(z)>0} \psi(z)\big(\beta \cdot \mathbb{1}[\mathsf{AND}_m(z) = -1] + \beta - (1-\eta)^m p_k(\mathbf{1}_m)\big) \\ &= \left(\sum_{z:\psi(z)>0} \psi(z) \right) \left(1 - \frac{(1-\eta)^m p_k(\mathbf{1}_m)}{\beta} \right) \\ &= \frac{1}{2} \left(1 - \frac{(1-\eta)^m p_k(\mathbf{1}_m)}{\beta} \right). \end{aligned}$$

The second-to-last equality uses the one-sidedness of ψ: whenever $\mathsf{AND}_m(z) = -1$, $\psi(z) < 0$. The last equality uses the fact that ψ is balanced.

Now let Π be a distribution on $\{-1,1\}^m$ where each coordinate is an independent, uniformly random element of $\{-1,1\}$. Then by multilinearity, $\mathbb{E}_{z\sim\Pi}[p_k(z)] = \beta$. By Lemma 10,

$$\begin{aligned} \frac{(1-\eta)^m p_k(\mathbf{1}_m)}{\beta} &\geq \left(1 + \frac{\eta^{k+1}}{(1-\eta)^m} \binom{m}{k+1} \right)^{-1} \\ &\geq 1 - \frac{\eta^{k+1}}{(1-\eta)^m} \binom{m}{k+1}. \end{aligned}$$

Plugging this back in and using Eq. 2,

$$\sum_{x:F(x)=-1,\Phi(x)>0} \Phi(x) \leq \frac{\eta^{k+1}}{2(1-\eta)^m} \binom{m}{k+1} \leq \frac{\varepsilon}{M}.$$ ◀

The ensuing corollaries apply Theorem 11 with specific parameters to obtain more "user-friendly" connections between the approximate degree of f and the SBQP degree of $\mathsf{AND}_m \circ f^m$.

▶ **Corollary 12.** *Let n be a function of m, and let $\{f : \{-1,1\}^{n(m)} \to \{-1,1\}\}_{m\in\mathbb{N}}$ be a family of functions. Let $c > 0$ be a constant, $d \geq 1$, and $\varepsilon \in [2^{-m-2}, 1/12]$ such that either $c \geq 1$ or $\varepsilon = 2^{-\Omega(m)}$. If $\widetilde{\deg}_{1-\frac{1}{m^c}}(f) > d$, then there exists a constant $\tau > 0$ such that for sufficiently large m, every (M,ε)-SBQP approximation to $F = \mathsf{AND}_m \circ f^m$ with $M \leq \varepsilon 2^{\tau\sqrt{m\log(1/\varepsilon)}\log m}$ must have degree $\Omega(d\sqrt{m\log(1/\varepsilon)})$. The constant τ depends only on c when $c \geq 1$, and on c as well as the constant in the Ω when $c \in (0,1)$ but $\varepsilon = 2^{-\Omega(m)}$.*

Proof. Let n, f, c, d, and ε be as in the corollary statement. Let $\tau > 0$ be a constant to be set later in this proof. We apply Theorem 11 with

$$\begin{aligned} \eta &= \frac{1}{m^c}, \\ \gamma &= 4\varepsilon, \\ M &= \varepsilon 2^{\tau\sqrt{m\log(1/\varepsilon)}\log m}, \text{ and} \\ k &= \text{the least even number greater than } \frac{\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)}}{2}. \end{aligned}$$

This will imply that every (M,ε)-SBQP approximation to $\mathsf{AND}_m \circ f^m$ has degree at least $\Omega(d \cdot (\delta_{\mathsf{AND}}\sqrt{m\log(1/\gamma)} - k)) = \Omega(d\sqrt{m\log(1/\varepsilon)})$. The lower bound extends to smaller M immediately since an (M',ε)-SBQP approximation with $M' < M$ is also an (M,ε)-SBQP approximation.

We now get an upper bound on the quantity $\frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1}$, which will allow us to show that the premises of Theorem 11 are satisfied. First, suppose $c \geq 1$. We have $(1-\eta)^m = ((1-1/m^c)^{m^c})^{m^{1-c}}$. For large m, this is at least $(1/2e)^{m^{1-c}} \geq 1/2e$. By Theorem 9 there is a universal constant C such that $k \geq C\sqrt{m\log(1/\gamma)}$. We have

$$\frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} \leq 2e \cdot m^{-c\,(k+1)} \cdot m^{k+1} = 2em^{-(c-1)(k+1)} = 2^{-\Omega(k\log m)}.$$

The last bound holds because $c \geq 1$ and hence $m^{-(c-1)(k+1)} = 2^{-\Omega(k\log m)}$.

Now suppose $c \in (0,1)$ but $\varepsilon = 2^{-\Omega(m)}$. Since $\gamma = 4\varepsilon$, there is a constant $C_1 \in (0,1]$ such that $\gamma = 2^{-C_1 m}$. Also, $k+1 \leq C_2 m$ for an appropriate constant $C_2 \in (0,1)$ that depends on δ_{AND} and C_1. Using Stirling's approximation,

$$\begin{aligned}
\frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} &\leq \frac{m^{-c\,(k+1)}}{(1-m^{-c})^m}\left(\frac{me}{k+1}\right)^{k+1} \\
&= \frac{(m^c)^{m-(k+1)}}{(m^c-1)^m}\left(\frac{e}{C_2}\right)^{C_2 m} \\
&= \exp_2\left((1-C_2)cm\log m - m\log(m^c-1) + C_2 m\log(e/C_2)\right) \\
&= 2^{-\Omega(m\log m)} = 2^{-\Omega(k\log m)},
\end{aligned}$$

where the last bound holds since for large m, $(1-C_2)cm\log m < m\log(m^c+1)$.

Thus in either case, it follows from $k \geq \frac{\delta_{\mathsf{AND}}}{2}\sqrt{m\log(1/\gamma)}$ that for sufficiently large m, $\frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} \leq \gamma/4$. Thus $\frac{\gamma}{2} - \frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} \geq \frac{\gamma}{4} = \varepsilon$, which establishes Eq. 1.

Eq. 2 holds similarly. Since $k \geq \frac{\delta_{\mathsf{AND}}}{2}\sqrt{m\log(1/\gamma)} \geq C_3\sqrt{m\log(1/\varepsilon)}$ (for an appropriate universal constant C_3), there exists τ' such that

$$\frac{\eta^{k+1}}{2(1-\eta)^m}\binom{m}{k+1} \leq 2^{-\tau' C_3\sqrt{m\log(1/\varepsilon)}\log m}.$$

Setting $\tau = \tau' \cdot C_3$, this quantity is at most ε/M. ◀

▶ **Corollary 13.** *Let n be a function of m, and let $\{f : \{-1,1\}^{n(m)} \to \{-1,1\}\}_{m\in\mathbb{N}}$ be a family of functions. If $\mathrm{adeg}_{1/3}(f) > d \geq 1$, then for sufficiently large m, every $(2^{m-1}, 2^{-m-2})$-SBQP approximation to $F = \mathsf{AND}_m \circ f^m$ must have degree $\Omega(dm)$.*

Proof. By standard error reduction techniques (see, e.g., [13, Section 3.4]), it is possible to transform any degree-k 0.99-approximation to f into a degree-$O(k)$ 1/3-approximation to f. Thus it suffices to prove the corollary assuming $\mathrm{adeg}_{0.99}(f) > d \geq 1$; the corollary as stated follows with at the cost of a constant factor in the degree lower bound.

We use the well-known fact that $\mathrm{adeg}^{-}_{2^{-m}}(\mathsf{AND}_m) \geq m$ (note the absence of a constant factor); the dual witness justifying this is simply the normalized parity function on m bits. Let $f : \{-1,1\}^n \to \{-1,1\}$ be such that $\mathrm{adeg}_{0.99}(f) > d \geq 1$. We apply Theorem 11 with

$$\eta = \frac{1}{100},$$
$$\gamma = 2^{-m},$$
$$\varepsilon = 2^{-m-2},$$
$$M = 2^{m-1}, \text{ and}$$
$$k = \frac{m}{2} - 1.^{2}$$

By Stirling's approximation,

$$\begin{aligned}\frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} &\le \frac{1}{100^{k+1}} \cdot \frac{100^m}{99^m} \cdot \left(\frac{me}{k+1}\right)^{k+1} \\ &= \frac{100^{m/2}}{99^m} \cdot (2e)^{m/2} && (k+1 = m/2) \\ &= \left(\frac{\sqrt{200e}}{99}\right)^m \le 2^{-2m}. && (\sqrt{200e}/99 \le 1/4)\end{aligned}$$

Thus for large m,

$$\frac{\gamma}{2} - \frac{\eta^{k+1}}{(1-\eta)^m}\binom{m}{k+1} \ge 2^{-m-1} - 2^{-2m} \ge 2^{-m-2} = \varepsilon \text{ and}$$
$$\frac{\eta^{k+1}}{2(1-\eta)^m}\binom{m}{k+1} \le 2^{-2m-1} = \frac{\varepsilon}{M},$$

satisfying the premises of Theorem 11. ◀

4 QMA$^{\mathrm{dt}}$ Lower Bounds

In this section, we apply Corollaries 12 and 13 to obtain various QMA$^{\mathrm{dt}}$ lower bounds. Our proofs rely on the following connection between QMA$^{\mathrm{dt}}$ protocols and SBQP approximations.

▶ **Lemma 14.** *If $f : \{-1,1\}^n \to \{-1,1\}$ has a* QMA$^{\mathrm{dt}}$ *protocol with witness length w and query complexity q, then for every $\varepsilon > 0$ there is an $O(q \cdot \max\{w, \log(1/\varepsilon)\})$-degree polynomial that $(2^{w+2}, \varepsilon)$-SBQP approximates f.*

Proof. Suppose f has a QMA$^{\mathrm{dt}}$ protocol with witness length w and query complexity q. Fix $\varepsilon > 0$, and let $\nu = \min\{\frac{\varepsilon}{2}, 2^{-2w}\}$. By Theorem 6, there is a QMA$^{\mathrm{dt}}$ protocol P for f with completeness and soundness error at most ν, witness length w, and query complexity $O(q \log(1/\nu))$. Let Q be the witness-free protocol that simulates P with the maximally mixed state as the witness and accepts if and only if P outputs -1. By soundness of P, whenever $f(x) = 1$, $\Pr_Q[Q^x \text{ accepts}] \le \nu$. On the other hand, on inputs where $f(x) = -1$, the probability that Q accepts is at least $(1-\nu)2^{-w} \ge (1-2^{-2w})2^{-w} \ge 2^{-w-1}$.

Let $\widetilde{p}$ be the $O(q \log(1/\nu))$-degree polynomial computing the acceptance probability of Q. Define $p = 1 - 2 \cdot 2^{w+1}\widetilde{p}$. The degree of p is $O(q \log(1/\nu)) = O(q \cdot \max\{w, \log(1/\varepsilon)\})$.

1. If $f(x) = -1$ then $2^{w+1}\widetilde{p}(x) \in [1, 2^{w+1}]$ and so $p(x) \in [1 - 2^{w+2}, -1] \subset [-2^{w+2} - 1 - \varepsilon, -1 + \varepsilon]$.
2. If $f(x) = 1$ then $2^{w+1}\widetilde{p}(x) \in [0, \nu]$ and so $p(x) \in [1 - 2\nu, 1] \subset [1 - \varepsilon, 1 + \varepsilon]$. ◀

[2] We assume for simplicity of exposition that k is an even natural number. The case where $m/2 - 1$ is not an even number follows by setting k to the least even number greater than $m/2 - 1$.

4.1 The QMA$^{\mathrm{dt}}$ Complexity of Batch Verification

Our first lower bound exhibits a smooth tradeoff between witness length and query complexity of batch verifying functions with high large-error approximate degree.

▶ **Theorem 15.** *Let n be a function of m, and $\{f : \{-1,1\}^{n(m)} \to \{-1,1\}\}_{m\in\mathbb{N}}$ be an arbitrary family of functions. For all sufficiently large m, if $\mathrm{adeg}_{1-1/m}(f) > d \geq 1$ then every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $w = O(m)$ requires $\Omega(d\sqrt{m/w})$ queries.*

Proof. Fix $n(m)$ and $f : \{-1,1\}^{n(m)} \to \{-1,1\}$. Assume $\mathrm{adeg}_{1-1/m}(f) > d \geq 1$. Suppose there is a $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $w = O(m)$ and query complexity q. We assume that $w \geq \log(12)$; the case where $w < \log(12)$ follows immediately with a smaller constant in the Ω. Let $\varepsilon = \max\{2^{-w}, 2^{-m-2}\} \in [2^{-m-2}, 1/12]$. The premises of Corollary 12 are satisfied with $c = 1$. Thus, by Corollary 12, there exists a constant τ such that every (M,ε)-SBQP approximation to $\mathsf{AND}_m \circ f^m$ with $M \leq \varepsilon 2^{\tau\sqrt{m\log(1/\varepsilon)\log m}}$ has degree $\Omega(d\sqrt{m\log(1/\varepsilon)})$.

By Lemma 14, there is a degree-$O(qw)$ polynomial p that $(2^{w+2},\varepsilon)$-SBQP approximates $\mathsf{AND}_m \circ f^m$. Since $\log(1/\varepsilon) = \Omega(w)$, we have $\tau\sqrt{m\log(1/\varepsilon)}\log m = \Omega(\sqrt{mw}\log m)$. Thus for large m, $2^{w+2} \leq \varepsilon 2^{\tau\sqrt{m\log(1/\varepsilon)\log m}}$. So $O(qw) = \deg(p) = \Omega(d\cdot\sqrt{m\log(1/\varepsilon)}) = \Omega(d\sqrt{mw})$. Rearranging yields the desired lower bound. ◀

Observe that the proof of Theorem 15 can actually tolerate witness lengths all the way up to $\tau' m\log m$ for some fixed constant $\tau' \in (0,1)$. Roughly speaking, it shows that whenever $w \leq \tau' m\log m$, any $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length w must have query complexity at least $\Omega(\mathrm{adeg}_{1-1/m}(f)\sqrt{m/w})$. Combining this observation with Theorem 5 we show that even a constant-factor improvement in witness length over the trivial protocol for batch verifying DNFs incurs a large increase in query complexity.

▶ **Corollary 16.** *For all sufficiently large n and all $\delta > 0$ there exists an explicitly given function $f : \{-1,1\}^n \to \{-1,1\}$ computable by a polynomial-size, constant-width DNF and a constant $k \in \mathbb{N}$ such that for all $m = n^{O(1)} \leq n$, the following hold.*

1. *There is a $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $km\log n$ and query complexity $O(k\sqrt{m})$.*
2. *There is a constant $\tau' \in (0,1)$ such that every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length at most $\tau' m\log n$ must have query complexity $\min\{\Omega(n^{1-\delta}\sqrt{m/w}), \Omega(n^{1-\delta}m/w)\}$.*

Proof. By Theorem 5 for every n and $\delta > 0$ there exists an explicitly given function $f : \{-1,1\}^n \to \{-1,1\}$ computable by a polynomial-size, constant-width DNF such that $\mathrm{adeg}_{1-1/n}(f) > n^{1-\delta}$. Let $k \in \mathbb{N}$ be a constant such that the DNF computing f has width at most k and size a most n^k. Henceforth we refer to f and the DNF computing it interchangeably. Let $m = n^a$ for some constant $a \in (0,1]$. Consider the following $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$.

1. The witness consists of indices $i_1, i_2, \ldots, i_m \in ([n^k])^m$ such that the i_jth term in the jth copy of f is satisfied.
2. The verifier, given query access to x and the witness $(i_1, i_2, \ldots, i_m)$ uses Grover's search algorithm to look for $j \in [m]$ such that the i_jth term of the jth copy of f is not satisfied.
3. If such an index is found, the verifier rejects. Else, it accepts.

If $(\mathsf{AND}_m \circ f^m)(x) = -1$ then the verifier, given the honest witness, will not find an unsatisfied term among those indicated by the witness. Thus the protocol has perfect completeness. If $(\mathsf{AND}_m \circ f^m)(x) = 1$ then there must be some j such that the jth copy of f is not satisfied.

It follows that none of the terms of f are satisfied in that copy, and so given any witness $(i_1, i_2, \ldots, i_m)$ the i_jth term of the jth copy of f will not be satisfied. By properties of Grover search, the verifier finds that index and rejects with probability $2/3$. Thus the protocol is sound. The witness length of the protocol is at most $km \log n$ since each term requires $k \log n$ bits to describe. Whether or not each term is satisfied can be computed exactly with k queries. One can thus implement the search over these checks with $O(k\sqrt{m})$ queries using standard techniques (see, e.g., [21]). This proves the upper bound.

The lower bound holds due to the proof of Theorem 15. Since $m \leq n$, $\widetilde{\deg}_{1-1/m}(f) \geq \widetilde{\deg}_{1-1/n}(f) > n^{1-\delta}$. For convenience, let $d = n^{1-\delta}$. Setting $\varepsilon = \max\{2^{-w}, 2^{-m-2}\}$, we can apply Corollary 12 as in the proof of Theorem 15. This yields a constant τ such that every (M, ε)-SBQP approximation to $\mathsf{AND}_m \circ f^m$ with $M \leq \varepsilon 2^{\tau\sqrt{m \log(1/\varepsilon)} \log m}$ must have degree $\Omega(d\sqrt{m \log(1/\varepsilon)}) = \Omega(d\sqrt{m \cdot \min\{m, w\}})$. Due to our setting of ε and m, there exists a constant $\tau'' > 0$ such that the degree lower bound holds for all (M, ε)-SBQP approximations with $M \leq 2^{\tau'' a\sqrt{m \cdot \min\{m-w\}} \log n}$. Set $\tau' = \tau'' a/2$. Suppose there is a $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $w \leq \tau' m \log n$. Then by Lemma 14 there is a $(2^{w+2}, \varepsilon)$-SBQP approximation to $\mathsf{AND}_m \circ f^m$ with degree $O(q \min\{w, \max\{w, m\}\}) = O(qw)$. Combining the upper and lower bounds, $q = \Omega(d\sqrt{m \cdot \min\{m, w\}}/w) = \min\{\Omega(d\sqrt{m/w}), \Omega(dm/w)\}$. ◀

Our techniques also allow us to get query complexity lower bounds for batch verifying f from lower bounds on $\widetilde{\deg}_{1-1/m^c}(f)$ for arbitrarily small constants $c > 0$ and on $\widetilde{\deg}_{1/3}(f)$.

▶ **Theorem 17.** *Let n be a function of m and let $\{f : \{-1,1\}^{n(m)} \to \{-1,1\}\}_{m \in \mathbb{N}}$ be an arbitrary family of functions. The following statements hold for all sufficiently large m.*

1. *For every $c > 0$, if $\widetilde{\deg}_{1-1/m^c}(f) > d \geq 1$ then every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length at most $O(m)$ requires $\Omega(d)$ queries.*
2. *If $\widetilde{\deg}_{1/3}(f) > d \geq 1$ then every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length at most $m - 3$ requires $\Omega(d)$ queries.*

Proof.

(Part 1). Fix $c > 0$ and $f : \{-1,1\}^{n(m)} \to \{-1,1\}$. Assume $\widetilde{\deg}_{1-1/m^c}(f) > d \geq 1$, and suppose there is a $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $w = O(m)$ and query complexity q. Assume that $w \geq \log(12)$; the complementary case follows immediately. Let $\varepsilon = 2^{-m-2}$. By Corollary 12 there is a constant τ that depends only on c such that every (M, ε)-SBQP approximation to $\mathsf{AND}_m \circ f^m$ with $M \leq \varepsilon 2^{\tau\sqrt{m \log(1/\varepsilon)} \log m}$ must have degree $\Omega(d\sqrt{m \log(1/\varepsilon)}) = \Omega(dm)$.

By Lemma 14 there is a degree-$O(qm)$ polynomial that $(2^{w+2}, \varepsilon)$-SBQP approximates $\mathsf{AND}_m \circ f^m$. We have $\varepsilon 2^{\tau\sqrt{m \log(1/\varepsilon)} \log m} = 2^{\Omega(m \log m)} \geq 2^{O(m)} = 2^{w+2}$. So we must have $O(qm) = \deg(p) = \Omega(dm)$, which implies $q = \Omega(d)$ as claimed.

(Part 2). Suppose $\widetilde{\deg}_{1/3}(f) > d \geq 1$, and suppose there is a $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ f^m$ with witness length $w \leq m - 3$ and query complexity q. By Corollary 13, for sufficiently large m, every $(M, 2^{-m-2})$-SBQP approximation with $M \leq 2^{m-1}$ must have degree $\Omega(dm)$. By Lemma 14 there is a degree-$O(qm)$ polynomial that $(2^{w+2}, 2^{-m-2})$-SBQP approximation to $\mathsf{AND}_m \circ f^m$. Since $w \leq m - 3$, $2^{w+2} \leq 2^{m-1}$. The query lower bound follows. ◀

4.2 $\mathrm{QMA}^{\mathrm{dt}}$ Lower Bounds for Natural Functions

Some functions have the property that they contain as sub-functions the problem of batch verifying smaller copies of themselves. For such functions, a lower bound on the $\mathrm{QMA}^{\mathrm{dt}}$ complexity of batch verification immediately implies a lower bound on the $\mathrm{QMA}^{\mathrm{dt}}$ complexity

of the functions themselves. In this subsection, we use this insight along with Theorem 17 to get $\mathrm{QMA}^{\mathrm{dt}}$ lower bounds for element distinctness, surjectivity, and k-distinctness. Our lower bounds make use of the observation by Sherstov and Thaler [41] that the surjectivity and k-element distinctness functions contain the batch verification problem for smaller instances of themselves as a subfunction.

▶ **Theorem 18.**

1. $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_m \circ \mathsf{OR}_n^m) = \Omega((mn)^{1/3})$. *Consequently,* $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_{n^{1/3}} \circ \mathsf{OR}_{n^{2/3}}^{n^{1/3}}) = \Omega(n^{1/3})$.
2. $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{SURJ}_n) = \widetilde{\Omega}(n^{3/7})$.
3. $\mathsf{QMA}^{\mathsf{dt}}(k\mathsf{ED}_n) = \widetilde{\Omega}(n^{\frac{3k-1}{7k-1}})$.

Before presenting the proof of our lower bounds, we briefly mention lower bounds that follow directly from prior work. For $\mathsf{AND}_{n^{1/3}} \circ \mathsf{OR}_{n^{2/3}}^{n^{1/3}}$, the best known lower bound is $\Omega(n^{1/6})$, which follows – via the connection between $\mathrm{QMA}^{\mathrm{dt}}$ complexity and threshold degree in [17] – from the fact that the threshold degree of $\mathsf{AND}_{n^{1/3}} \circ \mathsf{OR}_{n^{2/3}}^{n^{1/3}}$ is $\Omega(n^{1/3})$. For SURJ_n, the best known lower bound is $\widetilde{\Omega}(n^{1/4})$, due to a threshold degree lower bound of $\widetilde{\Omega}(\sqrt{n})$ [12]. Finally, for $k\mathsf{ED}_n$, we have $\mathsf{QMA}^{\mathsf{dt}}(k\mathsf{ED}_n) \geq \mathsf{QMA}^{\mathsf{dt}}(\mathsf{ED}_n) = \widetilde{\Omega}(n^{2/5})$, which was proved implicitly in [41].

Proof of Theorem 18.

(Part 1). Note that for every m' that divides m, $\mathsf{AND}_m \circ \mathsf{OR}_n^m = \mathsf{AND}_{m'} \circ (\mathsf{AND}_{m/m'} \circ \mathsf{OR}_n^{m/m'})^{m'}$. By a lower bound from Section 2, $\mathrm{adeg}(\mathsf{AND}_{m/m'} \circ \mathsf{OR}_n^{m/m'}) = \Omega(\sqrt{mn/m'})$. By Theorem 17, every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_{m'} \circ (\mathsf{AND}_{m/m'} \circ \mathsf{OR}_n^{m/m'})^{m'}$ with witness length at most $m' - 3$ must have query complexity $\Omega(\sqrt{mn/m'})$. Setting $m' = (mn)^{1/3}$, we get that every $\mathrm{QMA}^{\mathrm{dt}}$ protocol or $\mathsf{AND}_m \circ \mathsf{OR}_n^m$ either has witness length at least $(mn)^{1/3} - 3$ or query complexity $\Omega(\sqrt{mn/(mn)^{1/3}}) = \Omega((mn)^{1/3})$, which means $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_m \circ \mathsf{OR}_n^m) = \Omega((mn)^{1/3})$.

(Part 2). Set $R = N/2$ as is typical. Observe that for all m that divide $N/2$, $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{SURJ}_{N,R}) \geq \mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_m \circ \mathsf{SURJ}_{N/m,R/m}^m)$. This is due to the fact that evaluating $\mathsf{AND}_m \circ \mathsf{SURJ}_{N/m,R/m}^m$ amounts to evaluating whether m different lists of numbers from $[R/m]$ each contain every $r \in [R/m]$. One way to do this is to, for all $i \in [m]$, add $(i-1)R/m$ to each number in the ith list, concatenate all the lists, and then evaluate $\mathsf{SURJ}_{N,R}$ on the concatenated list. It follows that any $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{SURJ}_{N,R}$ can be used to solve $\mathsf{AND}_m \circ \mathsf{SURJ}_{N/m,R/m}^m$ with no additional cost. Recall that $\mathrm{adeg}(\mathsf{SURJ}_{N/m,R/m}) = \widetilde{\Omega}((N/m)^{3/4})$. By Theorem 17, every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ \mathsf{SURJ}_{N/m,R/m}^m$ with witness length at most $m - 3$ must have query complexity $\widetilde{\Omega}((N/m)^{3/4})$. Setting $m = N^{3/7}$, every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ \mathsf{SURJ}_{N/m,R/m}^m$ must either have witness length at least $N^{3/7} - 3$ or query complexity at $\widetilde{\Omega}(N^{3/7})$. Hence $\mathsf{QMA}^{\mathsf{dt}}(\mathsf{SURJ}_n) = \widetilde{\Omega}(N^{3/7})$, which is $\widetilde{\Omega}(n^{3/7})$.

(Part 3). Like in the case of surjectivity, $\mathsf{QMA}^{\mathsf{dt}}(k\mathsf{ED}_{N,N}) \geq \mathsf{QMA}^{\mathsf{dt}}(\mathsf{AND}_m \circ k\mathsf{ED}_{N/m,N/m}^m)$. When one shifts the input to $\mathsf{AND}_m \circ k\mathsf{ED}_{N/m,N/m}^m$ as in part 2, the shifted input is a "yes" instance of $k\mathsf{ED}_{N,N}$ if and only if the original instance is a "yes" instance of $\mathsf{AND}_m \circ k\mathsf{ED}_{N/m,N/m}^m$. Recall that $\mathrm{adeg}(k\mathsf{ED}_{N/m,N/m}) = \widetilde{\Omega}((N/m)^{3/4-1/(4k)})$. So Theorem 17 tells us that every $\mathrm{QMA}^{\mathrm{dt}}$ protocol for $\mathsf{AND}_m \circ k\mathsf{ED}_{N/m,N/m}^m$ with witness length at most $m - 3$ must have query complexity $\widetilde{\Omega}((N/m)^{3/4-1/(4k)})$. Setting $m = n^{\frac{3k-1}{7k-1}}$ yields the desired lower bound. ◀

References

1. Scott Aaronson. Impossibility of succinct quantum proofs for collision-freeness. *Quantum Info. Comput.*, 12(1–2):21–28, 2012. doi:10.26421/QIC12.1-2-3.
2. Scott Aaronson, Robin Kothari, William Kretschmer, and Justin Thaler. Quantum lower bounds for approximate counting via Laurent polynomials. In *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, 2020.
3. Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004. doi:10.1145/1008731.1008735.
4. Andris Ambainis. Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory Comput.*, 1(1):37–46, 2005. doi:10.4086/TOC.2005.V001A003.
5. Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM (JACM)*, 48(4):778–797, 2001. doi:10.1145/502090.502097.
6. Nir Bitansky, Chethan Kamath, Omer Paneth, Ron D Rothblum, and Prashant Nalini Vasudevan. Batch proofs are statistically hiding. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 435–443, 2024. doi:10.1145/3618260.3649775.
7. Harry Buhrman, Richard Cleve, Ronald de Wolf, and Christof Zalka. Bounds for small-error and zero-error quantum algorithms. In *40th Annual Symposium on Foundations of Computer Science, FOCS 1999, New York, NY, USA, October 17-18, 1999*, 1999. doi:10.1109/SFFCS.1999.814607.
8. Mark Bun, Mandar Juvekar, and Samuel King. QMA lower bounds for batch verification via approximate degree. *arXiv preprint*, 2026. arXiv:2607.08888.
9. Mark Bun, Robin Kothari, and Justin Thaler. The polynomial method strikes back: Tight quantum query bounds via dual polynomials. *Theory Comput.*, 16:1–71, 2020. doi:10.4086/TOC.2020.V016A010.
10. Mark Bun and Justin Thaler. Dual lower bounds for approximate degree and Markov–Bernstein inequalities. *Inf. Comput.*, 243:2–25, 2015. doi:10.1016/J.IC.2014.12.003.
11. Mark Bun and Justin Thaler. Hardness amplification and the approximate degree of constant-depth circuits. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, 2015. doi:10.1007/978-3-662-47672-7_22.
12. Mark Bun and Justin Thaler. The large-error approximate degree of ac^0. *Theory Comput.*, 17:1–46, 2021. doi:10.4086/TOC.2021.V017A007.
13. Mark Bun and Justin Thaler. Approximate degree in classical and quantum computing. *Foundations and Trends in Theoretical Computer Science*, 15(3–4):229–423, 2022. doi:10.1561/0400000107.
14. Jeffrey Champion and David J Wu. Non-interactive zero-knowledge from non-interactive batch arguments. In *Annual International Cryptology Conference*, pages 38–71. Springer, 2023. doi:10.1007/978-3-031-38545-2_2.
15. Arka Rai Choudhuri, Abhihsek Jain, and Zhengzhong Jin. SNARGs for P from LWE. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, 2022. doi:10.1109/FOCS52979.2021.00016.
16. Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Non-interactive batch arguments for NP from standard assumptions. In *Annual International Cryptology Conference*, 2021.
17. Marcel Dall'Agnol, Tom Gur, Subhayan Moulik, and Justin Thaler. Quantum proofs of proximity. *Quantum*, 6:834, 2022. doi:10.22331/Q-2022-10-13-834.
18. Lalita Devadas, Rishab Goyal, Yael Kalai, and Vinod Vaikuntanathan. Rate-1 non-interactive arguments for batch-NP and applications. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, 2022.

19 Dmitry Gavinsky and Alexander A. Sherstov. A separation of NP and coNP in multiparty communication complexity. *Theory Comput.*, 6(1):227–245, 2010. doi:10.4086/TOC.2010.V006A010.

20 Rishab Goyal, Aditya Jain, and Shashwatha Mitra GB. Succinct arguments for batch QMA and friends under 8 rounds. In *Annual International Cryptology Conference*, 2025.

21 Peter Høyer, Michele Mosca, and Ronald de Wolf. Quantum search on bounded-error inputs. In *International Colloquium on Automata, Languages, and Programming*, 2003.

22 Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Boosting batch arguments and ram delegation. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1545–1552, 2023. doi:10.1145/3564246.3585200.

23 Yael Tauman Kalai, Omer Paneth, and Lisa Yang. How to delegate computations publicly. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, pages 1115–1124, 2019. doi:10.1145/3313276.3316411.

24 Yael Tauman Kalai, Vinod Vaikuntanathan, and Rachel Yun Zhang. Somewhere statistical soundness, post-quantum security, and SNARGs. In *Theory of Cryptography Conference*, 2021.

25 Inbar Kaslasi, Guy N Rothblum, Ron D Rothblum, Adam Sealfon, and Prashant Nalini Vasudevan. Batch verification for statistical zero knowledge proofs. In *Theory of Cryptography Conference*, 2020.

26 Inbar Kaslasi, Ron D Rothblum, and Prashant Nalini Vasudevan. Public-coin statistical zero-knowledge batch verification against malicious verifiers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 219–246. Springer, 2021. doi:10.1007/978-3-030-77883-5_8.

27 Or Keret, Ron D Rothblum, and Prashant Nalini Vasudevan. Doubly-efficient batch verification in statistical zero-knowledge. In *Theory of Cryptography Conference*, 2024.

28 Hartmut Klauck. On Arthur Merlin games in communication complexity. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity, CCC 2011, San Jose, California, USA, June 8-10, 2011*, 2011. doi:10.1109/CCC.2011.33.

29 Nikhil S. Mande, Justin Thaler, and Shuchen Zhu. Improved approximate degree bounds for k-distinctness. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2020, Riga, Latvia, June 9-12, 2020*, 2020. doi:10.4230/LIPIcs.TQC.2020.2.

30 Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *Comput. Complex.*, 14(2):122–152, 2005. doi:10.1007/S00037-005-0194-X.

31 Changrui Mu, Shafik Nassar, Ron D Rothblum, and Prashant Nalini Vasudevan. Strong batching for non-interactive statistical zero-knowledge. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2024.

32 Shafik Nassar and Ronak Ramachandran. Quantum statistical witness indistinguishability. *arXiv preprint*, 2025. arXiv:2509.01945.

33 Ran Raz and Amir Shpilka. On the power of quantum proofs. In *19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA*, 2004. doi:10.1109/CCC.2004.1313849.

34 Omer Reingold, Guy N Rothblum, and Ron D Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 49–62, 2016. doi:10.1145/2897518.2897652.

35 Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Efficient batch verification for UP. In *33rd Computational Complexity Conference (CCC 2018)*, 2018. doi:10.4230/LIPIcs.CCC.2018.22.

36 Guy N Rothblum and Ron D Rothblum. Batch verification and proofs of proximity with polylog overhead. In *Theory of Cryptography Conference*, 2020.

37 Alexander A. Sherstov. The pattern matrix method. *SIAM J. Comput.*, 40(6):1969–2000, 2011. doi:10.1137/080733644.

38 Alexander A. Sherstov. Strong direct product theorems for quantum communication and query complexity. *SIAM J. Comput.*, 41(5):1122–1165, 2012. doi:10.1137/110842661.

39 Alexander A. Sherstov. Approximating the AND-OR tree. *Theory Comput.*, 9:653–663, 2013. doi:10.4086/TOC.2013.V009A020.

40 Alexander A. Sherstov. The approximate degree of DNF and CNF formulas. *SIAM J. Comput.*, 54(3):702–774, 2025. doi:10.1137/23M1557593.

41 Alexander A. Sherstov and Justin Thaler. Vanishing-error approximate degree and QMA complexity. *Chicago Journal of Theoretical Computer Science*, 2023(3), 2023.

42 Brent Waters and David J Wu. Batch arguments for NP and more from standard bilinear group assumptions. In *Annual International Cryptology Conference*, pages 433–463, 2022. doi:10.1007/978-3-031-15979-4_15.