A one-query lower bound for unitary synthesis and breaking quantum cryptography

Alex Lombardi* Fermi Ma† John Wright‡

Abstract

The Unitary Synthesis Problem (Aaronson-Kuperberg 2007) asks whether any n-qubit unitary U can be implemented by an efficient quantum algorithm A augmented with an oracle that *computes an arbitrary Boolean function* f. In other words, can the task of implementing any unitary be efficiently reduced to the task of implementing any Boolean function?

In this work, we prove a one-query lower bound for unitary synthesis. We show that there exist unitaries U such that no quantum polynomial-time oracle algorithm A^f can implement U, even approximately, if it only makes one (quantum) query to f. Our approach also has implications for quantum cryptography: we prove (relative to a random oracle) the existence of quantum cryptographic primitives that remain secure against all one-query adversaries A^f. Since such one-query algorithms can decide any language, solve any classical search problem, and even prepare any quantum state, our result suggests that implementing random unitaries and breaking quantum cryptography may be harder than all of these tasks.

To prove this result, we formulate unitary synthesis as an efficient challenger-adversary game, which enables proving lower bounds by analyzing the maximum success probability of an adversary A^f. Our main technical insight is to identify a natural spectral relaxation of the one-query optimization problem, which we bound using tools from random matrix theory.

We view our framework as a potential avenue to rule out polynomial-query unitary synthesis, and we state conjectures in this direction.

*Princeton. Email: alex.lombardi@princeton.edu. Research was done in part while the author was a Simons-Berkeley Postdoctoral Fellow.

†Simons Institute & UC Berkeley. Email: fermima1@gmail.com.

‡UC Berkeley. Email: jswright@berkeley.edu.

Contents

1 Introduction

This paper is about *unitary synthesis*, the task of implementing a given n-qubit unitary transformation U as a quantum circuit. Unitary synthesis is ubiquitous throughout quantum computing, since virtually any quantum computational task — be it preparing a state, performing a measurement, or transforming one state into another — can be done by implementing *some* unitary. Of course, not every unitary can be implemented efficiently. As a special case, consider the classical task of evaluating an $(n-1)$-bit Boolean function $f : \{0,1\}^{n-1} \to \{0,1\}$. This can be solved by implementing an n-qubit unitary transformation, namely the unitary $U : |x, b\rangle \mapsto |x, b \oplus f(x)\rangle$, and so Shannon's classic counting argument [Sha49] implies that even these unitaries require $\Omega(2^n/n)$ gates to implement. But are worst-case unitaries hard to compute only because they can solve hard classical problems? Or is it possible that unitaries could still be hard even if it were easy to solve all classical problems?

This question was first posed in 2006 in an influential work by Aaronson and Kuperberg [AK07], and it was later dubbed "the Unitary Synthesis Problem" by Aaronson in his 2016 Barbados lectures [Aar16]. Formally, they considered $\mathrm{poly}(n)$-size quantum oracle circuits $A^{(\cdot)}$ that have the ability to make quantum queries to an arbitrary Boolean function $f : \{0,1\}^{\ell} \to \{0,1\}$ on $\ell = \mathrm{poly}(n)$ bits. This gives these circuits the power to instantaneously compute any Boolean function of their choice, and Aaronson and Kuperberg asked if this power enables them to efficiently implement any unitary transformation as well. More concretely, they asked the following question.

> **The Unitary Synthesis Problem [AK07, Aar16]:** Is there a universal efficient oracle circuit $A^{(\cdot)}$ such that for any unitary U, there is a corresponding Boolean function f for which A^f implements U?

In other words, the Unitary Synthesis Problem asks whether the task of implementing an arbitrary unitary can be efficiently reduced to computing Boolean functions. Notably, if the answer turns out to be negative, this would give strong evidence (in the form of a black-box separation) that the hardest quantum problems are harder than the hardest classical problems.

Since it was first posed, the Unitary Synthesis Problem has become arguably *the* central open problem in the rapidly growing field of unitary complexity, which we will discuss in more detail in Sections 1.1 and 1.3 below. To date, there is no clear consensus on what the true complexity of unitary synthesis should be: for all we knew, it might require as little as one query to the oracle, or as many as $2^{\Omega(n)}$.

One reason this question is subtle is that algorithms that make just one query to an arbitrary Boolean function are already quite powerful. For example, it turns out that such algorithms can solve the *state synthesis problem*, in which the goal is to produce an arbitrary quantum state $|\psi\rangle$ [Aar16, INN+22, Ros23a] (see Section 1.3 for discussion). The state synthesis and unitary synthesis problems share a number of similarities, and there has been some speculation that extending state synthesis techniques could lead to positive results on unitary synthesis (for example, see [INN+22, Section 7.2]). An excellent treatment of these and related problems can be found in Aaronson's Barbados notes [Aar16], Rosenthal's Ph.D. thesis [Ros23b], as well as in recent course notes of Yuen [Yue22a, Yue22b].

How hard is unitary synthesis? There are several inefficient algorithms for the Unitary Synthesis Problem, the most basic of which queries an oracle $\widetilde{O}(2^{2n})$ times to learn a classical description of U and then implements it using $\widetilde{O}(2^{2n})$ gates. As noted by Yuen [Yue22a], this basic algorithm can be implemented with a single quantum query to f using the Bernstein-Vazirani algorithm [BV97],

at the expense of making the query extremely large: in particular, it requires a quantum query to a Boolean function $f : \{0,1\}^\ell \to \{0,1\}$ on inputs of length $\ell = \widetilde{O}(2^{2n})$. If we restrict to algorithms that only make efficient queries to f, i.e., queries that only evaluate f on $\ell = \text{poly}(n)$-length inputs, the best known query complexity is $O(2^{n/2})$, achieved by a Grover-style algorithm due to Rosenthal [Ros22].

On the other hand, prior to this work, no general query lower bound for the Unitary Synthesis Problem was known. There *is* a well-known lower bound due to Aaronson and Kuperberg [AK07] that rules out a certain class of one-query algorithms A^f, namely those that exactly implement a unitary operation on their first n qubits for all choices of f. More recently, Rosenthal [Ros22] proved a lower bound ruling out a different specialized class of many-query algorithms. We discuss both of these lower bounds further in Section 1.3. However, the problem of ruling out (or constructing) even *one-query* unitary synthesis algorithms has remained open since Aaronson and Kuperberg first posed it nearly two decades ago (cf. Open Problem 4 and footnote 13 in [AK07]).

In this work, we resolve this open question and prove the first one-query lower bound for the Unitary Synthesis Problem.

Theorem 1.1 (informal, see Theorem 4.18)**.** *There is no efficient oracle circuit $A^{(\cdot)}$ that approximately implements an arbitrary n-qubit unitary U by making one quantum query to a U-dependent Boolean function f.*

Our lower bound applies even if the oracle circuit is allowed to use an unbounded number of non-oracle gates and ancilla qubits. In fact, it even applies to circuits that are allowed to query f on inputs which are extremely long, but not *too* extremely long; technically, we require that $A^{(\cdot)}$ can only query Boolean functions $f : \{0,1\}^\ell \to \{0,1\}$ on $\ell = o(2^n)$ bits. Note that some restriction on the size of the queries is necessary, due to the one-query Bernstein-Vazirani-style algorithm mentioned above, which queries a Boolean function on $\ell = \widetilde{O}(2^{2n})$ bits. Finally, our lower bound also extends to circuits that query arbitrary functions f with $\text{poly}(n)$ bits of output, and to circuits that make $\text{poly}(n)$-many *non-adaptive* queries, i.e., queries of the form

$$|x_1\rangle |b_1\rangle \dots |x_t\rangle |b_t\rangle \mapsto |x_1\rangle |b_1 \oplus f_1(x_1)\rangle \dots |x_t\rangle |b_t \oplus f_t(x_t)\rangle .$$

This is because such queries can be simulated using a single query to a more complex function, via another Bernstein-Vazirani trick (see Remark 3.6).

We prove Theorem 1.1 by leveraging a connection between the unitary synthesis problem and *quantum cryptography*, as we discuss next.

1.1 Unitary synthesis and quantum cryptography

Background and motivation. The past few years have seen a surge of interest in so-called *inherently quantum problems*, which are computational tasks in which either the input is a quantum state, the output is a quantum state, or both. These include many of the most important tasks in quantum computing, such as breaking computationally-secure quantum bit commitments, performing quantum state tomography, preparing the ground state of a local Hamiltonian, and decoding black hole radiation. The central goal of this area is to classify these problems according to the computational resources needed to solve them. Normally, we would do so using the language of computational complexity theory. However, after initial classification attempts, a mysterious, recurring phenomenon has emerged: computational complexity theory appears to be completely unable to classify many of these problems at all.

As just one example of this phenomenon, let us look to the field of quantum cryptography, where some of the most exciting work involving inherently quantum problems is being done today. This is due to the remarkable discovery that certain quantum cryptographic primitives — such as pseudorandom states and quantum bit commitments — are sufficient for a wide array of cryptographic applications, and yet appear to be weaker than traditional "minimal" cryptographic assumptions such as one-way functions or pseudorandom generators (PRGs).

Pseudorandom states. Of these quantum primitives, we focus on single-copy pseudorandom states (PRSes), introduced by Ji, Liu, and Song [JLS18], which can be seen as a quantum analogue of PRGs.[1] Classically, a PRG is a set of $K \ll N := 2^n$ efficiently computable n-bit strings $\{x_k\}_{k\in[K]}$ in which a string $x_{\boldsymbol{k}}$ drawn uniformly at random from the set is computationally indistinguishable from a truly random n-bit string. Quantumly, a (single-copy) PRS is a set of $K \ll N$ efficiently computable n-qubit quantum states $\{|\psi_k\rangle\}_{k\in[K]}$ in which a state $|\psi_{\boldsymbol{k}}\rangle$ drawn uniformly at random from the set is computationally indistinguishable from a Haar random n-qubit state. Single-copy PRSes are known to imply the existence of quantum bit commitments [Yan22, MY22, BCQ23], which are a key ingredient in many cryptographic protocols, ranging from zero-knowledge proof systems [BG22, GJMZ23] to secure multiparty computation [GLSV21, BCKM21, AQY22].

With these definitions in mind, what can we say about the *computational complexity* of breaking cryptographic pseudorandomness? Classically, it is easy to see that secure PRGs do not exist if $\mathsf{P} = \mathsf{NP}$. In fact, there is a polynomial-time black-box (Turing, or even Karp) reduction $A^{(\cdot)}$ which can break PRGs given oracle access to a function $f : \{0,1\}^* \to \{0,1\}$ that decides an NP-complete language. This explains why proving the existence of unconditionally secure PRGs has so far been unsuccessful, as doing so would imply the breakthrough complexity theoretic lower bound $\mathsf{P} \neq \mathsf{NP}$.

In the quantum setting, what can we say about the computational complexity of breaking a PRS? Is there a complexity assumption that we can make, such as $\mathsf{BQP} = \mathsf{QMA}$, which would imply that PRSes can be broken in polynomial-time? The answer to this question is currently unknown, and the difficulty stems from the fact that the computational task associated with breaking a PRS is an inherently quantum problem. In particular, the adversary's goal is to distinguish between a pseudorandom state and a Haar random state, given one of the two at random—a quantum-input, classical-output task. On the other hand, traditional complexity classes such as P and PSPACE, and even quantum complexity classes such as QMA, only capture problems with classical inputs. For example, even though the witness for a QMA statement is a quantum state, the *input* to the problem is always a classical string, such as the description of a local Hamiltonian.

How hard is it to break quantum cryptography? As a result of this mismatch between classical-input and quantum-input problems, it is not at all clear how breaking a PRS is related to traditional complexity assumptions. For example, a recent work of Kretschmer, Qian, Sinha, and Tal [KQST23] has shown that the existence of PRSes is independent of the P-versus-NP question, at least in the oracle setting, by constructing an oracle relative to which PRSes exist but $\mathsf{P} = \mathsf{NP}$. However, [KQST23] derives security of their candidate PRS from the hardness of an oracle problem, $\mathsf{OR} \circ \mathsf{FORRELATION}$, which is easily solvable in PSPACE. Despite this, it is not clear whether a PSPACE oracle should be powerful enough to break every PRS — in fact, it is not even clear that an oracle for the *halting problem* would suffice.

[1] [JLS18] actually required a PRS to satisfy a stronger "many-copy" security notion, and subsequent works studied the weaker notion of single-copy security (e.g., [MY22]). We will not consider many-copy security in this work, but briefly mention that any polynomial-copy PRS can be broken given a *single query* to a PSPACE oracle [Kre21]. In contrast, our connection to the unitary synthesis problem makes essential use of the single-copy security notion.

This raises a tantalizing question: what if the existence of PRSes is independent of traditional complexity altogether? Could we show that breaking a PRS does not black-box reduce to deciding any language? Let us now relate this back to the Unitary Synthesis Problem. Given a PRS, there always exists a unitary U which one could use to break the PRS if one could implement it efficiently, namely any unitary which maps $\mathrm{span}\{|\psi_1\rangle, \ldots, |\psi_K\rangle\}$ to $\mathrm{span}\{|1\rangle, \ldots, |K\rangle\}$. If an efficient quantum oracle circuit $A^{(\cdot)}$ can synthesize such a U given oracle access to some Boolean function f, then the PRS can be efficiently broken relative to f.

Our second main result is to rule out any single-query algorithm for this task, relative to a random oracle.

Theorem 1.2 (informal, see Theorem 5.2)**.** *Relative to a random oracle, there exists a PRS (and a quantum bit commitment scheme) secure against all one-query oracle algorithms A^f for* every Boolean function f*.*

Theorem 1.2 offers the strongest evidence to date that the security of PRSes might be independent of *all* of traditional computational complexity. Our two results, when taken together, demonstrate the close connection between the Unitary Synthesis Problem and the security of PRSes; as we will see below, we essentially prove these two results simultaneously, because in constructing a PRS which cannot be broken with one query, we are implicitly constructing a unitary which cannot be synthesized with one query.

1.2 Our approach

We prove Theorems 1.1 and 1.2 by analyzing an oracle version of the single-copy PRS security game, which we call the "Oracle State Distinguishing Game" (see Section 3). To state this task, let us define two pieces of relevant notation. First, given a Boolean function $h : \{0,1\}^n \to \{\pm 1\}$, we define the corresponding *binary phase state* as

$$|\psi_h\rangle \coloneqq \frac{1}{\sqrt{N}} \cdot \sum_{x \in \{0,1\}^n} h(x) \cdot |x\rangle .$$

Next, a *function family* is a function $R : [K] \times \{0,1\}^n \to \{\pm 1\}$. We think of R as defining a family of K Boolean functions as follows: for each $1 \leq k \leq K$, we let $R_k : \{0,1\}^n \to \{\pm 1\}$ be the function $R_k(\cdot) \coloneqq R(k, \cdot)$. In general, we require $K \ll N$; a typical setting will be $K = N/2$.

Definition 1.3 (Oracle State Distinguishing Game)**.** Let $\boldsymbol{R} : [K] \times \{0,1\}^n \to \{\pm 1\}$ be a uniformly random function family. The *Oracle State Distinguishing Game* involves two parties, a challenger and an adversary. The adversary is modeled as an oracle circuit $A^{(\cdot)}$ which is allowed to query an arbitrary Boolean function f depending on $\boldsymbol{R}$. The game is played as follows.

1. The challenger samples a random bit $\boldsymbol{b} \in \{0,1\}$.
2. The challenger generates a random n-qubit state $|\boldsymbol{\psi}\rangle$ in one of two ways:
   - If $\boldsymbol{b} = 0$, the challenger samples a uniformly random $\boldsymbol{k} \sim [K]$ and generates $|\boldsymbol{\psi}\rangle \coloneqq |\psi_{R_{\boldsymbol{k}}}\rangle$, the binary phase state corresponding to the Boolean function $\boldsymbol{R_k}$.
   - If $\boldsymbol{b} = 1$, the challenger samples a uniformly random Boolean function $\boldsymbol{h} : \{0,1\}^n \to \{\pm 1\}$ and sets $|\boldsymbol{\psi}\rangle \coloneqq |\psi_{\boldsymbol{h}}\rangle$, the binary phase state corresponding to $\boldsymbol{h}$.
3. The challenger sends $|\boldsymbol{\psi}\rangle$ to the adversary.

4. The adversary runs the oracle circuit A^f on $|\boldsymbol{\psi}\rangle$ and outputs a bit $\boldsymbol{b}' \in \{0,1\}$.

5. If $\boldsymbol{b}' = \boldsymbol{b}$, then the adversary wins. Otherwise, they lose.

Intuitively, the function family $\boldsymbol{R}$ specifies a family of pseudorandom states $\{|\psi_{\boldsymbol{R}_k}\rangle\}_{k\in[K]}$, and the adversary's goal is to distinguish a randomly chosen state from this from a uniformly random binary phase state $|\psi_{\boldsymbol{h}}\rangle$. As discussed above, an algorithm for the Unitary Synthesis Problem yields a successful adversary for the Oracle State Distinguishing Game, and so a query lower bound for the Oracle State Distinguishing Game implies a query lower bound for the Unitary Synthesis Problem. We show the following lower bound for the Oracle State Distinguishing Game.

Theorem 1.4. *Suppose that A^f is a one-query oracle circuit that achieves advantage ε in the Oracle State Distinguishing Game. Then, A^f must make a query of size at least $\ell = \Omega(K\varepsilon^2)$ bits.*

This lower bound implies that for typical settings of K (such as $K = N/2$), to achieve a non-negligible distinguishing probability, the adversary's query must have length exponential in n; in particular, a superpolynomial-length query is required whenever $\varepsilon \geq n^{\omega(1)}/\sqrt{K}$. This dependence on K is optimal, as there are polynomial-time 1-query algorithms which do achieve distinguishing advantage $\Omega(1/\sqrt{K})$ (see Appendix C).

As discussed above, Theorem 1.4 immediately implies Theorem 1.1, our one-query lower bound for the Unitary Synthesis Problem. In fact, since we show that the adversary's distinguishing advantage is negligible, this gives a unitary $U_{\boldsymbol{R}}$ which is hard to synthesize even in an extremely weak sense: no efficient one-query algorithm A^f can correctly implement any unitary that even remotely approximates the behavior of $U_{\boldsymbol{R}}$. In addition, since the Oracle State Distinguishing Game is an oracle analogue of the security game for a single-copy PRS family, standard techniques (see Section 5) allow us to transform Theorem 1.4 into a proof that, relative to a random oracle, there exist PRS families and quantum bit commitment schemes secure against all one-query adversaries. This gives Theorem 1.2.

These results demonstrate the usefulness of the Oracle State Distinguishing Game as a means for studying the Unitary Synthesis Problem, and we believe that it is also a useful avenue for proving stronger lower bounds against algorithms which use more than one query. To this end, we make the following conjecture.

Conjecture 1.5 (Strong Non-Synthesis Conjecture)**.** *For all $K \geq n^{\omega(1)}$, any polynomial-query oracle algorithm A^f wins the Oracle State Distinguishing Game with advantage at most* $\mathsf{negl}(n)$.

A proof of Conjecture 1.5 would imply a negative resolution to the Unitary Synthesis Problem. In addition, it would imply the existence of single-copy PRSes (and thus, quantum bit commitments) secure against all efficient polynomial-query adversaries, relative to a random oracle. In other words, computationally secure quantum cryptography would not black-box imply the existence of any hard language. We note that the lower bound $K \geq n^{\omega(1)}$ in Conjecture 1.5 is necessary; as discussed above, if $K = \mathrm{poly}(n)$, then there is a simple attack that achieves $1/\sqrt{K} = 1/\mathrm{poly}(n)$ advantage.[2]

In Section 2.5, we state weaker conjectures which correspond to simpler cases of Conjecture 1.5. In particular, in Conjecture 2.7, we give a self-contained mathematical conjecture which corresponds to the simplest class of oracle adversaries that we do not know how to rule out.

[2] In fact, there is another attack that achieves advantage close to 1 in this regime, based on the LMR algorithm [LMR14, Yue22b]. The adversary can make a single call to its $\boldsymbol{R}$-dependent oracle f to generate $m = \mathrm{poly}(n)$ copies of each state $|\psi_{\boldsymbol{R}_k}\rangle$. Then for each $1 \leq k \leq K$, the adversary can test if the challenge state $|\boldsymbol{\psi}\rangle$ is equal to $|\psi_{\boldsymbol{R}_k}\rangle$ by measuring $|\boldsymbol{\psi}\rangle \otimes |\psi_{\boldsymbol{R}_k}\rangle^{\otimes m}$ with $\{\Pi_{\mathrm{sym}}, \mathsf{Id} - \Pi_{\mathrm{sym}}\}$, where Π_{sym} is the projector onto the symmetric subspace. If they are *not* equal, doing so will only perturb the state $|\boldsymbol{\psi}\rangle$ slightly, allowing the adversary to reuse it for further tests.

Additional remarks. We make two final observations about the Oracle State Distinguishing Game. First, note that the adversary's task is to perform a measurement $\{M_0, M_1\}$ which distinguishes between the two cases of the game. In particular, writing $U_{\boldsymbol{R}}$ for the unitary written above, the adversary would like to carry out the measurement specified by the two projectors

$$M_0 := U_{\boldsymbol{R}}^\dagger \cdot (|1\rangle\langle 1| + \cdots + |K\rangle\langle K|) \cdot U_{\boldsymbol{R}} \quad \text{and} \quad M_1 := \mathsf{Id} - M_0.$$

This is an example of a *measurement synthesis* task, an inherently quantum problem in which the input is quantum but the output is classical. Measurement synthesis has been discussed much less than state synthesis and unitary synthesis in the literature (the only work we are aware of that discusses it is [BEM+23]). However, our results suggest that it is measurement synthesis that is the hard problem at the core of unitary synthesis. Combined with the fact that state synthesis has efficient one-query algorithms [Ros23a], this suggests that the crucial distinction between classical problems and inherently quantum problems is whether the input, and not necessarily the output, is classical or quantum.

Second, we note that the Oracle State Distinguishing Game is fairly robust to the precise distribution of states used to specify it. For example, rather than specifying the game in terms of random binary phase states, we could have specified it using Haar random states. In this version of the game, K independent Haar random states $|\psi_1\rangle, \ldots, |\psi_K\rangle$ are sampled in advance. Then the adversary is given either ($b = 0$) one of these K states sampled uniformly at random, or ($b = 1$) a new Haar random state $|\psi\rangle$, and asked to distinguish between these two cases. Though we do not prove it here, our lower bound in Theorem 1.4 also holds for this variant of the Oracle State Distinguishing Game. One nice property of this distribution is that hardness of the Oracle State Distinguishing Game for this distribution directly implies hardness of Unitary Synthesis for a Haar-random unitary U. We refer the reader to Section 3.3 for further discussion.

1.3 Related Work

In this section, we elaborate on some works related to the Unitary Synthesis Problem and our results. We discuss (1) prior lower bounds, (2) positive results on the closely-related state synthesis problem, and (3) related work in unitary complexity theory.

1.3.1 Lower bounds for unitary synthesis

The best known prior lower bound for the Unitary Synthesis Problem comes from the original paper on this topic by Aaronson and Kuperberg [AK07]. To understand their lower bound, let us first make more explicit the computational model we are assuming for our oracle circuit $A^{(\cdot)}$. A general oracle circuit $A^{(\cdot)}$ may wish to make use of additional *ancilla qubits*, in which case it will be structured as follows: it will have an n-qubit input register and an input ancilla register initialized to $|0^a\rangle$, as well as an n-qubit output register and an a-qubit output "junk" register. Indeed, if $A^{(\cdot)}$ does *not* have ancillas, then it is unable to query any oracle f on inputs of length greater than n, which turns out to make $A^{(\cdot)}$ quite weak. This is because for such an $A^{(\cdot)}$, the number of possible unitaries you can synthesize when ranging over all functions f is bounded by 2^{2^n}, which is simply not enough to "cover all unitaries" by a counting argument. (See Appendix B for a simple lower bound along these lines.)

Now we can state the Aaronson and Kuperberg [AK07] lower bound. They showed a one-query lower bound against any oracle circuit $A^{(\cdot)}$ which has the following property: for every choice of oracle f, the oracle circuit A^f is required to *exactly implement* an n-qubit unitary on its first n

qubits. Mathematically, this means that for any n-qubit state $|\psi\rangle$, we must have that

$$A^f \cdot |\psi\rangle \otimes |0^a\rangle = (U_f \cdot |\psi\rangle) \otimes \left|\mathsf{junk}_f\right\rangle, \tag{1}$$

where U_f is some n-qubit unitary which depends on f, and $\left|\mathsf{junk}_f\right\rangle$ is some a-qubit junk state which depends on f. This defines a class of oracle algorithms that turns out to be highly restrictive, for several reasons. We list two.

1. The class excludes algorithms A^f such that Equation (1) only holds approximately, even with inverse-exponential precision.

2. The model requires that the circuit A^f implements a unitary for *every* oracle f. On the other hand, there are many examples of oracle circuits not belonging to this class which expect the oracle f to be "properly formatted", and do not synthesize any unitary if f is not properly formatted.

To elaborate on (2), consider the following simple attack: the oracle circuit A^f queries f to learn an ℓ-bit classical string s on an ancilla space, and then applies an n-qubit unitary U_s that depends on s. By the Bernstein-Vazirani trick, A^f can learn an ℓ-bit string s in a single query by first preparing the uniform superposition on ℓ qubits, then querying the Boolean function $f_s(x) \coloneqq s \cdot x$, and finally applying a Hadamard transform. Even though this oracle circuit A^f always implements a unitary on the first n qubits when f computes an inner-product function, this is not guaranteed in general: for arbitrary f, the oracle circuit may obtain a superposition over different s, in which case the operation on the first n-qubits is not guaranteed to be unitary.

Indeed, Aaronson and Kuperberg are able to prove their lower bound against this class by a *counting argument*: they prove that the number of distinct unitaries that a one-query oracle circuit $A^{(\cdot)}$ in this class can synthesize, ranging over all oracles f, is at most 4^{2^n} [AK07, Theorem 6.7], irrespective of the number of ancilla qubits a. Unfortunately, as we discuss in Section 2.1 and Appendix B, these types of counting arguments are insufficient to prove a general query lower bound.

A more recent lower bound, due to Rosenthal [Ros22], shows that unitary synthesis is hard relative to a *state synthesis* oracle. Roughly speaking, this lower bound states that synthesizing a unitary U requires roughly $2^{n/2}$ queries to an oracle that, on any classical input $|x\rangle$, for $x \in \{0,1\}^n$, outputs the state $|x\rangle \otimes U|x\rangle$. This shows that the power to produce any state of the form $U|x\rangle$ is insufficient to implement U efficiently. However, the technique says little about the problem of synthesizing U relative to an *arbitrary* function oracle f.

1.3.2 Relationship to state synthesis.

Let us contrast our one-query lower bound for the Unitary Synthesis Problem with the state of affairs for a related problem known as *state synthesis*. State synthesis is the task of implementing a quantum circuit that outputs a specified n-qubit quantum state $|\psi\rangle$ when run on the all-0's input. Alternatively, one can view state synthesis as an easier version of unitary synthesis, where the goal is merely to implement the unitary correctly on the all 0's input, rather than on *all* possible inputs.

Like unitary synthesis, state synthesis requires large quantum circuits: it can be shown via counting arguments that there exist worst-case states on n qubits that require circuits of size $\Omega(2^n/n)$ to compute approximately (see the excellent discussion of this in [Ros23b, Section 1.3.4]).

It turns out, however, that state synthesis becomes easy if Boolean functions are easy [Aar16, INN+22, Ros23a]. In particular, Rosenthal's state-of-the-art result [Ros23a] gives a quantum-polynomial time oracle algorithm $A^{(\cdot)}$ such that for any n-qubit pure state $|\psi\rangle$, there exists a

Boolean function $f : \{0,1\}^m \to \{\pm 1\}$ such that $A^f(1^n)$ makes *one* quantum query to f and outputs $|\psi\rangle$ up to inverse exponential precision. For some intuition behind this result, observe that *binary phase states* $\frac{1}{\sqrt{2^n}} \cdot \sum_{x\in\{0,1\}^n} f(x) \cdot |x\rangle$ are trivial to synthesize with one query: simply prepare the uniform superposition $\frac{1}{\sqrt{2^n}} \cdot \sum_{x\in\{0,1\}^n} |x\rangle$, and make one query to the phase oracle $\mathcal{O}_f : |x\rangle \to f(x) \cdot |x\rangle$. It turns out that worst-case states can then be synthesized via a careful reduction to the binary phase state case. This can be viewed as a one-query reduction from the task of state synthesis to the problem of computing an arbitrary Boolean function. In contrast, our main result shows that no such reduction is possible for unitary synthesis.

1.3.3 Quantum cryptography and unitary complexity

A connection between (plain model) quantum cryptography and the Unitary Synthesis Problem was recently discovered by Kretschmer [Kre23], who showed that if the Unitary Synthesis Problem is resolved in the positive, then showing the existence of a secure PRS implies that $\mathsf{BPP} \neq \mathsf{NEXP}$. This result says that traditional complexity theory *does* have something to say about the existence of PRSes, but only if unitaries are easy to synthesize.

Beyond "traditional complexity theory," a very recent and intriguing line of work has introduced a complexity theory of inherently quantum problems, with complexity classes corresponding to both state synthesis problems and unitary synthesis problems [RY22, INN+22, Ros23a, MY23, BEM+23, DGLM23]. As above, this line of work argues that traditional complexity theory is ill-equipped to address the complexity of inherently quantum problems, as traditional complexity theory is only about classical-input, classical-output problems, i.e., functions $f : \{0,1\}^* \to \{0,1\}$. In this new theory of unitary complexity, the existence of secure PRSes *does* have complexity theoretic implications (in particular, it implies the separation $\mathsf{unitaryBQP} \neq \mathsf{unitaryPSPACE}$).

An important open direction is to study the relationship between these new inherently quantum complexity theories and the traditional "classical" complexity theory. Interestingly, Kretschmer's result above [Kre23] suggests that these seemingly different complexity theories might be closer than they first appear, if the Unitary Synthesis Problem is resolved in the positive. In particular, his result, stated more broadly, is the following: suppose the Unitary Synthesis Problem is resolved in the positive. Then $\mathsf{unitaryBQP} \neq \mathsf{unitaryPSPACE}$ implies that $\mathsf{BPP} \neq \mathsf{NEXP}$. In this light, our Theorem 1.1, providing negative evidence for the Unitary Synthesis Problem, can also be interpreted as providing positive evidence that these complexity theories are in fact distinct.

1.4 Organization

The remainder of this paper is organized as follows. Section 2 gives a technical overview of our proofs. Section 3.1 includes preliminary details about oracle circuits, building towards a simple normal form for these circuits that we will use in our proofs. In Section 4, we give the proof of our main result, the one-query lower bound for the Oracle State Distinguishing Game, which we then use in Section 5 to show the existence of secure PRSes and quantum bit commitments relative to a random oracle. Appendix A includes a second proof of our main result with slightly worse parameters. In Appendix B, we show a counting lower bound against even many-query oracle circuits which can only compute a small number of distinct unitaries, generalizing the one-query lower bound of Aaronson and Kuperberg [AK07]. Finally, in Appendix C, we give a one-query algorithm to match our main lower bound (Theorem 4.18) in its dependence on K.

1.5 Acknowledgements

We thank Scott Aaronson, Prabhanjan Ananth, Zvika Brakerski, Ran Canetti, Lijie Chen, Shafi Goldwasser, Louis Golowich, Tarun Kathuria, William Kretschmer, Tony Metger, Sidhanth Mohanty, Anand Natarajan, Ryan O'Donnell, Luowen Qian, Prasad Raghavendra, Greg Rosenthal, Nick Spooner, Nikhil Srivastava, Vinod Vaikuntanathan, Ramon van Handel, Umesh Vazirani, Thomas Vidick, Henry Yuen, and Mark Zhandry for helpful discussions.

We thank Scott Aaronson and Greg Rosenthal for comments on an earlier draft of this paper.

2 Technical overview

We will sketch the proof of Theorem 1.4, beginning by describing our mathematical model for single-query adversaries in Section 2.1. Following this, we will develop our proof strategy in the context of three different and increasingly complicated types of adversaries. First, in Section 2.2, we will look at adversaries which use their one query to prepare a quantum advice state. Next, in Section 2.3, we will look at adversaries which have no ancilla qubits and do not apply any gates prior to their oracle query. Finally, in Section 2.4, we will look at general single-query adversaries.

2.1 Modeling the adversary

A single-query adversary can be modeled as a quantum circuit with an input register of n qubits and an ancilla register of a qubits, for a size of $m = n + a$ total qubits. Given an n-qubit input state $|\psi\rangle$, the adversary acts as follows.

1. The adversary will initialize its ancilla qubits to $|0^a\rangle$. Then, it applies a unitary U to $|\psi\rangle|0^a\rangle$. Equivalently, it applies the *isometry* $V \coloneqq U \cdot (\mathsf{Id} \otimes |0^a\rangle)$ to $|\psi\rangle$.

2. It then queries its oracle $f : \{0,1\}^m \to \{\pm 1\}$. This applies the unitary $\mathcal{O}_f$ to its state, where $\mathcal{O}_f$ is the unitary defined as

$$\mathcal{O}_f \cdot |z\rangle = f(z) \cdot |z\rangle\,, \quad \text{for all } z \in \{0,1\}^m.$$

3. Finally, the adversary performs a binary projective measurement $\{\Pi, \mathsf{Id} - \Pi\}$ on its state. This produces a measurement outcome $\boldsymbol{b'} \in \{0,1\}$, which it outputs as its guess.

After the oracle, the adversary's state is $\mathcal{O}_f \cdot V \cdot |\psi\rangle$. Thus, the probability it outputs $\boldsymbol{b'} = 0$ is

$$\|\Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi\rangle\|^2 = \langle\psi| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi\rangle\,. \tag{2}$$

Intuitively, one should think of the size m as "small", say $m = \mathrm{poly}(n)$. This is because m is also the length of the adversary's oracle query, and it is necessary for us to assume a bound on the query length so that the problem remains nontrivial. Otherwise, there is a simple attack based on the Bernstein-Vazirani algorithm [BV97] which solves the problem using a single extremely large query of length $K \cdot N$, which we describe below.

Example 2.1 (A one-query attack with exponential query size)**.** Define $f_R : \{0,1\}^{KN} \to \{\pm 1\}$ so that $f_R(z) = (-1)^{z \cdot r}$, where $r \in \{0,1\}^{KN}$ is a binary vector representation of R. Then if the adversary queries f_R on the uniform superposition over all $z \in \{0,1\}^{KN}$, it obtains the state

$$\frac{1}{\sqrt{KN}} \cdot \sum_{z \in \{0,1\}^{KN}} f_R(z) \cdot |z\rangle = \frac{1}{\sqrt{KN}} \cdot \sum_{z \in \{0,1\}^{KN}} (-1)^{z \cdot r} \cdot |z\rangle\,,$$

The state on the right-hand side is simply the Hadamard transform of $|r\rangle$, and thus the adversary can obtain the entire truth table of R.

As it turns out, once we assume our adversary has "small" query length, it can be converted to one with "small" size m as well (see Section 3.5). Hence, we may assume that the adversary's oracle is applied to all m qubits. We will now carry out the following change in notation that will be applied throughout the paper: to simplify notation, we will set $N \coloneqq 2^n$ and $M \coloneqq 2^m$ and

associate the set $\{1,\dots,N\}$ with $\{0,1\}^n$ and $\{1,\dots,M\}$ with $\{0,1\}^m$. As a result, a "Boolean" function is now formatted as $h : [N] \to \{\pm 1\}$ and is associated with the phase state

$$|\psi_h\rangle = \frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} h(x) \cdot |x\rangle \,,$$

a function family is formatted $R : [K] \times [N] \to \{\pm 1\}$, and an oracle function is formatted $f : [M] \to \{\pm 1\}$. With this change in notation, the input state space becomes $\mathbb{C}^N$ and the adversary's state space becomes $\mathbb{C}^M$, so that (i) the isometry V maps $\mathbb{C}^N$ to $\mathbb{C}^M$, and (ii) the oracle unitary $\mathcal{O}_f$ has dimension $M \times M$. Thus, for any particular function family $R : [K] \times [N] \to \{\pm 1\}$, and any one-query adversary $A^{(\cdot)}$, the maximum achievable distinguishing advantage is equal to

$$\max_{f:[M]\to\{\pm1\}} \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \Big[\mathbf{Pr}[A^f(|\psi_{R_{\boldsymbol{k}}}\rangle) \text{ outputs "0"}] \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[\mathbf{Pr}[A^f(|\psi_{\boldsymbol{h}}\rangle) \text{ outputs "0"}] \Big] \right|,$$

where here and throughout this section we are writing $\boldsymbol{h}$ for a uniformly random Boolean function $\boldsymbol{h} : [N] \to \{\pm 1\}$. Substituting in Equation (2), this is equal to

$$\max_{f} \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \Big[\langle\psi_{R_{\boldsymbol{k}}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{R_{\boldsymbol{k}}}\rangle \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[\langle\psi_{\boldsymbol{h}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle \Big] \right|. \tag{3}$$

Our goal is to prove Theorem 1.4, which can be phrased more formally as follows.

Theorem 2.2 (Theorem 1.4, rephrased)**.** *Let $A^{(\cdot)}$ be a single-query adversary for the Oracle State Distinguishing Game that acts on an M-dimensional Hilbert space. Then with high probability over the choice of $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$, $A^{(\cdot)}$ achieves maximum distinguishing advantage at most $O\Big(\sqrt{\frac{\log M}{K}}\Big)$.*

Now let us now briefly discuss one potential approach for proving Theorem 1.4: counting arguments. These are based on the simple observation is that the distinguishing advantage is easily upper-bounded for any *fixed* oracle f, which corresponds to an adversary that does not depend on $\boldsymbol{R}$. This can be argued using standard concentration of measure tools from probability theory, and the resulting concentration bound one can show is *extremely good*: in particular, the probability that a fixed A^f has distinguishing advantage at least ε is at most $2^{-\Omega(\varepsilon^2 KN)}$. Given this degree of concentration, it is tempting to simply union bound over all choices of f to upper-bound the maximum distinguishing advantage; this is known as a counting argument. Unfortunately, this approach quickly begins to fail as the adversary's space grows: the number of possible functions f is 2^M, where M is potentially much larger than KN. Recall that $KN \leq N^2 = 2^{2n}$, while M could be (at least) $2^{\mathrm{poly}(n)}$, for an arbitrary $\mathrm{poly}(n)$. Thus, this type of counting argument cannot give a general one-query lower bound. That said, it *can* rule out some interesting special cases of adversaries, which we discuss in Appendix B. Finally, we note that there is a more powerful version of counting arguments known as *chaining* (cf. [Ver18, Chapter 8]), but we were unable to successfully apply chaining arguments to this problem.

In the next few subsections, we will describe an alternative approach for bounding the maximum distinguishing advantage across all choices of f simultaneously via matrix concentration inequalities.

2.2 Adversaries with quantum advice

We begin with the simple but conceptually useful special case of one-query adversaries, namely those that use the query to f to synthesize an f-dependent *advice state*. In other words, the adversary acts as follows.

1. First, it applies an isometry V that acts by appending a fixed m-qubit state $|\phi\rangle$. Thus, the n-qubit input state $|\psi\rangle$ is mapped to the $(n+m)$-qubit state $|\psi\rangle \otimes |\phi\rangle$. (We are abusing notation in this subsection by writing m only for the qubits in the advice state, rather than for all of the qubits. We will return to the normal definition of m in Sections 2.3 and 2.4 below.)

2. Next, it makes an oracle query $\mathcal{O}_f$ that acts as the identity on the input state $|\psi\rangle$ and only modifies $|\phi\rangle$. Then the adversary's state becomes $|\psi\rangle \otimes |\phi_f\rangle$, where $|\phi_f\rangle$ is some f-dependent state.

In total, for such an adversary, $\mathcal{O}_f \cdot V \cdot |\psi\rangle = |\psi\rangle \otimes |\phi_f\rangle$. Attacks of this form can synthesize many kinds of states: for example, if $|\phi\rangle$ is a uniform superposition, then $|\phi_f\rangle$ can be any binary phase state. (We remark that there are techniques in the cryptography literature for proving lower bounds against quantum advice [HXY19, CLQ20, CGLQ20, Liu23]. However, the techniques seem to be highly tailored to the advice setting and are not related to our approach.)

Supposing the adversary works in this manner, we can compute its maximum distinguishing advantage on a uniformly random $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ as

$$\max_{f:[M]\to\{\pm1\}} \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[\langle\psi_{\boldsymbol{R_k}}| \langle\phi_f| \cdot \Pi \cdot |\psi_{\boldsymbol{R_k}}\rangle |\phi_f\rangle \right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\langle\psi_{\boldsymbol{h}}| \langle\phi_f| \cdot \Pi \cdot |\psi_{\boldsymbol{h}}\rangle |\phi_f\rangle \right] \right|,$$

by Equation (3). The benefit of focusing on advice states is that we can factor out the f-dependent term $|\phi_f\rangle$ from each expectation. To do so, for any Boolean function $h : [N] \to \{\pm 1\}$, let Π_h denote the $M \times M$-dimensional matrix

$$\Pi_h := (\langle\psi_h| \otimes \mathsf{Id}) \cdot \Pi \cdot (|\psi_h\rangle \otimes \mathsf{Id}).$$

Note that $0 \le \Pi_h \le \mathsf{Id}$, since $|\psi_h\rangle$ is a unit vector and Π is a projection. Then we can rewrite the distinguishing advantage as

$$\begin{aligned} &\max_{f:[M]\to\{\pm1\}} \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[\langle\phi_f| \cdot \Pi_{\boldsymbol{R_k}} \cdot |\phi_f\rangle \right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\langle\phi_f| \cdot \Pi_{\boldsymbol{h}} \cdot |\phi_f\rangle \right] \right| \\ = &\max_{f:[M]\to\{\pm1\}} \left| \langle\phi_f| \cdot \left(\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[\Pi_{\boldsymbol{R_k}}\right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}}\right] \right) \cdot |\phi_f\rangle \right|. \end{aligned} \tag{4}$$

Since $|\phi_f\rangle$ is a unit vector, we can upper bound this by a maximum over all unit vectors, i.e.

$$\begin{aligned} (4) &\le \max_{\||v\rangle\|=1} \left| \langle v| \cdot \left(\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[\Pi_{\boldsymbol{R_k}}\right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}}\right] \right) \cdot |v\rangle \right| \\ &= \left\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[\Pi_{\boldsymbol{R_k}}\right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}}\right] \right\|_{\mathrm{op}} \\ &= \left\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[Z_{\boldsymbol{R_k}}\right] \right\|_{\mathrm{op}}, \quad \text{for} \quad Z_h := \Pi_h - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}}\right]. \end{aligned}$$

Here, we are writing $\|\cdot\|_{\mathrm{op}}$ for the operator norm. Thus, we have reduced our problem to bounding the operator norm of the average of K random matrices $Z_{\boldsymbol{R}_1}, \ldots, Z_{\boldsymbol{R}_K}$.

We will bound this operator norm using the technique of *matrix concentration*, which generalizes scalar concentration bounds (such as Chernoff-Hoeffding bounds) to the random matrix setting. Specifically, the matrix Hoeffding inequality (roughly) says the following (see [Tro12], Theorem 1.3 or Theorem A.18 for the precise statement).

Theorem 2.3 (Matrix Hoeffding (informal))**.** *If K independent and identically distributed mean-zero random $D \times D$ Hermitian matrices $\boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_K$ always have bounded operator norm, then with high probability,*

$$\Big\| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \boldsymbol{Z}_{\boldsymbol{k}} \Big\|_{\text{op}} \leq O\Big(\sqrt{\frac{\log(D \cdot K)}{K}}\Big).$$

(Note that the scalar Hoeffding bound can be recovered by taking $D = 1$ above.) To apply the matrix Hoeffding inequality to our problem, we need to verify that when $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is uniformly random, our matrices $Z_{\boldsymbol{R}_1}, \ldots, Z_{\boldsymbol{R}_K}$ satisfy these properties. Indeed:

- $Z_{\boldsymbol{R}_1}, \ldots, Z_{\boldsymbol{R}_K}$ are independent and identically distributed since each $\boldsymbol{R}_k$ is an independent, uniformly random Boolean function $\boldsymbol{R}_k : [N] \to \{\pm 1\}$.
- For each $1 \leq k \leq K$, $Z_{\boldsymbol{R}_k}$ has expectation zero:

$$\mathop{\mathbf{E}}_{\boldsymbol{R}} \left[Z_{\boldsymbol{R}_k} \right] = \mathop{\mathbf{E}}_{\boldsymbol{R}} \left[\Pi_{\boldsymbol{R}_k} \right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}} \right] = 0.$$

- For each $1 \leq k \leq K$, the operator norm $\| Z_{\boldsymbol{R}_k} \|_{\text{op}}$ is always bounded by 2, since

$$\| Z_{\boldsymbol{R}_k} \|_{\text{op}} = \Big\| \Pi_{\boldsymbol{R}_k} - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}} \right] \Big\|_{\text{op}} \leq \| \Pi_{\boldsymbol{R}_k} \|_{\text{op}} + \Big\| \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\Pi_{\boldsymbol{h}} \right] \Big\|_{\text{op}} \leq 2.$$

As a result, since in our setting $D = M$, an ε-distinguisher requires $\log(M) = \Omega(K\varepsilon^2)$, as claimed. In other words, the adversary needs a *huge* advice state to win the distinguishing game.

In summary, our strategy involved identifying a well-behaved quantity that governs the advantage of A^f across all choices of f simultaneously. As we have seen, the operator norm is an example of such a quantity: although bounding the quadratic form $\langle v| \cdot (\mathbf{E}_{\boldsymbol{k} \sim [K]} Z_{\boldsymbol{R}_{\boldsymbol{k}}}) \cdot |v\rangle$ for *all* vectors $|v\rangle$ would naively require the concentration of $O(1/\varepsilon)^M$ different scalars (corresponding to an ε-net over $\mathbb{C}^M$), matrix concentration shows that the operator norm behaves as if it has M, rather than 2^M, "independent degrees of freedom".

2.3 Adversaries with a trivial isometry

Let us recall Equation (3), our expression for the adversary's maximum distinguishing advantage:

$$\max_{f:[M] \to \{\pm 1\}} \Big| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big[\langle \psi_{R_{\boldsymbol{k}}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{R_{\boldsymbol{k}}}\rangle \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[\langle \psi_{\boldsymbol{h}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle \Big] \Big|.$$

The advice state case above suggests the following approach to bounding this expression:

1. Factor the dependence on the oracle $\mathcal{O}_f$ to the "outside" of the expression, and
2. Rely on a matrix concentration inequality to bound the advantage for all f simultaneously.

Unfortunately, the advice state case does not tell us whether this approach is possible, or how to carry it out, in general. To gain some intuition, we will analyze another simple special case, the case where $V = \mathsf{Id}$, in which the adversary does not use any ancilla qubits and only applies the identity unitary. In this case, $M = N$, and we will allow the adversary to query an arbitrary oracle

$f : [N] \to \{\pm 1\}$. Then the because $V = \mathsf{Id}$, the adversary's maximum distinguishing advantage on a uniformly random $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is given by

$$\max_{f:\{0,1\}^n \to \{\pm 1\}} \Big| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big[\langle \psi_{\boldsymbol{R_k}} | \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot |\psi_{\boldsymbol{R_k}}\rangle \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[\langle \psi_{\boldsymbol{h}} | \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot |\psi_{\boldsymbol{h}}\rangle \Big] \Big|. \tag{5}$$

Towards "factoring out" the $\mathcal{O}_f$ dependence to the outside of the expression, we make use of the fact that any binary phase state $|\psi_h\rangle$ can be written as the product of a diagonal $\{\pm 1\}$-matrix and the uniform superposition state:

$$|\psi_h\rangle = \frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} h(x) \cdot |x\rangle = \underbrace{\left(\sum_{x=1}^{N} h(x) \cdot |x\rangle\langle x| \right)}_{D_h} \cdot \underbrace{\left(\frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} |x\rangle \right)}_{|+_N\rangle} = D_h \cdot |+_N\rangle \,.$$

The key benefit of this "diagonal decomposition" is that the diagonal matrices D_h and $\mathcal{O}_f$ *commute*, which allows us to rewrite the state $\mathcal{O}_f \cdot |\psi_h\rangle$ as follows:

$$\mathcal{O}_f \cdot |\psi_h\rangle = \mathcal{O}_f \cdot D_h \cdot |+_N\rangle = D_h \cdot \mathcal{O}_f \cdot |+_N\rangle = D_h \cdot |\phi_f\rangle \,,$$

where $|\phi_f\rangle \coloneqq \mathcal{O}_f \cdot |+_N\rangle$ is the binary phase state corresponding to f. Plugging this back into our expression for the maximum distinguishing advantage, we can again employ a spectral relaxation:

$$\begin{aligned}
(5) &= \max_{f:\{0,1\}^n \to \{\pm 1\}} \Big| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big[\langle \phi_f | \cdot D_{\boldsymbol{R_k}} \cdot \Pi \cdot D_{\boldsymbol{R_k}} \cdot |\phi_f\rangle \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[\langle \phi_f | \cdot D_{\boldsymbol{h}} \cdot \Pi \cdot D_{\boldsymbol{h}} \cdot |\phi_f\rangle \Big] \Big| \\
&= \max_{f:\{0,1\}^n \to \{\pm 1\}} \Big| \langle \phi_f | \cdot \Big(\mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big[D_{\boldsymbol{R_k}} \cdot \Pi \cdot D_{\boldsymbol{R_k}} \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[D_{\boldsymbol{h}} \cdot \Pi \cdot D_{\boldsymbol{h}} \Big] \Big) \cdot |\phi_f\rangle \Big| \\
&\leq \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big[D_{\boldsymbol{R_k}} \cdot \Pi \cdot D_{\boldsymbol{R_k}} \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[D_{\boldsymbol{h}} \cdot \Pi \cdot D_{\boldsymbol{h}} \Big] \Big\|_{\mathrm{op}} \\
&= \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} Z_{\boldsymbol{R_k}} \Big\|_{\mathrm{op}}, \quad \text{for} \quad Z_{\boldsymbol{R}_k} \coloneqq D_{\boldsymbol{R_k}} \cdot \Pi \cdot D_{\boldsymbol{R_k}} - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[D_{\boldsymbol{h}} \cdot \Pi \cdot D_{\boldsymbol{h}} \Big].
\end{aligned}$$

As in the advice state case, our problem has again reduced to bounding the operator norm of $\mathbf{E}_{\boldsymbol{k} \sim [K]} Z_{\boldsymbol{R_k}}$ for a uniformly random $\boldsymbol{R}$. And just like before, the matrices $Z_{\boldsymbol{R}_1}, \ldots, Z_{\boldsymbol{R}_K}$ are mean-zero, independent and identically distributed, and their norm is bounded by 2, since

$$\|Z_{\boldsymbol{R}_k}\|_{\mathrm{op}} \leq \|D_{\boldsymbol{R_k}} \cdot \Pi \cdot D_{\boldsymbol{R_k}}\|_{\mathrm{op}} + \Big\| \mathop{\mathbf{E}}_{\boldsymbol{h}} \big[D_{\boldsymbol{h}} \cdot \Pi \cdot D_{\boldsymbol{h}} \big] \Big\|_{\mathrm{op}} \leq 2,$$

where the second inequality uses the fact that the D_h is a unitary matrix for any Boolean function $h : [N] \to \{\pm 1\}$, so $\|D_h \cdot M \cdot D_h\|_{\mathrm{op}} \leq 1$. Thus, we can apply the matrix Hoeffding inequality as before. Since the $Z_{\boldsymbol{R}_1}, \ldots, Z_{\boldsymbol{R}_K}$ are $N \times N$ matrices, matrices, we obtain a bound on the maximum distinguishing advantage of

$$O\left(\sqrt{\frac{\log(N)}{K}} \right).$$

To summarize, the key new idea in this special case was to introduce a diagonal decomposition which holds for arbitrary phase states $|\psi_h\rangle$.

2.4 The general one-query bound

Now we consider the case of a general adversary. Let us recall one last time Equation (3), our expression for the adversary's maximum distinguishing advantage:

$$\max_{f:[M]\to\{\pm 1\}} \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[\langle \psi_{R_{\boldsymbol{k}}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{R_{\boldsymbol{k}}}\rangle \right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[\langle \psi_{\boldsymbol{h}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle \right] \right|.$$

The previous special case suggests the following strategy for bounding this expression:

1. First, for any Boolean function $h : [N] \to \{\pm 1\}$, find a "diagonal decomposition" of the state $V \cdot |\psi_h\rangle$ of the form

$$V \cdot |\psi_h\rangle = \big(h\text{-dependent diagonal matrix}\big) \cdot |\text{fixed state}\rangle\,,$$

2. Next, use this decomposition to re-express $\mathcal{O}_f \cdot V \cdot |\psi_h\rangle$ as

$$\begin{aligned} \mathcal{O}_f \cdot V \cdot |\psi_h\rangle &= \mathcal{O}_f \cdot \big(h\text{-dependent diagonal matrix}\big) \cdot |\text{fixed state}\rangle \\ &= \big(h\text{-dependent diagonal matrix}\big) \cdot \mathcal{O}_f \cdot |\text{fixed state}\rangle \\ &= \big(h\text{-dependent diagonal matrix}\big) \cdot |f\text{-dependent state}\rangle\,. \end{aligned}$$

Importantly, the diagonal decomposition should satisfy the following two properties:

1. the fixed state should have *unit norm*, so that we can perform a spectral relaxation, and
2. the h-dependent diagonal matrix should have *bounded operator norm*, so that we can apply the matrix Hoeffding inequality.

Unfortunately, it turns out that for a general isometry V, a diagonal decomposition satisfying the above requirements does not exist. Consider the following example.

Example 2.4 (No nice diagonal decomposition)**.** Let $M = N$ and let V be the $N \times N$ Hadamard transform $V = H^{\otimes n}$. For all Boolean vectors $r \in \{0,1\}^n$, let $h_r : \{0,1\}^n \to \{0,1\}$ denote the inner product function $h_r(x) = r \cdot x \pmod 2$. Then

$$|\psi_{h_r}\rangle = \frac{1}{\sqrt{N}} \cdot \sum_{x\in\{0,1\}^n} (-1)^{r\cdot x} \cdot |x\rangle = H^{\otimes n} \cdot |r\rangle\,,$$

and thus $V \cdot |\psi_{h_r}\rangle = H^{\otimes n} \cdot |\psi_{h_r}\rangle = |r\rangle$. Suppose that we try to write each $|\psi_{h_r}\rangle$ as the product of an h_r-dependent diagonal matrix and the uniform superposition state $|+_N\rangle := (1/\sqrt{N}) \cdot \sum_{x\in\{0,1\}^n} |x\rangle$. That is, for each $r \in \{0,1\}^n$:

$$V \cdot |\psi_{h_r}\rangle = |r\rangle = D_r \cdot \Big(\frac{1}{\sqrt{N}} \cdot \sum_{x\in\{0,1\}^n} |x\rangle \Big).$$

Then the only choice of D_r satisfying the above is $D_r = \sqrt{N} \cdot |r\rangle\langle r|$, which has *exponentially large* operator norm. Moreover, there is nothing special about the uniform superposition — no matter what fixed state we use, there will exist h_r such that D_r has exponentially large operator norm.

Thus, Example 2.4 shows that we cannot hope for a diagonal decomposition that satisfies our desired conditions for *all* binary phase states $|\psi_h\rangle$. Nevertheless, we will show that a meaningful diagonal decomposition is still possible in the general case. The key insight, which we will show next, is that for any isometry V, there exists a diagonal decomposition of $V \cdot |\psi_{\boldsymbol{h}}\rangle$ in which the $\boldsymbol{h}$-dependent diagonal matrix has bounded operator norm with extremely high probability over the choice of $\boldsymbol{h} : [N] \to \{\pm 1\}$.

2.4.1 The weight vector decomposition

Our goal is to find a diagonal decomposition of the form

$$V \cdot |\psi_{\boldsymbol{h}}\rangle = D_{V,\boldsymbol{h}} \cdot |\phi\rangle\,,$$

in which $D_{V,\boldsymbol{h}}$ has a "small" operator norm, with high probability over a uniformly random $\boldsymbol{h}$. Let us consider what would be implied if such a decomposition were to exist, and then work backwards to construct the decomposition.

If such a decomposition exists, then for each $1 \le i \le M$, let us consider the i-th coordinates of the left-hand and right-hand sides, which are given by

$$\langle i| \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle = \langle i| \cdot D_{V,\boldsymbol{h}} \cdot |\phi\rangle = (D_{V,\boldsymbol{h}})_{i,i} \cdot \phi_i. \tag{6}$$

Because $D_{V,\boldsymbol{h}}$ has "small" operator norm for a "typical" $\boldsymbol{h}$, this means that $(D_{V,\boldsymbol{h}})_{i,i}$ is "small" for a "typical" $\boldsymbol{h}$. Hence, for such an $\boldsymbol{h}$, the right-hand side of Equation (6) must be roughly equal to ϕ_i, the magnitude of the i-th coordinate in $|\phi\rangle$. (More correctly, it must be not too much *larger* than ϕ_i.) This, in turn, implies that the left-hand side of Equation (6) must be roughly equal to ϕ_i as well, at least for a "typical" $\boldsymbol{h}$. This motivates studying the magnitude of the i-th coordinate of $V \cdot |\psi_{\boldsymbol{h}}\rangle$ for a "typical" Boolean function $\boldsymbol{h}$. We can do so by looking at its average squared magnitude

$$p_i \coloneqq \mathop{\mathbf{E}}_{\boldsymbol{h}}[|\langle i| \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle|^2]. \tag{7}$$

In other words, p_i denotes the probability that measuring the state $V|\psi_{\boldsymbol{h}}\rangle$ in the standard basis results in an outcome of i. Then we expect the i-th coordinate of $V \cdot |\psi_{\boldsymbol{h}}\rangle$ to have magnitude roughly $\sqrt{p_i}$, and that suggests the following choice for our fixed state in the diagonal decomposition:

$$|\mathsf{wt}_V\rangle \coloneqq \sum_{i=1}^{M} \sqrt{p_i} \cdot |i\rangle\,,$$

which we refer to as the *weight vector* for V. We observe that $|\mathsf{wt}_V\rangle$ is indeed a unit vector because

$$\langle \mathsf{wt}_V|\mathsf{wt}_V\rangle = \sum_{i=1}^{M} p_i = \mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[\sum_{i=1}^{M} |\langle i| \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle|^2\Big] = \mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[\|V \cdot |\psi_{\boldsymbol{h}}\rangle\|^2\Big] = 1,$$

where the last equality holds because V is an isometry. Intuitively, the state $|\mathsf{wt}_V\rangle$ encodes how much *weight* the isometry V places on each individual coordinate $1 \le i \le M$.

To compute the full diagonal decomposition, we write the isometry $V : \mathbb{C}^N \to \mathbb{C}^M$ explicitly as

$$V = \sum_{i=1}^{M}\sum_{x=1}^{N} v_{i,x} \cdot |i\rangle\langle x| = \sum_{i=1}^{M} |i\rangle \cdot \Big(\sum_{x=1}^{N} v_{i,x}\,\langle x|\Big) = \sum_{i=1}^{M} |i\rangle\langle v_i|\,,$$

so that the ith coordinate of $V \cdot |\psi_{\boldsymbol{h}}\rangle$ is given by

$$\alpha_{\boldsymbol{h},i} \coloneqq \langle i| \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle = \langle i| \cdot \Big(\sum_{i=1}^{M} |i\rangle\langle v_i|\Big) \cdot |\psi_{\boldsymbol{h}}\rangle = \langle v_i|\psi_{\boldsymbol{h}}\rangle\,,$$

yielding the decomposition

$$V \cdot |\psi_{\boldsymbol{h}}\rangle = \sum_{i=1}^{M} \alpha_{\boldsymbol{h},i} \cdot |i\rangle = \underbrace{\Big(\sum_{i=1}^{M} \frac{\alpha_{\boldsymbol{h},i}}{\sqrt{p_i}} \cdot |i\rangle\langle i|\Big)}_{D_{V,\boldsymbol{h}}} \cdot \Big(\sum_{i=1}^{M} \sqrt{p_i} \cdot |i\rangle\Big) = D_{V,\boldsymbol{h}} \cdot |\mathsf{wt}_V\rangle\,.$$

2.4.2 Bounding the operator norm of $D_{V,\boldsymbol{h}}$.

Our next step is to determine whether the random matrix $D_{V,\boldsymbol{h}}$ actually *has* bounded operator norm with high probability. Its operator norm is given by

$$\|D_{V,\boldsymbol{h}}\|_{\mathrm{op}} = \max_{1\le i\le M} \frac{|\alpha_{\boldsymbol{h},i}|}{\sqrt{p_i}},$$

and we know from Equation (7) that for every $1 \le i \le M$,

$$\mathop{\mathbf{E}}_{\boldsymbol{h}}[|\alpha_{\boldsymbol{h},i}|^2] = p_i.$$

Therefore, if each coordinate $\alpha_{\boldsymbol{h},i}$ has good enough (scalar) concentration, we can bound $\|D_{\boldsymbol{h}}\|_{\mathrm{op}}$ with high probability.

Thus, we have reduced the problem to understanding the concentration of the random variables

$$\frac{\alpha_{\boldsymbol{h},i}}{\sqrt{p_i}} = \frac{1}{\sqrt{p_i}} \cdot \langle v_i|\psi_{\boldsymbol{h}}\rangle = \sum_{x=1}^{N} \frac{v_{i,x}}{\sqrt{p_i}} \cdot \boldsymbol{h}(x).$$

To see that this expression is small with high probability, we observe that it is a weighted linear combination of the N i.i.d. $\{\pm 1\}$ random variables $\{\boldsymbol{h}(x)\}_{x\in[N]}$, has mean zero, and has variance

$$\mathop{\mathbf{E}}_{\boldsymbol{h}}\left[\left|\frac{\alpha_{\boldsymbol{h},i}}{\sqrt{p_i}}\right|^2\right] = \frac{\mathbf{E}_{\boldsymbol{h}}[|\alpha_{\boldsymbol{h},i}|^2]}{p_i} = \frac{p_i}{p_i} = 1.$$

Therefore, standard (scalar) concentration tools (see Theorem 4.21) tell us that this random variable exhibits "sub-Gaussian concentration," implying (in this case) that it is larger than any t with probability at most $2 \cdot \exp(-t^2/2)$. Union bounding over all M coordinates, we conclude that $\|D_{V,\boldsymbol{h}}\|_{\mathrm{op}} > t$ with probability at most $2M \cdot \exp(-t^2/2)$, and so it is, for example, unlikely to be much larger than $O(\sqrt{\log M})$.

2.4.3 Putting everything together

With the weight-vector decomposition in hand, we can proceed to bounding the adversary's maximum distinguishing advantage along similar lines as in Section 2.3. To begin, we can rewrite the state $\mathcal{O}_f \cdot V \cdot |\psi_h\rangle$ as follows:

$$\mathcal{O}_f \cdot V \cdot |\psi_h\rangle = \mathcal{O}_f \cdot D_{V,h} \cdot |\mathsf{wt}_V\rangle = D_{V,h} \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle\,,$$

for every function h. In other words, the choice of $\boldsymbol{R}$-dependent function f here is accounted for as the vector $\mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$, which is a unit vector for all functions f.

By an argument similar to the one in Section 2.3, we can then bound the adversary's maximum distinguishing advantage by the operator norm

$$\left\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} Z_{\boldsymbol{R}_{\boldsymbol{k}}} \right\|_{\mathrm{op}}, \quad \text{where } Z_h \coloneqq D_{V,h}^{\dagger} \cdot \Pi \cdot D_{V,h} - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[D_{V,\boldsymbol{h}}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{h}} \right].$$

Thus, we have again reduced our problem to bounding the operator norm of an average of K independent and identically distributed matrices $Z_{\boldsymbol{R}_k}$ whose operator norms are bounded with high probability. In particular, since the operator norm of $\|D_{V,\boldsymbol{h}}\|$ is usually no more than $O(\sqrt{\log M})$,

over a uniformly random $\boldsymbol{h}$, the operator norm of $D^{\dagger}_{V,\boldsymbol{h}} \cdot \Pi \cdot D_{V,\boldsymbol{h}}$ is usually no more than $O(\log M)$. We would like to then conclude that $\left\|\mathbf{E}_{\boldsymbol{k}\sim[K]} Z_{\boldsymbol{R_k}}\right\|_{\mathrm{op}}$ is at most

$$O\Big(\sqrt{\frac{\log M}{K}}\Big)$$

with high probability, which would imply our claimed result.

Unfortunately, we cannot quite apply the matrix Hoeffding inequality directly, which requires that the matrices have bounded operator norm with probability 1, not just with high probability. Getting around this issue requires some additional technical ideas, and we give two ways of handling it in the main body of the paper.

1. Our first approach is to truncate the diagonal matrices $D_{V,\boldsymbol{R}_k}$ so that any entries whose magnitude exceeds some number B are scaled down so that their magnitude is equal to B. The result is that all matrices now have bounded operator norm, which means we are in fact able to apply the matrix Hoeffding inequality. Ultimately, this results in a bound of $O(1/K^{1/4})$ on the adversary's distinguishing advantage for reasonably small values of M (say, $M \leq \exp\big(K^{1/8}/4\big)$), which is more than enough to prove the one-query lower bound for the Unitary Synthesis Problem. That said, this bound is not quite strong enough to prove the bound claimed in Theorem 1.4.

2. To prove the precise bound claimed in Theorem 1.4 and thereby achieve the correct asymptotic dependence on K, we give a somewhat different analysis.

   (a) First, we show that it suffices to bound the *expected* distinguishing advantage on a random $\boldsymbol{R}$, rather than proving a bound with high probability. To show this, we show that the maximum distinguishing advantage concentrates extremely well around its expectation (see Lemma 3.18).

   (b) To bound the expected distinguishing advantage, we use a different technique called "decoupling", which is common in the random matrix theory literature [Ver11, vH17]. At a high level, the technique (when combined with the ideas from this technical overview) allows us to reduce to bounding the expected operator norm of the random matrix

   $$\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[D^{\dagger}_{V,\boldsymbol{R_k}} \cdot \Pi \cdot D_{V,\boldsymbol{R}'_{\boldsymbol{k}}}],$$

   where $\boldsymbol{R}$ and $\boldsymbol{R}'$ are independent and uniformly random function families. This is easier to give a sharp bound on because the dependence on each of $\boldsymbol{R}$ and $\boldsymbol{R}'$ is *linear* rather than quadratic, allowing us to prove an optimal bound on the expected value by applying a different matrix concentration inequality for matrix Rademacher series (Theorem 4.10). Unlike the matrix Hoeffding inequality, this matrix concentration inequality does not require the matrices to have bounded operator norm with probability 1 (though it does require the matrices to be random Rademacher matrices), which is why it gives stronger bounds than we achieve using our first approach.

The second proof, which achieves the optimal dependence on K, is presented in Section 4. The first proof is presented in Appendix A. We believe that theoretical computer scientists might find the first proof more straightforward to follow.

2.5 Future directions: beyond one query

Theorem 1.4 proves that efficient one-query oracle algorithms achieve at most negligible advantage in the Oracle State Distinguishing Game (and thus cannot synthesize arbitrary unitaries). We conjecture (see Conjecture 1.5) that efficient oracle circuits making $\mathrm{poly}(n)$-many sequential queries cannot win our distinguishing game. Towards resolving the full conjecture, we believe it may be useful to focus on the special case of *two-query* adversaries. In this subsection, we present several conjectures — all weaker than Conjecture 1.5 — that capture the simplest unresolved special cases of two-query attacks.

First, we will need the following observation about the power of *classical* oracle queries. Let us fix a function family $R : [K] \times [N] \to \{\pm 1\}$ and a projective measurement $\mathcal{P} = \{\Pi_i\}_{i\in[M]}$ with M outcomes. Suppose the adversary, upon receiving $|\boldsymbol{\psi}\rangle$, applies an isometry $V : \mathbb{C}^N \to \mathbb{C}^M$ followed by an oracle query $\mathcal{O}_f$. Next, it performs the measurement $\mathcal{P}$, obtaining some outcome in $\{1, \ldots, D\}$. Depending on whether the adversary's input state is sampled from the "pseudorandom" or "random" distribution, the outcome of measuring $\mathcal{P}$ is distributed as either:

- Dist_0: The result of applying $\mathcal{P}$ to $\mathcal{O}_f \cdot V \cdot |\psi_{R_{\boldsymbol{k}}}\rangle$, for a random $\boldsymbol{k} \sim [K]$.
- Dist_1: The result of applying $\mathcal{P}$ to $\mathcal{O}_f \cdot V \cdot |\psi_{\boldsymbol{h}}\rangle$, for a random $\boldsymbol{h} : [N] \to \{\pm 1\}$.

We observe that if the total variation distance (or statistical distance) between Dist_0 and Dist_1 is ε, then, by making one classical oracle query, the adversary can distinguish the two cases with advantage ε. This is because the second query can be made to the Boolean function $g : [M] \to \{\pm 1\}$, defined as $g(i) = \mathsf{Sign}(\mathbf{Pr}[i \sim \mathsf{Dist}_0] - \mathbf{Pr}[i \sim \mathsf{Dist}_1])$. If the output of g is $+1$, the adversary guesses that it was in the pseudorandom case, and if the output of g is -1, the adversary guesses that it was in the Haar random case, and attains distinguishing advantage ε.

The "1.5-query" conjecture. We conjecture that adversaries that make one quantum query followed by one classical query cannot win our distinguishing game.

Conjecture 2.5 (The 1.5-query conjecture.)**.** *Fix an isometry $V : \mathbb{C}^N \to \mathbb{C}^M$ and an M-outcome projective measurement $\mathcal{P} = \{\Pi_i\}_{i\in[M]}$ acting on $\mathbb{C}^M$. For any subset $S \subseteq [M]$, let $\Pi_S \coloneqq \sum_{i\in S} \Pi_i$. With high probability over a uniformly random $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$, the "1.5-query" adversary's distinguishing advantage,*

$$\max_{S\subseteq[M]} \max_{f:[M]\to\{\pm1\}} \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \Big[\langle \psi_{\boldsymbol{R_k}} | V^\dagger \cdot \mathcal{O}_f \cdot \Pi_S \cdot \mathcal{O}_f \cdot V | \psi_{\boldsymbol{R_k}} \rangle \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[\langle \psi_{\boldsymbol{h}} | V^\dagger \cdot \mathcal{O}_f \cdot \Pi_S \cdot \mathcal{O}_f \cdot V | \psi_{\boldsymbol{h}} \rangle \Big],$$

is at most $\mathsf{negl}(n)$.

One potential approach towards bounding this expression is to observe that if we fix a subset S, the remaining expression has the same form as the maximum distinguishing advantage for a one-query adversary. We can therefore apply the same weight-vector decomposition described in Section 2.4.1 and invoke a spectral relaxation. The result is that the following is an upper bound for the adversary's distinguishing advantage:

$$\max_{S\subseteq[M]} \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k}} \Big[D^\dagger_{V,\boldsymbol{R_k}} \cdot \Pi_S \cdot D_{V,\boldsymbol{R_k}} \Big] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big[D^\dagger_{V,\boldsymbol{h}} \cdot \Pi_S \cdot D_{V,\boldsymbol{h}} \Big] \Big\|_{\mathrm{op}}, \tag{8}$$

where the definitions of $D_{V,\boldsymbol{h}}$ and $D_{V,\boldsymbol{R_k}}$ are the same as in Section 2.4.1.

The central difficulty we face is that matrix concentration inequalities are not sufficient to bound Eq. (8). Indeed, they can be applied for any *fixed* choice of S, but it is unclear how to bound the operator norm of 2^M matrices simultaneously, one for each S. Nevertheless, we believe that the expression (8) is in fact negligible with high probability over $\boldsymbol{R}$.

The "$(1+\varepsilon)$-query" conjecture. Finally, we highlight a sub-class of 1.5-query adversaries that we do not know how to rule out, which we refer to as $(1+\varepsilon)$-query adversaries. Instead of making an arbitrary first query to the oracle, these adversaries use their first query to synthesize an advice state $|v\rangle \in \mathbb{C}^{M/N}$ (similar to the adversaries we considered in Section 2.2); note that while $|v\rangle$ is technically restricted to states of a certain type, treating it as an arbitrary unit vector is essentially without loss of generality.[3]

Conjecture 2.6 (The $(1+\varepsilon)$-query conjecture.)**.** *Fix any M-outcome projective measurement $\mathcal{P} = \{\Pi_i\}_{i\in[M]}$ acting on $\mathbb{C}^M$. With high probability over a uniformly random $\boldsymbol{R} : [K]\times[N] \to \{\pm 1\}$, the adversary's distinguishing advantage*

$$\max_{S\subseteq[M]} \left\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[(\langle \psi_{\boldsymbol{R_k}}| \otimes \mathsf{Id}) \cdot \Pi_S \cdot (|\psi_{\boldsymbol{R_k}}\rangle \otimes \mathsf{Id}) \right] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \left[(\langle \psi_{\boldsymbol{h}}| \otimes \mathsf{Id}) \cdot \Pi_S \cdot (|\psi_{\boldsymbol{h}}\rangle \otimes \mathsf{Id}) \right] \right\|_{\mathrm{op}} \tag{9}$$

is at most $\mathsf{negl}(n)$.

Again, the difficulty we face in bounding Eq. (9) is that matrix concentration inequalities only seem to apply when the subset S is fixed, and not when we maximize over all S.

A simple mathematical conjecture. Finally, in order to state the simplest mathematical conjecture that captures this "simultaneous matrix concentration" problem, we give a slightly different version of the above $(1+\varepsilon)$-query conjecture (which corresponds to the case where $|\psi_{\boldsymbol{R}_k}\rangle$ and $|\psi_{\boldsymbol{h}}\rangle$ are Haar random).

Conjecture 2.7. *Let $c > 0$ be any constant. Set parameters $N = 2^n$, $K = N/2$ and $M = 2^n \cdot 2^{n^c}$. Let $\{\Pi_i\}_{i\in[M]}$ be projectors acting on $\mathbb{C}^M$ such that $\sum_{i\in[M]} \Pi_i = \mathsf{Id}_M$. Sample K Haar-random unit vectors $|\boldsymbol{\psi}_1\rangle, \ldots, |\boldsymbol{\psi}_K\rangle \in \mathbb{C}^N$. Then with probability $1 - \mathsf{negl}(n)$,*

$$\max_{S\subseteq[M]} \left\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \left[(\langle \boldsymbol{\psi}_{\boldsymbol{k}}| \otimes \mathsf{Id}) \cdot \Pi_S \cdot (|\boldsymbol{\psi}_{\boldsymbol{k}}\rangle \otimes \mathsf{Id}) \right] - \mathop{\mathbf{E}}_{|\boldsymbol{\psi}\rangle} \left[(\langle \boldsymbol{\psi}| \otimes \mathsf{Id}) \cdot \Pi_S \cdot (|\boldsymbol{\psi}\rangle \otimes \mathsf{Id}) \right] \right\|_{\mathrm{op}} = \mathsf{negl}(n).$$

where each Id is $M/N \times M/N$-dimensional and $|\boldsymbol{\psi}\rangle$ is Haar-random.

We note that it would be extremely surprising to us if Conjecture 2.7 turns out to be false, since that would imply that a two-query algorithm can win (the Haar-random state version of) the Oracle State Distinguishing Game.

[3] For example, the adversary can implement Rosenthal's state synthesis algorithm [Ros23a] to prepare an arbitrary quantum advice state.

3 The Oracle State Distinguishing Game

The purpose of this section is to define the Oracle State Distinguishing Game and prove several fundamental properties about it. We note that our main proof in Section 4 can be understood without reading Sections 3.3 to 3.5.

The section is organized as follows. In Section 3.1, we introduce some notation and formalism for oracle algorithms. In Section 3.2, we define the Oracle State Distinguishing Game. In Section 3.3 we show that hardness of the Oracle State Distinguishing Game for T-query adversaries implies hardness of T-query unitary synthesis for any parameter T.

In Section 3.4, we appeal to concentration of measure to give (for any oracle adversary A) an upper tail inequality on the optimal distinguishing advantage in the oracle state distinguishing game, which implies that it suffices to bound the adversary's expected distinguishing advantage over the choice of $\boldsymbol{R}$.

In Section 3.5, we show that two complexity measures of an oracle algorithm — *query length* and *space complexity* — are tightly related in the oracle state distinguishing game. The assumption that our adversaries are space-efficient as well as query-efficient will be crucial in both proofs of the one-query lower bound.

Finally, in Section 3.6, we give an explicit "normal form" for one-query adversaries (using Section 3.5), setting up simplified notation that suffices for Section 4.

3.1 Preliminary notation

We will use **boldface** to denote random variables. We will write $\ln(\cdot)$ for the natural logarithm and $\log_2(\cdot)$ for the base-2 logarithm.

Notation 3.1 (Register size versus dimension)**.** A quantum register consisting of m qubits has dimension $M = 2^m$. Viewing it as a space of m qubits, it is natural to index the basis by binary strings $x \in \{0,1\}^m$. On the other hand, viewing it as a space of dimension M, it is natural to index the basis by integers $1 \leq i \leq M$. We can associate these two indexing schemes by associating the number i with the string x that is the m-bit binary representation of $i-1$. We will typically prefer the second indexing scheme, and will therefore typically represent m-qubit states as

$$|\psi\rangle = \sum_{i=1}^{M} \psi_i \cdot |i\rangle \,.$$

Throughout this work, we will consider algorithms which take as input a quantum state. We will typically reserve n for the length of the input register in qubits and $N \coloneqq 2^n$ for the dimension of this register.

Most of our quantum state inputs will come in the form of binary phase states.

Definition 3.2 (Binary phase state)**.** A *Boolean function* is a function $h : \{0,1\}^n \to \{\pm 1\}$. Due to the association between $\{0,1\}^n$ and $[N]$ given in Notation 3.1, we will typically prefer to write such a function as $h : [N] \to \{\pm 1\}$, and we will elect to still refer to such a function as a "Boolean function". The corresponding *binary phase state* is

$$|\psi_h\rangle \coloneqq \frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} h(x) \cdot |x\rangle \,.$$

Definition 3.3 (Phase oracle)**.** Let $f : [L] \to \{\pm 1\}$ be a Boolean function. Then the corresponding *phase oracle* is the $L \times L$ diagonal unitary matrix $\mathcal{O}_f$ given by

$$\mathcal{O}_f = \sum_{i=1}^{L} f(i) \cdot |i\rangle\langle i| \,.$$

Operationally, for any $1 \leq i \leq L$, the oracle acts as $\mathcal{O}_f \cdot |i\rangle = f(i) \cdot |i\rangle$. If $L = 2^\ell$ for some integer ℓ, then we refer to ℓ as the *input length* of the phase oracle and L as the *dimension* of the phase oracle.

Phase oracles can be contrasted with *bit flip oracles*. A bit flip oracle $\mathcal{O}_g^{\mathsf{flip}}$ is specified by a function $g : [M] \to \{0,1\}$ and is defined as follows: for each $1 \leq i \leq M$ and $b \in \{0,1\}$, it acts as $\mathcal{O}_g^{\mathsf{flip}} \cdot |i, b\rangle = |i, b \oplus g(i)\rangle$. In general, a bit flip oracle can always be used to implement a phase oracle, but the reverse is only partially true: implementing a bit flip oracle requires a *controlled* phase oracle. However, we will see below that the class of phase oracles we consider are actually powerful enough to implement controlled phase oracles, and hence can be converted to bit flip oracles if desired.

Throughout this work, we will consider a class of circuits which take as input a quantum state and are allowed to perform several queries to a phase oracle. We define these formally as follows.

Definition 3.4 (Oracle circuit)**.** A t-query *oracle circuit* $A^{(\cdot)}$ begins with an input register of n qubits and an ancilla register of a qubits, each initialized to $|0\rangle$, for a total of $m = n + a$ qubits. It then performs the m-qubit unitaries $U_1, \ldots, U_{t+1}$. In addition, between each pair of unitaries, it performs a query to a phase oracle of input length ℓ, which acts on the first ℓ qubits. We write $A^{(\cdot)} = (n, m, \ell, U_1, \ldots, U_{t+1})$ in order to specify these parameters.

The precise execution of the oracle circuit depends on which Boolean function it is given query access to. Given a Boolean function $f : \{0,1\}^\ell \to \{\pm 1\}$, we write A^f for the oracle circuit given access to f. On input an n-qubit state, it computes the state

$$U_{t+1} \cdot \mathcal{O}_f \cdot U_t \cdots \mathcal{O}_f \cdot U_2 \cdot \mathcal{O}_f \cdot U_1 \cdot |\psi\rangle \otimes |0^a\rangle \,.$$

This is illustrated in Figure 1.

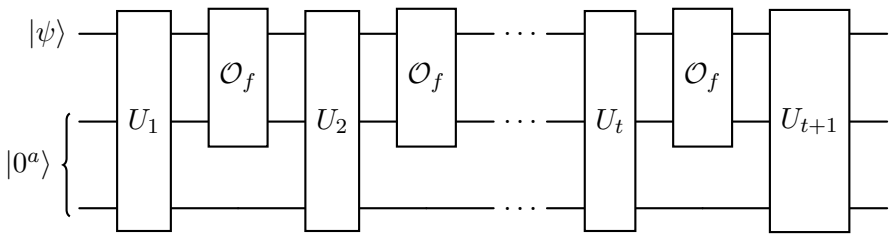

Figure 1: The execution of A^f on input $|\psi\rangle$.

Remark 3.5 (Querying multiple functions)**.** Note that we have defined our oracle circuits so that every application of the oracle gate queries the same Boolean function f. One can consider an alternative model of t-query oracle circuits which are instead allowed to query a different Boolean function f_i for each oracle call $1 \leq i \leq t$. However, one can simulate access to these t Boolean functions using a single Boolean function f defined as $f(\mathrm{bin}(i), x) := f_i(x)$, where $\mathrm{bin}(i)$ is the

$a = \lceil \log_2(t) \rceil$-bit binary encoding of i. Hence, an adversary which queries t different Boolean functions can be simulated by an adversary which queries one Boolean function and has a small a-qubit overhead. Thus, it is essentially without loss of generality to focus on adversaries which query a single function, as we do. We note that this transformation is standard and appears, for example, at the top of page 5 in Rosenthal's Ph.D. thesis [Ros23b].

Remark 3.6 (Querying many-bit functions)**.** Yet another model of t-query oracle circuits allows for making bit-flip queries to d-bit output functions of the form $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus f(x)\rangle$ for $x \in [M], y \in \{0,1\}^d$. As pointed out in [Ros23b] Section 2.1, such queries can be simulated by a single quantum query to a 1-bit function:

$$|x\rangle|r\rangle \mapsto (-1)^{r \cdot f(x)}|x\rangle|r\rangle .$$

The new function $g(x, r) = r \cdot f(x)$ has domain $[M \cdot 2^d]$. This also allows us to simulate *parallel queries* of the form

$$|x_1\rangle|b_1\rangle \ldots |x_t\rangle|b_t\rangle \mapsto |x_1\rangle|b_1 \oplus f(x_1)\rangle \ldots |x_t\rangle|b_t \oplus f(x_t)\rangle$$

by defining $x = (x_1, \ldots, x_t)$. Therefore, our one-query lower bounds imply lower bounds against a bounded (e.g., polynomial or sub-exponential) number of parallel queries.

3.2 Defining the Oracle State Distinguishing Game

In this section, we define the Oracle State Distinguishing Game. To begin, every such game is parameterized by a particular family of functions, which is defined as follows.

Definition 3.7 (Function families)**.** Let K and N be integers. A *function family* is a function $R : [K] \times [N] \to \{\pm 1\}$. We think of R as defining a family of K Boolean functions as follows: for each $1 \leq k \leq K$, we let $R_k : [N] \to \{\pm 1\}$ be the function $R_k(\cdot) \coloneqq R(k, \cdot)$.

We have chosen the letter "R" for function families as shorthand for the word "Random", as our function families will often (though not always) be random variables.

Definition 3.8 (Oracle State Distinguishing Game)**.** Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. The *Oracle State Distinguishing Game on* R, denoted Game^R, involves two parties, a challenger and an adversary. It is played as follows.

1. The challenger samples a random bit $\boldsymbol{b} \in \{0, 1\}$.
2. The challenger generates a random n-qubit state $|\boldsymbol{\psi}\rangle$ in one of two ways:
   - If $\boldsymbol{b} = 0$, the challenger samples a uniformly random $\boldsymbol{k} \sim [K]$ and generates $|\boldsymbol{\psi}\rangle \coloneqq |\psi_{R_{\boldsymbol{k}}}\rangle$.
   - If $\boldsymbol{b} = 1$, the challenger samples a uniformly random $\boldsymbol{x} \sim [N]$ and sets $|\boldsymbol{\psi}\rangle \coloneqq |\boldsymbol{x}\rangle$.
3. The challenger sends $|\boldsymbol{\psi}\rangle$ to the adversary.
4. The adversary outputs a bit $\boldsymbol{b}' \in \{0, 1\}$.
5. If $\boldsymbol{b}' = \boldsymbol{b}$, then the adversary wins. Otherwise, they lose.

The *Oracle State Distinguishing Game*, denoted $\mathsf{Game}_{N,K}$ is played as follows. A uniformly random function family $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is sampled, and then the challenger plays $\mathsf{Game}^{\boldsymbol{R}}$ with the adversary.

Remark 3.9 (Computational complexity of the challenger)**.** In the "$\boldsymbol{b} = 1$ case", the view of the adversary is that it receives a maximally mixed state Id_N/N. Hence, we can equivalently view the challenger as sampling a random state from *any* distribution, so long as an average state drawn from this distribution is maximally mixed. For example, we can equivalently view the challenger as sampling a uniformly random Boolean function $\boldsymbol{h} : [N] \to \{\pm 1\}$ and setting $|\boldsymbol{\psi}\rangle \coloneqq |\psi_{\boldsymbol{h}}\rangle$, or sampling $|\boldsymbol{\psi}\rangle$ as an N-dimensional Haar-random state. We will typically prefer the first of these points of view throughout this work.

We have chosen to have the challenger sample a random basis state $|\boldsymbol{x}\rangle$ in this case to emphasize that the challenger is computationally efficient in our construction. Note that they can also efficiently construct the state $|\psi_{R_{\boldsymbol{k}}}\rangle$ in the "$\boldsymbol{b} = 0$ case" given oracle access to R. In particular, they need only query the oracle $R(\boldsymbol{k}, \cdot)$ on the uniform superposition state $|+_N\rangle \coloneqq \frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} |x\rangle$.

We will model our adversary as an oracle circuit $A^{(\cdot)}$ with an N-dimensional input register. Intuitively, the adversary will be allowed to select its own preferred oracle f to give it the best chance of winning the Oracle State Distinguishing Game on R. When the game is played on a uniformly random choice of the function family $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$, the adversary will be allowed to select an oracle $f_{\boldsymbol{R}}$ which depends on $\boldsymbol{R}$.

Definition 3.10 (Adversary)**.** An *adversary* is specified by an oracle circuit $A^{(\cdot)}$. Let L be the dimension of the queries the oracle circuit makes. Given oracle access to a Boolean function $f : [L] \to \{\pm 1\}$, the adversary acts as follows. On input the quantum state $|\psi\rangle$, it applies A^f, and then it measures the first qubit in the standard basis. It outputs the measurement outcome $\boldsymbol{b}' \in \{0, 1\}$.

Now we introduce several pieces of notation which will help us describe the adversary's winning probability in the Oracle State Distinguishing Game.

Notation 3.11 (Adversary's acceptance probability)**.** Let $A^{(\cdot)}$ be an adversary which has an N-dimensional input register and makes L-dimensional queries. Let $h : [N] \to \{\pm 1\}$ be a Boolean function, and let $f : [L] \to \{\pm 1\}$ be another Boolean function. We will use the notation

$$p_A(h \mid f) \coloneqq \mathbf{Pr}[A^f \text{ outputs "0" on } |\psi_h\rangle].$$

Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. Then in the "$\boldsymbol{b} = 0$ case", the probability the adversary wins Game^R can be expressed in this notation as $\mathbf{E}_{\boldsymbol{k} \sim [K]}[p_A(R_{\boldsymbol{k}} \mid f)]$. As for the "$\boldsymbol{b} = 1$ case" let us follow Remark 3.9 and view the challenger as sampling a uniformly random Boolean function $\boldsymbol{h} : [N] \to \{\pm 1\}$ and setting $|\boldsymbol{\psi}\rangle \coloneqq |\psi_{\boldsymbol{h}}\rangle$; given this state, the adversary wins with probability $1 - p_A(\boldsymbol{h} \mid f)$. Putting these two together, the probability the adversary wins Game^R is

$$\frac{1}{2} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]}[p_A(R \mid f)] + \frac{1}{2} \cdot \mathop{\mathbf{E}}_{\boldsymbol{h}}[1 - p_A(\boldsymbol{h} \mid f)]. \tag{10}$$

Note that the adversary can trivially win with probability $1/2$ by always outputting $\boldsymbol{b}' = 0$. Thus, we care about the amount by which the adversary's acceptance probability differs from $1/2$, which is known as its *advantage*.

Definition 3.12 (Distinguishing advantage)**.** Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. Let $A^{(\cdot)}$ be an adversary with an N-dimensional input register. Let L be the dimension of $A^{(\cdot)}$'s queries, and let $f : [L] \to \{\pm 1\}$ be a Boolean function. Then the *distinguishing advantage of* A^f *on* Game^R is defined as

$$\Delta_A(R \mid f) \coloneqq 2 \cdot \left| \mathbf{Pr}[A^f \text{ wins on } \mathsf{Game}^R] - \frac{1}{2} \right|.$$

The factor of 2 in front was chosen so that the distinguishing advantage is a number between 0 and 1 and is equal to 1 if the adversary always wins (or loses). If we plug in Equation (10) for the adversary's winning probability, we can rewrite the distinguishing advantage as

$$\Delta_A(R \mid f) = \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_A(\boldsymbol{h} \mid f)] \right|,$$

where $\boldsymbol{h} : [N] \to \{\pm 1\}$ is a uniformly random Boolean function. This equation is the form that we will most typically express the distinguishing advantage in, and it explains why we refer to this as the *distinguishing* advantage, which is because it expresses how well the adversary's output can be used to distinguish between the two cases.

The adversary's goal is to maximize the distinguishing advantage, and it can do so by picking the best possible function $f : [L] \to \{\pm 1\}$ to perform oracle queries to. This motivates the following quantity, which is the main quantity we will be studying throughout this paper.

Definition 3.13 (Maximum distinguishing advantage)**.** Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. Let $A^{(\cdot)}$ be an adversary which has an N-dimensional input register and makes L-dimensional queries. The *maximum distinguishing advantage of* $A^{(\cdot)}$ *on* Game^R is defined as

$$\Delta_A(R) \coloneqq \max_{f:[L]\to\{\pm 1\}} \{\Delta_A(R \mid f)\}.$$

Finally, the *maximum distinguishing advantage of* $A^{(\cdot)}$ *on* $\mathsf{Game}_{K,N}$ is equal to

$$\Delta_A^{\mathrm{avg}} \coloneqq \mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})],$$

where $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is a uniformly random function family.

The goal of this work is to show that Δ_A^{avg} is small for any adversary $A^{(\cdot)}$ which makes a single query of length $\ell = o(K)$. Moreover, we will prove that in the same parameter regime, with high probability over $\boldsymbol{R}$, $\Delta_A(\boldsymbol{R})$ is small.

3.3 Relationship to the Unitary Synthesis Problem

In this section, we formalize the Unitary Synthesis Problem and its relationship to the Oracle State Distinguishing Game. Or, rather, we will suggest one possible way of formalizing the Unitary Synthesis Problem, as there seems to be no generally agreed upon precise formulation of the problem. For example, the task is to approximate a general n-qubit unitary U, but there are many different ways of defining what it means to approximate a unitary. This was addressed by Scott Aaronson in a comment on the Shtetl-Optimized blog [Aar21], in which he said the following.

> *"The unitary synthesis problem is interesting for any reasonable notion of approximating U. In other words, we lack a positive result even for the loosest notions of approximation you mentioned, or a negative result even for the most stringent ones! Once we have some results, then we can start worrying about these distinctions."*

The last few years have seen increasing interest in fundamentally quantum tasks, and as a result we now do have some results on problems related to unitary synthesis [RY22, Ros22, BEM+23], and these have given several ways of precisely formalizing unitary synthesis.

Let us first recall several standard notions from quantum information theory. Given two n-qubit density matrices ρ_1 and ρ_2, their *trace distance* is given by

$$\mathrm{D}_{\mathrm{tr}}(\rho_1, \rho_2) \coloneqq \frac{1}{2} \cdot \|\rho_1 - \rho_2\|_1,$$

where $\|\cdot\|_1$ is the trace norm. Given two quantum channels Φ_1, Φ_2, both with n-qubit inputs and outputs, their *diamond distance* is given by

$$\mathrm{D}_{\diamond}(\Phi_1, \Phi_2) \coloneqq \max_{\rho} \{\mathrm{D}_{\mathrm{tr}}((\Phi_1 \otimes \mathsf{Id})(\rho), (\Phi_2 \otimes \mathsf{Id})(\rho))\},$$

where the maximization is over all $2n$-qubit density matrices ρ, and both Id operators refer to the n-qubit identity channel. For more background these distances, see [Wat18, Chapter 3].

Following [BEM+23], we will define what it means to approximate a unitary in terms of the diamond distance.

Definition 3.14 (Approximating a unitary)**.** Let U be an n-qubit unitary, and let Φ_U be the associated quantum channel. Let Φ_{approx} be a quantum channel with n-qubit input and output registers. Then Φ_{approx} is said to ε-*approximate* U if $\mathrm{D}_{\diamond}(\Phi_{\mathrm{approx}}, \Phi_U) \leq \varepsilon$.

We will also define the channel associated with an oracle circuit $A^{(\cdot)}$ in the natural way.

Definition 3.15 (Channel implemented by an oracle circuit)**.** Given a t-query oracle circuit $A^{(\cdot)} = (n, m, \ell, U_1, \ldots, U_t, U_{t+1})$ and a Boolean function $f : \{0,1\}^{\ell} \to \{\pm 1\}$, the associated n-*qubit channel* Φ_{A^f} is defined as follows:

1. Given an n qubit input $|\psi\rangle$, compute the state

$$U_{t+1} \cdot \mathcal{O}_f \cdot U_t \cdot \cdots \cdot \mathcal{O}_f \cdot U_2 \cdot \mathcal{O}_f \cdot U_1 \cdot |\psi\rangle \, |0^{m-n}\rangle$$

2. Return the first n qubits as the output (and discard the rest).

With these definitions in hand, we can give a formal statement of the Unitary Synthesis Problem.

Definition 3.16 (The Unitary Synthesis Problem)**.** Fix an error parameter $\varepsilon(n) = 1/2^{\Omega(n)}$. Does there exist a $\mathrm{poly}(n)$-query oracle circuit $A^{(\cdot)}$ computable by a $\mathrm{poly}(n)$-sized quantum circuit such that for all n-qubit unitaries U, there exists a Boolean function $f : \{0,1\}^* \to \{\pm 1\}$ such that $\mathrm{D}_{\diamond}(\Phi_{A^f}, \Phi_U) \leq \varepsilon(n)$?

As discussed in Section 1.1, a bound on the maximum distinguishing advantage in the Oracle Distinguishing Game immediately implies a lower bound for the worst-case version of the Unitary Synthesis Problem, since there always *exists* an information-theoretic distinguisher that wins the corresponding Oracle State Distinguishing Game for R. In fact, if we make a slight modification to the Oracle State Distinguishing Game, then a the distinguishing advantage bound would imply a slightly stronger claim, namely that Unitary Synthesis Problem is hard for a Haar-random U (we note that this is technically the version of the problem stated by Aaronson and Kuperberg [AK07]).

In more detail, one can consider a variant of the Oracle State Distinguishing Game where every $|\psi_{\boldsymbol{R}_k}\rangle$ is sampled as a Haar random state, rather than as a binary phase state (we do not give a separate analysis for the version of Oracle State Distinguishing with Haar random states, but our proof technique can easily be adapted to handle it). Next, suppose that there exists an oracle circuit $A^{(\cdot)}$ that can synthesize an n-qubit Haar random unitary U. Then for a random U and any $K < N$, there exists (with high probability) a choice of f such that A^f implements the channel corresponding

to U. In particular, this means that for a Haar-random subspace $S = \mathrm{span}\{U^\dagger |1\rangle, \ldots, U^\dagger |K\rangle\}$, there exists f such that A^f maps S to $\mathrm{span}\{|1\rangle, \ldots, |K\rangle\}$. Such an oracle circuit $A^{(\cdot)}$ can be used to win the Oracle State Distinguishing Game, since the subspace $\mathrm{span}\{|\psi_{\boldsymbol{R}_1}\rangle, \ldots, |\psi_{\boldsymbol{R}_K}\rangle\}$ is distributed as a K-dimensional Haar-random subspace, and the ability to map this subspace to $\mathrm{span}\{|1\rangle, \ldots, |K\rangle\}$ immediately yields a distinguisher for the game.

To summarize, we have argued that a lower bound for breaking a (single-copy) pseudorandom state family — in an oracle setting where the K pseudorandom states are distributed as Haar random states — directly implies hardness of synthesizing the first K columns of a Haar-random unitary. Thus, we have the following claim.

Claim 3.17. *If the maximum distinguishing advantage of any efficient t-query adversary in the oracle distinguishing game is $o(1)$, then there is no efficient t-query oracle algorithm for the Unitary Synthesis Problem on a Haar-random unitary U.*

3.4 Upper tail inequality for the maximum distinguishing advantage

Throughout this subsection, we will write $A^{(\cdot)} = (n, m, \ell, U_1, \ldots, U_{t+1})$ for a t-query adversary with an $(N := 2^n)$-dimensional input register and $(L := 2^\ell)$-dimensional queries which is playing Game^R for function families of the form $R : [K] \times [N] \to \{\pm 1\}$.

Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. In this section, we consider the random variable $\Delta_A(\boldsymbol{R})$ corresponding to the maximum distinguishing advantage of A (Definition 3.13), and we show that it has strong one-sided concentration around its mean Δ_A. Our main result is as follows.

Lemma 3.18 (Upper tail for $\Delta_A(\boldsymbol{R})$)**.** *There exists a constant $c > 0$ such that the following is true. Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Then for all $\varepsilon \geq 0$,*

$$\Pr_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \geq \Delta_A^{\mathrm{avg}} + \varepsilon] \leq 4 \cdot \exp\left(-c \cdot \varepsilon^2 KN\right).$$

The main technical lemma we will need to prove this is the following version of Talagrand's concentration inequality, which is stated in [Ver18, Theorem 5.2.16].

Definition 3.19 (Lipschitz)**.** Let $g : [-1, 1]^n \to \mathbb{R}$ be a function. It has *Lipschitz constant C* if for all $u, v \in [-1, 1]^n$,

$$|g(u) - g(v)| \leq C \cdot \|u - v\|_2.$$

Theorem 3.20 (Talagrand's concentration inequality)**.** *There exists a constant $c > 0$ such that the following is true. Let $g : [-1, 1]^d \to \mathbb{R}$ be a convex function with Lipschitz constant C. Let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_d$ be independent random variables satisfying $|\boldsymbol{v}_i| \leq 1$ for all $1 \leq i \leq d$. Then for all $t \geq 0$,*

$$\mathbf{Pr}\left[|g(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_d) - \mathbf{E}[g(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_d)]| \geq t\right] \leq 2 \cdot \exp\left(-\frac{c \cdot t^2}{C^2}\right).$$

To derive Lemma 3.18 using Talagrand's concentration inequality, we will view a uniformly random function family $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ as a collection of KN independent $\{\pm 1\}$ random variables. We would then like to apply Talagrand's concentration inequality with the "g" function set to the maximum distinguishing advantage $\Delta_A(\cdot)$, interpreted as a function of an input R. However, doing so faces two difficulties: first, $\Delta_A(\cdot)$ is defined only for $\{\pm 1\}$-valued inputs, wheres the "g" function in Talagrand's concentration inequality must be defined over $[-1, 1]$ inputs. Second, $\Delta_A(\cdot)$ is not convex. The first difficulty is straightforward to address, and we begin to do so in the following definition.

Definition 3.21 (Expanding the acceptance probability to bounded inputs)**.** Let us fix a Boolean function $f : [L] \to \{\pm 1\}$. Given as input the state $|\psi\rangle$, the adversary applies the oracle circuit A^f and then measures the first qubit of the resulting state. We can therefore view the adversary as applying a POVM measurement $E^f := \{E_0^f, E_1^f\}$ to $|\psi\rangle$, where

$$E_0^f := (A^f)^\dagger \cdot (|0\rangle\langle 0| \otimes \mathsf{Id}_2^{\otimes m-1}) \cdot A^f,$$

and $E_1^f := \mathsf{Id}_N - E_0^f$. As a result, for any Boolean function $h : [N] \to \{\pm 1\}$, we can write

$$p_A(h \mid f) = \langle \psi_h | \cdot E_0^f \cdot |\psi_h\rangle . \tag{11}$$

We will now extend this expression to functions which are $[-1, 1]$-valued rather than $\{\pm 1\}$-valued. For a bounded function $\underline{h} : [N] \to [-1, 1]$ and a bit $b \in \{0, 1\}$, we define

$$\left|\psi_{\underline{h}}\right\rangle := \frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} \underline{h}(x) \cdot |x\rangle , \quad \text{and} \quad p_{A,b}(\underline{h} \mid f) := \left\langle \psi_{\underline{h}} \right| \cdot E_b^f \cdot \left|\psi_{\underline{h}}\right\rangle .$$

Note that $\left|\psi_{\underline{h}}\right\rangle$ is *sub-normalized*, meaning that $\left\langle \psi_{\underline{h}} | \psi_{\underline{h}} \right\rangle \leq 1$, and so it is no longer necessarily a quantum state. In addition, if $h : [N] \to \{\pm 1\}$ is a Boolean function, then by Equation (11), $p_{A,b}(\underline{h} \mid f)$ still recovers our traditional definition of $p_A(\cdot \mid f)$ when $b = 0$. As for the $b = 1$ case, note that because h is a Boolean function,

$$\begin{aligned} p_{A,1}(h \mid f) = \langle \psi_h | \cdot E_1^f \cdot |\psi_h\rangle &= \langle \psi_h | \cdot (\mathsf{Id}_N - E_1^f) \cdot |\psi_h\rangle \\ &= 1 - \langle \psi_h | \cdot E_0^f \cdot |\psi_h\rangle = 1 - p_{A,0}(h \mid f). \end{aligned} \tag{12}$$

However, this is not necessarily true of bounded functions $\underline{h}$.

Now we address the second issue, that of $\Delta_A(\cdot)$ not being convex. To do so, we will have to define two variants of $\Delta_A(\cdot)$ called $\Delta_{A,0}(\cdot)$ and $\Delta_{A,1}(\cdot)$ which we will eventually show *are* convex. This motivates the following definition, which will only be used in this subsection.

Definition 3.22 (Modifying the distinguishing advantage)**.** Let $f : [L] \to \{\pm 1\}$ be a Boolean function and $\underline{R} : [K] \times [N] \to [-1, 1]$ be a bounded function. For $b \in \{0, 1\}$, we define

$$\begin{aligned} \Delta_{A,b}(\underline{R} \mid f) &:= \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}} [p_{A,b}(\boldsymbol{h} \mid f)], \\ \Delta_{A,b}(\underline{R}) &:= \max_{f:[L]\to\{\pm 1\}} \{\Delta_{A,b}(\underline{R} \mid f)\}. \end{aligned}$$

We note that unlike in the definition of $\Delta_A(\cdot)$, there is no absolute value in the definition of $\Delta_{A,b}(\cdot)$. (This is needed so that we can later show that it is convex.) Finally, we define

$$\Delta_{A,b}^{\text{avg}} := \mathop{\mathbf{E}}_{\boldsymbol{R}} [\Delta_{A,b}(\boldsymbol{R})],$$

where $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is a uniformly random function family.

We will now make some observations about these definitions. Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. Then by Equation (12),

$$\begin{aligned} \Delta_{A,1}(R \mid f) &= \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,1}(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}} [p_{A,1}(\boldsymbol{h} \mid f)] \\ &= \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [1 - p_{A,0}(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}} [1 - p_{A,0}(\boldsymbol{h} \mid f)] \\ &= - \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,0}(R_{\boldsymbol{k}} \mid f)] + \mathop{\mathbf{E}}_{\boldsymbol{h}} [p_{A,0}(\boldsymbol{h} \mid f)] = -\Delta_{A,0}(R \mid f). \end{aligned}$$

(As before, this is not necessarily true of bounded functions $\underline{R}$.) Thus, we have that

$$\begin{aligned}\Delta_A(R \mid f) = |\Delta_{A,0}(R \mid f)| &= \max\{\Delta_{A,0}(R \mid f), -\Delta_{A,0}(R \mid f)\} \\ &= \max\{\Delta_{A,0}(R \mid f), \Delta_{A,1}(R \mid f)\}.\end{aligned}$$

As a result,

$$\begin{aligned}\Delta_A(R) &= \max_{f:[L]\to\{\pm1\}} \{\Delta_A(R \mid f)\} \\ &= \max_{f:[L]\to\{\pm1\}} \big\{ \max\{\Delta_{A,0}(R \mid f), \Delta_{A,1}(R \mid f)\}\big\} \\ &= \max\Big\{ \max_{f:[L]\to\{\pm1\}} \{\Delta_{A,0}(R \mid f)\}, \max_{f:[L]\to\{\pm1\}} \{\Delta_{A,1}(R \mid f)\}\Big\} \\ &= \max\{\Delta_{A,0}(R), \Delta_{A,1}(R)\}. \end{aligned} \tag{13}$$

We will show the following concentration bound for these two variants of the maximum distinguishing advantage.

Lemma 3.23 (Concentration of the modified distinguishing advantages)**.** *There exists an absolute constant $c > 0$ such that the following is true. Let $b \in \{0,1\}$. Let $\boldsymbol{R} : [K] \times [N] \to \{\pm1\}$ be a uniformly random function family. Then for all $\varepsilon \geq 0$,*

$$\mathop{\mathbf{Pr}}_{\boldsymbol{R}}[|\Delta_{A,b}(\boldsymbol{R}) - \Delta^{\mathrm{avg}}_{A,b}| \geq \varepsilon] \leq 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big).$$

Before proving Lemma 3.23, let us see how it implies the main result of this subsection, Lemma 3.18.

Proof of Lemma 3.18. First, we note that by Equation (13), $\Delta_A(R) \geq \Delta_{A,b}(R)$ for any function family $R : [K] \times [N] \to \{\pm1\}$ and any $b \in \{0,1\}$. Thus, we have that $\Delta^{\mathrm{avg}}_A \geq \Delta^{\mathrm{avg}}_{A,b}$. Next, for a uniformly random function family $\boldsymbol{R} : [K] \times [N] \to \{\pm1\}$,

$$\begin{aligned}
&\mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_A + \varepsilon] \\
&= \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\max\{\Delta_{A,0}(\boldsymbol{R}), \Delta_{A,1}(\boldsymbol{R})\} \geq \Delta^{\mathrm{avg}}_A + \varepsilon] && \text{(by Equation (13))} \\
&= \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,0}(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_A + \varepsilon, \text{ or } \Delta_{A,1}(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_A + \varepsilon] \\
&\leq \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,0}(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_A + \varepsilon] + \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,1}(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_A + \varepsilon] && \text{(by the union bound)} \\
&\leq \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,0}(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_{A,0} + \varepsilon] + \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,1}(\boldsymbol{R}) \geq \Delta^{\mathrm{avg}}_{A,1} + \varepsilon] && \text{(because } \Delta^{\mathrm{avg}}_A \geq \Delta^{\mathrm{avg}}_{A,0} \text{ and } \Delta^{\mathrm{avg}}_{A,1}\text{)} \\
&\leq \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[|\Delta_{A,0}(\boldsymbol{R}) - \Delta^{\mathrm{avg}}_{A,0}| \geq \varepsilon] + \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[|\Delta_{A,1}(\boldsymbol{R}) - \Delta^{\mathrm{avg}}_{A,1}| \geq \varepsilon] \\
&\leq 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big) + 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big) && \text{(by Lemma 3.23)} \\
&= 4 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big).
\end{aligned}$$

This completes the proof. □

Now we focus on proving Lemma 3.23. To do so, we would like to show that $\Delta_{A,0}(\cdot)$ and $\Delta_{A,1}(\cdot)$ are convex and Lipschitz. Prior to doing so, however, we will first prove this for the $p_{A,b}(\cdot \mid f)$ function.

Lemma 3.24 (The $p_{A,b}$ functions are convex and Lipschitz)**.** *Let $f : [L] \to \{\pm1\}$ be a Boolean function. Let $b \in \{0,1\}$. Then $p_{A,b}(\cdot \mid f)$ is convex and $(2/\sqrt{N})$-Lipschitz.*

Proof. Consider two bounded functions $\underline{h}, \underline{h}' : [N] \to [-1, 1]$. Let $0 \le t \le 1$. Then

$$\begin{aligned} p_{A,b}(t \cdot \underline{h} + (1-t) \cdot \underline{h}' \mid f) &= \langle \psi_{t \cdot \underline{h} + (1-t) \cdot \underline{h}'} | \cdot E_b^f \cdot |\psi_{t \cdot \underline{h} + (1-t) \cdot \underline{h}'} \rangle \\ &= \| (E_b^f)^{1/2} \cdot |\psi_{t \cdot \underline{h} + (1-t) \cdot \underline{h}'} \rangle \|_2^2 \\ &= \| (E_b^f)^{1/2} \cdot (t \cdot |\psi_{\underline{h}}\rangle + (1-t) \cdot |\psi_{\underline{h}'}\rangle) \|_2^2. \end{aligned}$$

Now, because $\| \cdot \|_2$ is convex and $x \mapsto x^2$ is convex, we also have that $\| \cdot \|_2^2$ is convex. Hence, by Jensen's inequality, this is at most

$$\begin{aligned} & t \cdot \| (E_b^f)^{1/2} |\psi_{\underline{h}}\rangle \|_2^2 + (1-t) \cdot \| (E_b^f)^{1/2} \cdot |\psi_{\underline{h}'}\rangle \|_2^2 \\ &= t \cdot \langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}}\rangle + (1-t) \cdot \langle \psi_{\underline{h}'} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle \\ &= t \cdot p_{A,b}(\underline{h} \mid f) + (1-t) \cdot p_{A,b}(\underline{h}' \mid f). \end{aligned}$$

And so $p_{A,b}(\cdot \mid f)$ is convex. Next,

$$\begin{aligned} & |p_{A,b}(\underline{h} \mid f) - p_{A,b}(\underline{h}' \mid f)| \\ &= |\langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}}\rangle - \langle \psi_{\underline{h}'} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle| \\ &= |\langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}}\rangle - \langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle + \langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle - \langle \psi_{\underline{h}'} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle| \\ &\le |\langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}}\rangle - \langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle| + |\langle \psi_{\underline{h}} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle - \langle \psi_{\underline{h}'} | \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle| \\ &= |\langle \psi_{\underline{h}} | \cdot E_b^f \cdot (|\psi_{\underline{h}}\rangle - |\psi_{\underline{h}'}\rangle)| + |(\langle \psi_{\underline{h}} | - \langle \psi_{\underline{h}'} |) \cdot E_b^f \cdot |\psi_{\underline{h}'}\rangle|. \end{aligned}$$

By Cauchy-Schwarz, we can bound the first term by

$$\begin{aligned} \| E_b^f \cdot |\psi_{\underline{h}}\rangle \|_2 \cdot \| |\psi_{\underline{h}}\rangle - |\psi_{\underline{h}'}\rangle \|_2 &\le \| |\psi_{\underline{h}}\rangle \|_2 \cdot \| |\psi_{\underline{h}}\rangle - |\psi_{\underline{h}'}\rangle \|_2 && \text{(because } 0 \preceq E_b^f \preceq I\text{)} \\ &\le \| |\psi_{\underline{h}}\rangle - |\psi_{\underline{h}'}\rangle \|_2 && \text{(because } |\psi_{\underline{h}}\rangle \text{ is sub-normalized)} \\ &= \frac{1}{\sqrt{N}} \cdot \|h - h'\|_2. && \text{(by definition of } |\psi_h\rangle \text{ and } |\psi_{h'}\rangle\text{)} \end{aligned}$$

A similar argument shows that the second term is also bounded by $\|h - h'\|_2 / \sqrt{N}$. Putting these together, this shows that $p_{A,b}(\cdot \mid f)$ is $(2/\sqrt{N})$-Lipschitz. □

Next, we use this lemma to show that $\Delta_{A,b}(\cdot \mid f)$ is also convex and Lipschitz.

Lemma 3.25 (The $\Delta_{A,b}(\cdot \mid f)$ functions are convex and Lipschitz)**.** *Let $f : [L] \to \{\pm 1\}$ be a Boolean function. Let $b \in \{0, 1\}$. Then the map*

$$\underline{R} \mapsto \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f)] \tag{14}$$

is convex and $(2/\sqrt{KN})$-Lipschitz. In addition, $\Delta_{A,b}(\cdot \mid f)$ is also convex and $(2/\sqrt{KN})$-Lipschitz.

Proof. We first prove the lemma for the map in Equation (14). Consider two bounded functions $\underline{R}, \underline{R}' : [K] \times [N] \to [-1, 1]$. Let $0 \le t \le 1$. Then by Lemma 3.24,

$$\begin{aligned} \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,b}(t \cdot \underline{R}_{\boldsymbol{k}} + (1-t) \cdot \underline{R}'_{\boldsymbol{k}} \mid f)] &\le \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [t \cdot p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f) + (1-t) \cdot p_{A,b}(\underline{R}'_{\boldsymbol{k}} \mid f)] \\ &= t \cdot \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f)] + (1-t) \cdot \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A,b}(\underline{R}'_{\boldsymbol{k}} \mid f)]. \end{aligned}$$

Thus, this map is convex. Next,

$$\left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_{A,b}(\underline{R}'_{\boldsymbol{k}} \mid f)] \right| \le \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} |p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f) - p_{A,b}(\underline{R}'_{\boldsymbol{k}} \mid f)|.$$

Now, we apply Lemma 3.24, which states that $p_{A,b}(\cdot \mid f)$ is $(2/\sqrt{N})$-Lipschitz. Hence, we can upper-bound this by

$$\begin{aligned}\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \frac{2}{\sqrt{N}} \cdot \|\underline{R}_{\boldsymbol{k}} - \underline{R}'_{\boldsymbol{k}}\|_2 &= \frac{2}{\sqrt{N}\cdot K} \cdot \sum_{k=1}^{K} \|\underline{R}_k - \underline{R}'_k\|_2 \\ &\le \frac{2}{\sqrt{N}\cdot K} \cdot \sqrt{K \cdot \sum_{k=1}^{K} \|\underline{R}_k - \underline{R}'_k\|_2^2} = \frac{2}{\sqrt{KN}} \cdot \|\underline{R} - \underline{R}'\|_2,\end{aligned}$$

where the inequality is due to Cauchy-Schwarz. Thus, this map is $(2/\sqrt{KN})$-Lipschitz. As for $\Delta_{A,b}(\cdot \mid f)$, we recall that it is defined as follows:

$$\Delta_{A,b}(\underline{R} \mid f) \coloneqq \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_{A,b}(\underline{R}_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_{A,b}(\boldsymbol{h} \mid f)].$$

This is just the map in Equation (14), offset by a constant. Hence, it too is convex and $(2/\sqrt{KN})$-Lipschitz. This completes the proof. □

We have finally reached our goal, which is to show that the $\Delta_{A,b}(\cdot) = \max_f \Delta_{A,b}(\cdot \mid f)$ functions are convex and Lipschitz.

Lemma 3.26 (The $\Delta_{A,b}(\cdot)$ functions are convex and Lipschitz)**.** *Let $b \in \{0,1\}$. Then $\Delta_{A,b}(\cdot)$ is convex and $(2/\sqrt{KN})$-Lipschitz.*

Proof. Consider two bounded functions $\underline{R}, \underline{R}' : [K] \times [N] \to [-1,1]$. Let $0 \le t \le 1$. Then

$$\Delta_{A,b}(t \cdot \underline{R} + (1-t) \cdot \underline{R}') = \max_{f:[L]\to\{\pm1\}} \{\Delta_{A,b}(t \cdot \underline{R} + (1-t) \cdot \underline{R}' \mid f)\}$$

By Lemma 3.25, the function $\Delta_{A,b}(\cdot \mid f)$ is convex. Hence, this is at most

$$\begin{aligned}&\max_{f:[L]\to\{\pm1\}} \{t \cdot \Delta_{A,b}(\underline{R} \mid f) + (1-t) \cdot \Delta_{A,b}(\underline{R}' \mid f)\} \\ &\le \max_{f:[L]\to\{\pm1\}} \{t \cdot \Delta_{A,b}(\underline{R} \mid f) + (1-t) \cdot \Delta_{A,b}(\underline{R}' \mid f)\} \\ &\le t \cdot \max_{f:[L]\to\{\pm1\}} \{\Delta_{A,b}(\underline{R} \mid f)\} + (1-t) \cdot \max_{f:[L]\to\{\pm1\}} \{\Delta_{A,b}(\underline{R}' \mid f)\} \\ &= t \cdot \Delta_{A,b}(\underline{R}) + (1-t) \cdot \Delta_{A,b}(\underline{R}').\end{aligned}$$

Hence, $\Delta_{A,b}(\cdot)$ is convex.

Now we show that $\Delta_{A,b}(\cdot)$ is Lipschitz. To do so, we will show that for any two bounded functions $\underline{R}, \underline{R}' : [K] \times [N] \to [-1,1]$,

$$\Delta_{A,b}(\underline{R}) - \Delta_{A,b}(\underline{R}') \le \frac{2}{\sqrt{KN}}.$$

This will show that $\Delta_{A,b}(\cdot)$ is $(2/\sqrt{KN})$-Lipschitz, as

$$|\Delta_{A,b}(\underline{R}) - \Delta_{A,b}(\underline{R}')| = \max\{\Delta_{A,b}(\underline{R}) - \Delta_{A,b}(\underline{R}'), \Delta_{A,b}(\underline{R}') - \Delta_{A,b}(\underline{R})\} \le \frac{2}{\sqrt{KN}}.$$

To begin,

$$\Delta_{A,b}(\underline{R}) - \Delta_{A,b}(\underline{R}') = \max_{f:[L]\to\{\pm1\}}\{\Delta_{A,b}(\underline{R} \mid f)\} - \max_{f':[L]\to\{\pm1\}}\{\Delta_{A,b}(\underline{R}' \mid f')\}.$$

Let f be function maximizing the first expression. Then this is equal to

$$\begin{aligned}\Delta_{A,b}(\underline{R} \mid f) - \max_{f':[L]\to\{\pm1\}}\{\Delta_{A,b}(\underline{R}' \mid f')\} &\le \Delta_{A,b}(\underline{R} \mid f) - \Delta_{A,b}(\underline{R}' \mid f) \\ &\le \frac{2}{\sqrt{KN}} \cdot \|\underline{R} - \underline{R}'\|_2. \qquad \text{(by Lemma 3.25)}\end{aligned}$$

This completes the proof. □

With this in hand, we can finally prove Lemma 3.23.

Proof of Lemma 3.23. By Lemma 3.26, $\Delta_{A,b}(\cdot)$ is convex and has Lipschitz constant $(2/\sqrt{KN})$. Let $\boldsymbol{R} : [k] \times [N] \to \{\pm1\}$ be a uniformly random function family. Viewing $\boldsymbol{R}$ as a collection of KN independent $\{\pm1\}$ random variables, we can apply Talagrand's concentration inequality, which states that there exists an absolute constant $c > 0$ such that

$$\Pr_{\boldsymbol{R}}[|\Delta_{A,b}(\boldsymbol{R}) - \mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_{A,b}(\boldsymbol{R})]| \ge \varepsilon] \le 2 \cdot \exp\Big(-\frac{c \cdot \varepsilon^2}{(2/\sqrt{KN})^2}\Big) = 2 \cdot \exp\Big(-\Big(\frac{c}{4}\Big) \cdot \varepsilon^2 KN\Big).$$

Recalling that $\Delta^{\mathrm{avg}}_{A,b} = \mathbf{E}_{\boldsymbol{R}}[\Delta_{A,b}(\boldsymbol{R})]$, this completes the proof. □

3.5 The adversary's space is bounded without loss of generality

In this subsection, we will show that if $A^{(\cdot)}$ is an oracle circuit that makes t queries, each of which has size at most ℓ, then we can assume without loss of generality that $A^{(\cdot)}$ uses at most $t \cdot \ell$ ancilla qubits, in addition to the n qubits in its input register. We prove this by showing that for any such oracle circuit (that potentially uses unbounded space), there is an oracle circuit $B^{(\cdot)}$ that *simulates* $A^{(\cdot)}$ using only $t \cdot \ell$ ancilla qubits. This will allow us to restrict our attention to adversaries that are space-efficient when proving our one-query lower bounds, which is necessary given the technical tools we apply. We begin by defining in what sense $B^{(\cdot)}$ simulates $A^{(\cdot)}$.

Notation 3.27 (Query register)**.** In this subsection, we will assume that every oracle circuit makes an oracle call on a register of exactly ℓ qubits. We will write $L = 2^\ell$ for the dimension of this register, and we will write $\mathcal{H}_{\mathsf{query}} \coloneqq \mathbb{C}^L$ for the vector space corresponding to this register.

Definition 3.28 (Oracle circuit simulation)**.** Let $m_B \le m_A$ be integers. Consider two t-query oracle circuits

$$\begin{aligned}A^{(\cdot)} &= (n, m_A = \log_2(M_A), \ell, U^A_1, \ldots, U^A_{t+1}), \text{ and} \\ B^{(\cdot)} &= (n, m_B = \log_2(M_B), \ell, U^B_1, \ldots, U^B_{t+1})\end{aligned}$$

with ancilla dimensions $D_A \coloneqq 2^{m_A - \ell}$ and $D_B \coloneqq 2^{m_B - \ell}$, respectively. Then $B^{(\cdot)}$ *simulates* $A^{(\cdot)}$ if there exists an isometry $T : \mathbb{C}^{M_B} \to \mathbb{C}^{M_A}$ such that for all Boolean functions $f : [L] \to \{\pm1\}$,

$$\begin{aligned}&U^A_{t+1} \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U^A_t \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U^A_1 \cdot (\mathsf{Id}_N \otimes |0^{m_A - n}\rangle) \\ = T \cdot &U^B_{t+1} \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{D_B}) \cdot U^B_t \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_B}) \cdot U^B_1 \cdot (\mathsf{Id}_N \otimes |0^{m_B - n}\rangle).\end{aligned}$$

Now we state the main lemma of this section, namely that an oracle circuit that makes t queries of size ℓ can be converted to one of space $n + t \cdot \ell$. Typical values for these parameters are $t, \ell = \mathrm{poly}(n)$, in which case this results in an oracle circuit of $\mathrm{poly}(n)$ space.

Lemma 3.29 (Space reduction for oracle circuits)**.** *Consider a t-query oracle circuit*

$$A^{(\cdot)} = (n, m_A, \ell, U_1^A, \ldots, U_{t+1}^A).$$

Then $A^{(\cdot)}$ can be simulated by a t-query oracle circuit

$$B^{(\cdot)} = (n, m_B, \ell, U_1^B, \ldots, U_{t+1}^B).$$

that uses $m_B = (n + t \cdot \ell)$ qubits of space.

The key technical ingredient we will use in the proof of this lemma is the following method for compressing an isometry with a large output dimension into an isometry with a small output dimension.

Definition 3.30 (Compression of an isometry)**.** Let $V : \mathbb{C}^D \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^S$ be an isometry. Then the *compression of V* is the isometry $\mathsf{compress}(V) : \mathbb{C}^D \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^D$ defined as

$$\mathsf{compress}(V) \coloneqq \sum_{z=1}^{L} |z\rangle \otimes \sqrt{M_z},$$

where $M_z \coloneqq V^\dagger \cdot (|z\rangle\langle z| \otimes \mathsf{Id}_S) \cdot V$.

To get intuition for this definition, note that the operators $\{M_z\}$ correspond to the following measurement: first apply the original isometry V, and then measure the resulting query register to obtain an outcome z. As a result, $\mathsf{compress}(V)$ is the natural isometry that corresponds to the $\{M_z\}$ measurement. We note that $\mathsf{compress}(V)$ *is* indeed an isometry, because

$$\begin{aligned}\mathsf{compress}(V)^\dagger \cdot \mathsf{compress}(V) &= \Big(\sum_{z=1}^{L} \langle z| \otimes \sqrt{M_z}\Big) \cdot \Big(\sum_{z=1}^{L} |z\rangle \otimes \sqrt{M_z}\Big) \\ &= \sum_{z=1}^{L} M_z = \sum_{z=1}^{L} V^\dagger \cdot (|z\rangle\langle z| \otimes \mathsf{Id}_S) \cdot V = V^\dagger \cdot \mathsf{Id}_{\mathrm{query}} \otimes \mathsf{Id}_S \cdot V = \mathsf{Id}_B,\end{aligned}$$

where the last step used the assumption that V is an isometry. The following technical lemma gives one sense in which $\mathsf{compress}(V)$ does indeed compress V, in that whenever V is used to temporarily transition into $\mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^S$ in order to query an oracle, we can use $\mathsf{compress}(V)$ to move into $\mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^D$ instead with the exact same results.

Lemma 3.31 ($\mathsf{compress}$ compresses)**.** *Let $V : \mathbb{C}^D \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^S$ be an isometry. Then for every function $f : [L] \to \{\pm 1\}$,*

$$\mathsf{compress}(V)^\dagger \cdot \Big(\mathcal{O}_f \otimes \mathsf{Id}_D\Big) \cdot \mathsf{compress}(V) = V^\dagger \cdot \Big(\mathcal{O}_f \otimes \mathsf{Id}_S\Big) \cdot V.$$

Proof. The proof is via a straightforward calculation:

$$\begin{aligned}
&\mathsf{compress}(V)^\dagger \cdot \Big(\mathcal{O}_f \otimes \mathsf{Id}_D\Big) \cdot \mathsf{compress}(V) \\
&= \Big(\sum_{z=1}^{L} \langle z| \otimes \sqrt{M_z}\Big) \cdot \Big(\sum_{z=1}^{L} f(z) \cdot |z\rangle\langle z| \otimes \mathsf{Id}_D\Big) \cdot \Big(\sum_{z=1}^{L} |z\rangle \otimes \sqrt{M_z}\Big) \\
&= \sum_{z=1}^{L} f(z) \cdot M_z \\
&= \sum_{z=1}^{L} f(z) \cdot V^\dagger \cdot (|z\rangle\langle z| \otimes \mathsf{Id}_S) \cdot V \\
&= V^\dagger \cdot \Big(\sum_{z=1}^{L} f(z) \cdot |z\rangle\langle z| \otimes \mathsf{Id}_S\Big) \cdot V \\
&= V^\dagger \cdot \Big(\mathcal{O}_f \otimes \mathsf{Id}_S\Big) \cdot V.
\end{aligned}$$

That completes the proof. □

Now we use this technical lemma to show that the action of $\mathsf{compress}(V)$ followed by an oracle is actually *equivalent* to the action of V followed by an oracle, up to an isometry.

Lemma 3.32 (Equivalence of compressed and uncompressed isometry)**.** *Let $V : \mathbb{C}^D \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^S$ be an isometry. Then there exists an isometry $T : \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^D \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^S$ such that for all Boolean functions $f : [L] \to \{\pm 1\}$,*

$$T \cdot (\mathcal{O}_f \otimes \mathsf{Id}_D) \cdot \mathsf{compress}(V) = (\mathcal{O}_f \otimes \mathsf{Id}_S) \cdot V.$$

Prior to proving this lemma, we will establish the following linear-algebraic proposition. We expect that this proposition is well-known, although we were unable to find a reference for it.

Proposition 3.33 (Matching inner products implies an isometry)**.** *Let $d_1 \leq d_2$ be integers. Consider two sets of m vectors $|x_1\rangle, \ldots, |x_m\rangle \in \mathbb{C}^{d_1}$ and $|y_1\rangle, \ldots, |y_m\rangle \in \mathbb{C}^{d_2}$. Suppose that these sets have the same pairwise inner products, meaning that*

$$\langle x_i | x_j \rangle = \langle y_i | y_j \rangle,$$

for all $1 \leq i, j \leq m$. Then there exists an isometry $T : \mathbb{C}^{d_1} \to \mathbb{C}^{d_2}$ such that $T \cdot |x_i\rangle = |y_i\rangle$, for all $1 \leq i \leq m$.

Proof. Define $X \coloneqq \sum_{i=1}^m |x_i\rangle\langle i|$ and $Y = \sum_{i=1}^m |y_i\rangle\langle i|$. Because the two sets of vectors have matching inner products,

$$\begin{aligned}
X^\dagger \cdot X = \Big(\sum_{i=1}^{m} |i\rangle\langle x_i|\Big) \cdot \Big(\sum_{j=1}^{m} |x_j\rangle\langle j|\Big) &= \sum_{i,j=1}^{m} \langle x_i | x_j \rangle \cdot |i\rangle\langle j| \qquad (15) \\
&= \sum_{i,j=1}^{m} \langle y_i | y_j \rangle \cdot |i\rangle\langle j| = \Big(\sum_{i=1}^{m} |i\rangle\langle y_i|\Big) \cdot \Big(\sum_{j=1}^{m} |y_j\rangle\langle j|\Big) = Y^\dagger \cdot Y.
\end{aligned}$$

Given a complex matrix A, we will denote by A^+ the *Moore-Penrose pseudo-inverse* of A. The one fact we will use about the pseudo-inverse, which can be found in [Pet12, Proposition 4.9.2], is that

$A^+ \cdot A$ is the projector onto the image of $A^\dagger$. Multiplying both sides of Equation (15) by $(Y^\dagger)^+$ yields

$$(Y^\dagger)^+ \cdot X^\dagger \cdot X = (Y^\dagger)^+ \cdot Y^\dagger \cdot Y.$$

From our pseudo-inverse fact, $(Y^\dagger)^+ \cdot Y^\dagger$ is the projector onto the image of $(Y^\dagger)^\dagger = Y$. Hence, $(Y^\dagger)^+ \cdot Y^\dagger \cdot Y = Y$.

$$(Y^\dagger)^+ \cdot X^\dagger \cdot X = Y.$$

Now, let us define $T \coloneqq (Y^\dagger)^+ \cdot X^\dagger$, so that $T \cdot X = Y$. Note that for all $1 \leq i \leq m$, this implies that

$$T \cdot |x_i\rangle = T \cdot X \cdot |i\rangle = Y \cdot |i\rangle = |y_i\rangle\,, \tag{16}$$

as desired. Next, write

$$\mathsf{span}_X \coloneqq \mathsf{span}\{|x_1\rangle\,, \ldots, |x_m\rangle\} \quad \text{and} \quad \mathsf{span}_Y \coloneqq \mathsf{span}\{|y_1\rangle\,, \ldots, |y_m\rangle\}.$$

Then we claim (i) T maps any vector in $\mathsf{span}_X^\perp$ to 0, and (ii) T is an isometry from span_X to span_Y. We prove these as follows.

(i) Let $|u\rangle \in \mathsf{span}_X^\perp$. Because $X^\dagger = \sum_{i=1}^m |i\rangle\langle x_i|$, we have that $X^\dagger \cdot |u\rangle = 0$. Thus,

$$T \cdot |u\rangle = (Y^\dagger)^+ \cdot X^\dagger \cdot |u\rangle = 0.$$

(ii) To show that T is an isometry mapping span_X to span_Y, it suffices to show that it maps any vector in span_X to span_Y, and that it preserves lengths. Let $|v\rangle \in \mathsf{span}_X$. Then $|v\rangle = \alpha_1 \cdot |x_1\rangle + \cdots + \alpha_m \cdot |x_m\rangle$ for some complex coefficients $\alpha_1, \ldots, \alpha_m$. By Equation (16),

$$T \cdot |v\rangle = T \cdot \Big(\sum_{i=1}^m \alpha_i \cdot |x_i\rangle\Big) = \sum_{i=1}^m \alpha_i \cdot T \cdot |x_i\rangle = \sum_{i=1}^m \alpha_i \cdot |y_i\rangle\,,$$

which is indeed an element of span_Y. Next, the squared length of $|v\rangle$ is

$$\begin{aligned}
\langle v|v\rangle &= \Big(\sum_{i=1}^m \alpha_i^\dagger \cdot \langle x_i|\Big) \cdot \Big(\sum_{j=1}^m \alpha_j \cdot |x_j\rangle\Big) \\
&= \sum_{i,j=1}^m \alpha_i^\dagger \alpha_j \cdot \langle x_i|x_j\rangle \\
&= \sum_{i,j=1}^m \alpha_i^\dagger \alpha_j \cdot \langle y_i|y_j\rangle \\
&= \Big(\sum_{i=1}^m \alpha_i^\dagger \cdot \langle y_i|\Big) \cdot \Big(\sum_{j=1}^m \alpha_j^\dagger \cdot |y_j\rangle\Big) = (\langle v| \cdot T^\dagger) \cdot (T \cdot |v\rangle),
\end{aligned}$$

which is the squared length of $T \cdot |v\rangle$. This proves the claim.

Hence, T is an isometry mapping span_X to span_Y, acts as 0 outside of span_X, and satisfies $T \cdot |x_i\rangle = |y_i\rangle$, for all $1 \leq i \leq m$. As a result, it can be extended to an isometry mapping $\mathbb{C}^{d_1}$ to $\mathbb{C}^{d_2}$ which satisfies this property by picking any isometry that maps $\mathsf{span}_X^\perp$ to $\mathsf{span}_Y^\perp$. This gives the desired construction. □

Now we prove Lemma 3.32.

Proof of Lemma 3.32. For each $f : [L] \to \{\pm 1\}$ and $1 \le x \le D$, define

$$|\Phi_{f,x}\rangle := (\mathcal{O}_f \otimes \mathsf{Id}_S) \cdot V \cdot |x\rangle\,, \text{ and } |\widehat{\Phi}_{f,x}\rangle := (\mathcal{O}_f \otimes \mathsf{Id}_D) \cdot \mathsf{compress}(V) \cdot |x\rangle\,.$$

We will prove that there exists an isometry $T : \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^D \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^S$ such that

$$T \cdot |\widehat{\Phi}_{f,z}\rangle = |\Phi_{f,x}\rangle\,,$$

for all $1 \le x \le D$ and Boolean functions $f : [L] \to \{\pm 1\}$. This will in turn imply the desired claim by linearity. By Proposition 3.33, it suffices to show that $\{|\Phi_{f,x}\rangle\}_{f,x}$ and $\{|\widehat{\Phi}_{f,x}\rangle\}_{f,x}$ have the same pairwise inner products, i.e.

$$\langle\widehat{\Phi}_{f,x}|\widehat{\Phi}_{g,y}\rangle = \langle\Phi_{f,x}|\Phi_{g,y}\rangle\,, \tag{17}$$

for all $1 \le x, y \le D$ and Boolean functions $f, g : [L] \to \{\pm 1\}$. To complete the proof, we verify this by direct calculation for all x, y, f, g:

$$\begin{aligned}
\langle\widehat{\Phi}_{f,x}|\widehat{\Phi}_{g,y}\rangle &= \langle x| \cdot \mathsf{compress}(V)^\dagger \cdot \Big(\mathcal{O}_f \otimes \mathsf{Id}_D\Big) \cdot \Big(\mathcal{O}_g \otimes \mathsf{Id}_D\Big) \cdot \mathsf{compress}(V) \cdot |y\rangle \\
&= \langle x| \cdot \mathsf{compress}(V)^\dagger \cdot \Big(\mathcal{O}_{f\cdot g} \otimes \mathsf{Id}_D\Big) \cdot \mathsf{compress}(V) \cdot |y\rangle \\
&= \langle x| \cdot V^\dagger \cdot \Big(\mathcal{O}_{f\cdot g} \otimes \mathsf{Id}_S\Big) \cdot V \cdot |y\rangle && \text{(by Lemma 3.31)} \\
&= \langle x| \cdot V^\dagger \cdot \Big(\mathcal{O}_f \otimes \mathsf{Id}_S\Big) \cdot \Big(\mathcal{O}_g \otimes \mathsf{Id}_S\Big) \cdot V \cdot |y\rangle \\
&= \langle\Phi_{f,x}|\Phi_{g,y}\rangle\,,
\end{aligned}$$

where $f \cdot g$ is the Boolean function defined as $(f \cdot g)(x) = f(x) \cdot g(x)$, for all $1 \le x \le D$. This completes the proof. □

With this in hand, we can finally prove the main result of this section, Lemma 3.29.

Proof of Lemma 3.29. In this proof, we will construct a sequence of isometries $V_1, \ldots, V_t$ in which for each $1 \le i \le t$,

$$V_i : \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i-1} \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i-1} = \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i}. \tag{18}$$

Given a Boolean function $f : [L] \to \{\pm 1\}$, we will use the shorthand

$$\begin{aligned}
\mathsf{Prod}_{U,f,i} &:= (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_i^A \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_2^A \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_1^A \cdot (\mathsf{Id}_N \otimes |0^{m_A - n}\rangle), \\
\mathsf{Prod}_{V,f,i} &:= (\mathcal{O}_f \otimes \mathsf{Id}_{NL^{i-1}}) \cdot V_i \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{NL}) \cdot V_2 \cdot (\mathcal{O}_f \otimes \mathsf{Id}_N) \cdot V_1 \cdot \mathsf{Id}_N.
\end{aligned}$$

Operationally, $\mathsf{Prod}_{U,f,i}$ corresponds to alternating between i unitaries and oracle calls, and similarly for $\mathsf{Prod}_{V,f,i}$.

We will first prove the following statement: for each $0 \le i \le t$, there exists an isometry $T_i : \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i-1} \to \mathbb{C}^{M_A}$, such that for every Boolean function $f : [L] \to \{\pm 1\}$,

$$T_i \cdot \mathsf{Prod}_{V,f,i} = \mathsf{Prod}_{U,f,i}. \tag{19}$$

At the end, we will derive Lemma 3.29 from this statement.

The proof is by induction on t, the base case being $t = 0$. In this case, the statement follows from setting $T_0 : \mathbb{C}^N \to \mathbb{C}^{M_A}$ as $T_0 := \mathsf{Id}_N \otimes |0^{m_A - n}\rangle$. This is because

$$T_0 \cdot \mathsf{Prod}_{V,f,0} = T_0 \cdot \mathsf{Id}_N = \mathsf{Id}_N \otimes |0^{m_A - n}\rangle = \mathsf{Prod}_{U,f,0},$$

as desired. As for the induction step we suppose it is true for $i \le t-1$ and prove that it holds for $i+1$. By the induction hypothesis, we have that

$$\mathsf{Prod}_{U,f,i+1} = (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_{i+1}^A \cdot \mathsf{Prod}_{U,f,i} = (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_{i+1}^A \cdot T_i \cdot \mathsf{Prod}_{V,f,i}. \tag{20}$$

Note that $U_{i+1}^A \cdot T_i$ is an isometry mapping $\mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i}$ to $\mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^{D_A}$. Thus, if we set $V_{i+1} \coloneqq \mathsf{compress}(U_{i+1}^A \cdot T_i)$, then

$$V_{i+1} : \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i} \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i},$$

as desired. Applying Lemma 3.32, there exists an isometry

$$T_{i+1} : \mathcal{H}_{\mathsf{query}} \otimes C^N \otimes (\mathbb{C}^L)^{\otimes i} \to \mathcal{H}_{\mathsf{query}} \otimes \mathbb{C}^{D_A}$$

such that for all Boolean functions $f : [L] \to \{\pm 1\}$,

$$T_{i+1} \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{NL^i}) \cdot V_{i+1} = (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_{i+1}^A \cdot T_i.$$

Plugging this into Equation (20), we have that

$$\mathsf{Prod}_{U,f,i+1} = T_{i+1} \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{NL^i}) \cdot V_{i+1} \cdot \mathsf{Prod}_{V,f,i} = T_{i+1} \cdot \mathsf{Prod}_{V,f,i+1}.$$

Thus, the $(i+1)$ case of the statement is also true, completing the proof by induction.

It remains to show that the existence of isometries $V_1, \ldots, V_t$ and $T_0, \ldots, T_t$ satisfying Eqs. (18) and (19) implies Lemma 3.29. Recall that our goal is to construct an oracle circuit $B^{(\cdot)} = (n, m_B, \ell, U_1^B, \ldots, U_{t+1}^B)$ that uses $m_B = (n + t \cdot \ell)$ qubits of space and *simulates* $A^{(\cdot)}$ in the sense of Definition 3.28: namely, there exists an isometry $T : \mathbb{C}^{M_B} \to \mathbb{C}^{M_A}$ such that for all Boolean functions $f : [L] \to \{\pm 1\}$,

$$\begin{aligned} &U_{t+1}^A \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_t^A \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_1^A \cdot (\mathsf{Id}_N \otimes |0^{m_A - n}\rangle) \\ &= T \cdot U_{t+1}^B \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{D_B}) \cdot U_t^B \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_B}) \cdot U_1^B \cdot (\mathsf{Id}_N \otimes |0^{m_B - n}\rangle), \end{aligned} \tag{21}$$

where $D_A = 2^{m_A - \ell}$ and $D_B = 2^{m_B - \ell}$. To this end, for $1 \le i \le t$, we will extend each isometry $V_i : \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i-1} \to \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes i}$ to a unitary U_i^B acting on $m_B = n + t \cdot \ell$ qubits as follows. First, for $1 \le i \le t$, define $\widetilde{U}_i^B$ to be an extension of the isometry V_i to a unitary on $n + \ell \cdot i$ qubits, i.e.,

$$\widetilde{U}_i^B \cdot (\mathsf{Id}_{NL^{i-1}} \otimes |0^\ell\rangle) = V_i.$$

We then extend this to a unitary on $n + t \cdot \ell$ qubits by setting $U_i^B \coloneqq \widetilde{U}_i^B \otimes \mathsf{Id}_{L^{t-i}}$. Finally, we pick $U_{t+1}^B \coloneqq \mathsf{Id}_{NL^t}$.

To put everything together, we need to prove the existence of an isometry T satisfying Eq. (21). Plugging in our definitions for U_i^B into Eq. (19), there exists an isometry $T_t : \mathbb{C}^{M_B} = \mathbb{C}^N \otimes (\mathbb{C}^L)^{\otimes t} \to \mathbb{C}^{M_A}$ such that for all Boolean functions $f : [L] \to \{\pm 1\}$,

$$\begin{aligned} &(\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_t^A \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_A}) \cdot U_1^A \cdot (\mathsf{Id}_N \otimes |0^{m_A - n}\rangle) && (22) \\ &= T_t \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{NL^{i-1}}) \cdot V_t \cdots (\mathcal{O}_f \otimes \mathsf{Id}_N) \cdot V_1 \cdot \mathsf{Id}_N \\ &= T_t \cdot U_{t+1}^B \cdot (\mathcal{O}_f \otimes \mathsf{Id}_{D_B}) \cdot U_t^B \cdots (\mathcal{O}_f \otimes \mathsf{Id}_{D_B}) \cdot U_1^B \cdot (\mathsf{Id}_N \otimes |0^{m_B - n}\rangle). && (23) \end{aligned}$$

The equation (22) = (23) shows that T_t *almost* satisfies the desired properties of the isometry $T : \mathbb{C}^{M_B} \to \mathbb{C}^{M_A}$ that we want (Eq. (21)), except that in Eq. (21), there is an additional U_{t+1}^A unitary applied on the left-hand side. To complete the proof, we set $T = U_{t+1}^A \cdot T_t$. □

Simulating the measurement. Our notion of what it means for $B^{(\cdot)}$ to *simulate* $A^{(\cdot)}$ only guarantees that there exists an isometry T such that (on any input) running $B^{(\cdot)}$ and then applying T produces the same state as running $A^{(\cdot)}$. However, our aim is to use $B^{(\cdot)}$ in place of $A^{(\cdot)}$ as an adversary in the Oracle State Distinguishing Game; recall that an adversary (Definition 3.10) in this game first applies the oracle circuit on a given input state, and then *measures the first qubit* of the resulting state to produce a guess bit $\boldsymbol{b}'$. Thus, what we need is a way to run a low-space oracle circuit $B^{(\cdot)}$ so that when we measure the first qubit of the resulting state, the outcome distribution is the same as if we had run $A^{(\cdot)}$ and measured its first qubit.

Fortunately, we can resolve this issue with standard techniques from quantum information (namely Naimark dilation; see, e.g., page 94 of [NC10]). First, define the m_B-qubit binary-outcome POVM $\{E_0, E_1 = \mathsf{Id} - E_0\}$ where

$$E_0 \coloneqq T^\dagger \cdot (|0\rangle\langle 0| \otimes \mathsf{Id}_2^{\otimes m_A - 1}) \cdot T.$$

Now, observe that if we run $B^{(\cdot)}$ and then measure $\{E_0, E_1 = \mathsf{Id} - E_0\}$, the resulting outcome $\boldsymbol{b}'$ is distributed exactly the same as it would be if we had instead run $B^{(\cdot)}$, then applied T, and measured the first qubit (and the latter is equivalent to running $A^{(\cdot)}$ and measuring the first qubit).

To implement this POVM as a measurement of the *first qubit* of the adversary's state, we will define the isometry $V_{\text{guess}} : \mathbb{C}^{M_B} \to \mathbb{C}^2 \otimes \mathbb{C}^{M_B}$ as

$$V_{\text{guess}} \coloneqq \sum_{b \in \{0,1\}} |b\rangle \otimes \sqrt{E_b}.$$

We note that V_{guess} is in fact an isometry, since

$$V_{\text{guess}}^\dagger \cdot V_{\text{guess}} = \Big(\sum_{b \in \{0,1\}} \langle b| \otimes \sqrt{E_b} \Big) \cdot \Big(\sum_{b \in \{0,1\}} |b\rangle \otimes \sqrt{E_b} \Big) = \sum_{b \in \{0,1\}} E_b = \mathsf{Id}_{M_B}.$$

Moreover, applying V_{guess} and measuring the first qubit of the resulting state produces the same distribution as measuring $\{E_0, E_1\}$, since for any state $|\psi\rangle \in \mathbb{C}^{M_B}$ and any $b \in \{0,1\}$,

$$\langle\psi| \cdot V_{\text{guess}}^\dagger \cdot |b\rangle\langle b| \cdot V_{\text{guess}} \cdot |\psi\rangle = \langle\psi| \cdot \sqrt{E_b} \cdot \sqrt{E_b} \cdot |\psi\rangle = \langle\psi| E_b |\psi\rangle .$$

Thus, given any circuit $B^{(\cdot)} = (n, m_B, \ell, U_1^B, \ldots, U_{t+1}^B)$ that simulates $A^{(\cdot)}$ in the sense of Definition 3.28, we can easily modify $B^{(\cdot)}$ to obtain another low-space t-query oracle circuit $C^{(\cdot)} = (n, m_B + 1, \ell, U_1^C, \ldots, U_{t+1}^C)$ that has the additional guarantee that running $C^{(\cdot)}$ and measuring its first qubit produces a guess from the correct output distribution.

Concretely, define U_{guess} to be an $(m_B + 1)$-qubit unitary that extends the isometry V_{guess} in the sense that $U_{\text{guess}} \cdot (|0\rangle \otimes \mathsf{Id}_{M_B}) = V_{\text{guess}}$. Then define

$$U_{t+1}^C \coloneqq U_{\text{guess}} \cdot (\mathsf{Id}_2 \otimes U_{t+1}^B).$$

The unitaries corresponding to $1 \le i \le t$ are defined as they are in $B^{(\cdot)}$ except that they act on one additional qubit, i.e., $U_i^C \coloneqq \mathsf{Id}_2 \otimes U_i^B$. By the preceding discussion, these definitions guarantee that running $C^{(\cdot)}$ and measuring its first qubit yields the outcome distribution of the original oracle adversary $A^{(\cdot)}$. Thus, we have the following corollary of Lemma 3.29.

Corollary 3.34. *Without loss of generality, any t-query adversary in the Oracle State Distinguishing Game uses an oracle circuit $(n, m, \ell, U_1, \ldots, U_t, U_{t+1})$ that requires $m \le n + t \cdot \ell + 1$ qubits of space.*

3.6 One-query adversary model, final problem setup

In this section, we give a "normal form" for one-query adversaries $A^{(\cdot)}$ with bounded query length. By Corollary 3.34, for any $A^{(\cdot)}$ with query length bounded by ℓ, we may assume without loss of generality that $A^{(\cdot)}$ uses at most $a \leq \ell + 1$ ancilla qubits, for a total number of $m = n + a$ qubits. As a result, following Definition 3.4, we may assume that A^f operates as follows, for some choice of unitaries U_1, U_2:

1. Given an n qubit input $|\psi\rangle$, compute the state

$$U_2 \cdot \mathcal{O}_f \cdot U_1 \cdot |\psi\rangle \ |0^a\rangle$$

2. Measure the first qubit of the resulting state in the standard basis.

In this special case, we simplify our notation slightly with the following definitions:

- Let $M := 2^m$ denote the dimension of the adversary's final Hilbert space.
- Let $V = U_1 \cdot (\mathsf{Id}_N \otimes |0^a\rangle)$ denote the isometry describing A's behavior prior to the query.
- Let $\Pi = U_2^\dagger \cdot (|0\rangle\langle 0| \otimes \mathsf{Id}) \cdot U_2$ be the projection describing the adversary's measurement applied to $\mathcal{O}_f \cdot V \cdot |\psi\rangle$.

To summarize, we have modeled the adversary as $A^{(\cdot)} = (M, V, \Pi)$, where M is an integer, $V : \mathbb{C}^N \to \mathbb{C}^M$ is an isometry, and $\Pi \in \mathbb{C}^{M \times M}$ is a projection. In this language, the adversary's probability of outputting "0" on a binary phase state $|\psi_h\rangle$ is given by

$$p_A(h \mid f) = \langle \psi_h | \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_h\rangle \,.$$

Our goal in Section 4 will be to prove an upper bound on $\Delta_A(\boldsymbol{R})$ (see Definition 3.12) of the form

$$\Delta_A(\boldsymbol{R}) = O\left(\sqrt{\frac{\log M}{K}}\right)$$

with high probability over $\boldsymbol{R}$, which will establish a lower bound of $m = \log(M) = \Omega(K\varepsilon^2)$ for adversaries with distinguishing advantage ε.

4 Proof of the one-query lower bound

In this section, we will consider a single-query adversary A and show that its advantage in the oracle state distinguishing game $\Delta_A(\boldsymbol{R})$ is very small with overwhelming probability over $\boldsymbol{R}$. By Lemma 3.18, it suffices to bound the expectation $\mathbf{E}[\Delta_A(\boldsymbol{R})] = \Delta_A^{\mathrm{avg}}$ for every A.

We will bound this expected value as follows: in Section 4.1, we will apply a standard decoupling trick to the expression for the adversary's distinguishing advantage. Next, in Section 4.2 we will develop a natural spectral relaxation of this decoupled distinguishing advantage. Following that, in Section 4.3 we will use a matrix concentration inequality to bound the expectation of the spectral relaxation in terms of a quantity that we call the width of a collection of binary phase states. Then, in Section 4.4, we show how to bound the expected width of a random family of binary phase states. Finally, in Section 4.5 we will combine these ingredients and complete the proof of the one-query lower bound.

Notation. We will first fix some notation to use throughout the section. By Section 3.6, we can model the adversary as $A = (M, V, \Pi)$, where M is an integer (a typical value of which is $M = 2^{\mathrm{poly}(n)}$), $V : \mathbb{C}^N \to \mathbb{C}^M$ is an isometry, and $\Pi \in \mathbb{C}^{M\times M}$ is a projection. Writing $v_{i,x}$ for the (i,x)-th entry of V, we can express it as

$$V = \sum_{i=1}^{M}\sum_{x=1}^{N} v_{i,x} \cdot |i\rangle\langle x| = \sum_{i=1}^{M} |i\rangle \cdot \Big(\sum_{x=1}^{N} v_{i,x} \langle x|\Big) = \sum_{i=1}^{M} |i\rangle\langle v_i|\,,$$

where $|v_i\rangle \in \mathbb{C}^N$ is the vector

$$|v_i\rangle = \sum_{x=1}^{N} v_{i,x}^{\dagger} |x\rangle\,.$$

Note that

$$\sum_{i=1}^{M} \langle v_i|v_i\rangle = \mathrm{tr}\Big(V^{\dagger}V\Big) = \mathrm{tr}(\mathsf{Id}_{N\times N}) = N.$$

This motivates the following definition.

Definition 4.1 (Isometry weights)**.** The *isometry weights* are the numbers

$$\mathsf{wt}_{V,i} \coloneqq \tfrac{1}{N} \cdot \langle v_i|v_i\rangle\,,$$

for $1 \le i \le M$. Note that these sum to one and therefore form a probability distribution.[4]

4.1 Decoupling the quadratic form

Our overall goal is to bound the adversary's distinguishing advantage. Writing $\boldsymbol{h} : [N] \to \{\pm 1\}$ for a uniformly random Boolean function, the distinguishing advantage can be written as

$$\begin{aligned}
&\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{R}}\Big[\max_{f:[M]\to\{\pm1\}}\Big|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R}_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_A(\boldsymbol{h} \mid f)]\Big|\Big] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{R}}\Big[\max_{f}\Big|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[\,\langle\psi_{\boldsymbol{R}_{\boldsymbol{k}}}| \cdot V^{\dagger}\cdot\mathcal{O}_f\cdot\Pi\cdot\mathcal{O}_f\cdot V\cdot|\psi_{\boldsymbol{R}_{\boldsymbol{k}}}\rangle] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[\,\langle\psi_{\boldsymbol{h}}| \cdot V^{\dagger}\cdot\mathcal{O}_f\cdot\Pi\cdot\mathcal{O}_f\cdot V\cdot|\psi_{\boldsymbol{h}}\rangle]\Big|\Big].
\end{aligned} \tag{24}$$

[4] It turns out that $\mathsf{wt}_{V,i}$ is the probability that a standard basis measurement of $V|\psi_{\boldsymbol{h}}\rangle$, for uniformly random $\boldsymbol{h}$, results in an outcome of i.

The coefficients of the vector $|\psi_{\boldsymbol{R}_k}\rangle$ are independent $\{\pm 1\}$ Rademacher random variables, and indeed there are tools from random matrix theory which allow us to prove concentration bounds on matrices whose entries are linear combinations of Rademachers. However, the first term in Equation (24) is quadratic in the $|\psi_{\boldsymbol{R}_k}\rangle$ vector, and so these tools cannot be immediately applied. What we would like to do is *decouple* the left random vector $\langle\psi_{\boldsymbol{R}_k}|$ from the right random vector $|\psi_{\boldsymbol{R}_k}\rangle$ so that this expression becomes a function of two independent random vectors, and is linear in both of them, rather than being quadratic in a single random vector. This motivates the following definition, a natural decoupled analogue of the distinguishing advantage.

Definition 4.2 (The decoupled distinguishing advantage)**.** Let $R, R' : [K] \times [N] \to \{\pm 1\}$ be two function families. Let $f : [M] \to \{\pm 1\}$ be a function, and let A denote a one-query adversary. Then the corresponding *decoupling distinguishing advantage* is given by

$$\Delta_{\mathsf{Decup}}(R, R' \mid f) \coloneqq \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [\, \langle\psi_{R_{\boldsymbol{k}}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{R'_{\boldsymbol{k}}}\rangle] \Big|.$$

Unlike the normal distinguishing advantage, the decoupled distinguishing advantage has no natural operational interpretation. However, it still gives a convenient upper bound to the normal distinguishing advantage, as shown in the following lemma.

Lemma 4.3. *Let $\boldsymbol{R}, \boldsymbol{R}' : [K] \times [N] \to \{\pm 1\}$ be two independent and uniformly random function families. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})] \leq 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'} \max_f \Big\{ \Delta_{\mathsf{Decup}}(\boldsymbol{R}, \boldsymbol{R}' \mid f) \Big\}.$$

Decoupling inequalities are standard in the random matrix theory literature. Our proof follows an outline similar to other decoupling arguments, for example those in [vH17, Lemma 5.2] and [Ver11].

Proof of Lemma 4.3. Throughout this proof, we will adopt the following shorthand for convenience: given an oracle $\mathcal{O}$ acting on $\mathbb{C}^M$, we will write

$$W_{\mathcal{O}} \coloneqq V^\dagger \cdot \mathcal{O} \cdot \Pi \cdot \mathcal{O} \cdot V.$$

For example, if $h : [N] \to \{\pm 1\}$ is a Boolean function, and $R : [K] \times [N] \to \{\pm 1\}$ is a function family, then

$$\begin{aligned} p_A(h \mid f) &= \langle\psi_h| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_h\rangle \\ &= \langle\psi_h| \cdot W_{\mathcal{O}_f} \cdot |\psi_h\rangle \,. \end{aligned}$$

Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Let $\boldsymbol{R}'$ be an independent copy of $\boldsymbol{R}$. Then for each $1 \leq k \leq K$, $\boldsymbol{R}'_k$ is distributed as a uniformly random function, even

conditioned on $\boldsymbol{R}$. As a result, the average distinguishing advantage is given by

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})] &= \mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R_k} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_A(\boldsymbol{h} \mid f)]\right|\right] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R_k} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{R}'}\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R'_k} \mid f)]\right|\right] \\
&\le \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\left[\max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R_k} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R'_k} \mid f)]\right|\right] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\left[\max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(\boldsymbol{R_k} \mid f) - p_A(\boldsymbol{R'_k} \mid f)]\right|\right] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\left[\max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[\langle\psi_{\boldsymbol{R_k}}| \cdot W_{\mathcal{O}_f} \cdot |\psi_{\boldsymbol{R_k}}\rangle - \langle\psi_{\boldsymbol{R'_k}}| \cdot W_{\mathcal{O}_f} \cdot |\psi_{\boldsymbol{R'_k}}\rangle]\right|\right] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\left[\max_f \left| \mathrm{Re}\left(\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[(\langle\psi_{\boldsymbol{R_k}}| + \langle\psi_{\boldsymbol{R'_k}}|) \cdot W_{\mathcal{O}_f} \cdot (|\psi_{\boldsymbol{R_k}}\rangle - |\psi_{\boldsymbol{R'_k}}\rangle)]\right)\right|\right] \\
&\le \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\left[\max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[(\langle\psi_{\boldsymbol{R_k}}| + \langle\psi_{\boldsymbol{R'_k}}|) \cdot W_{\mathcal{O}_f} \cdot (|\psi_{\boldsymbol{R_k}}\rangle - |\psi_{\boldsymbol{R'_k}}\rangle)]\right|\right]. \qquad (25)
\end{aligned}$$

For each $1 \le k \le K$, consider the two vectors

$$|\psi_{\boldsymbol{R}_k}\rangle \pm |\psi_{\boldsymbol{R}'_k}\rangle = \frac{1}{\sqrt{N}} \cdot \sum_{x=1}^{N} (\boldsymbol{R}_k(x) \pm \boldsymbol{R}'_k(x)) \cdot |x\rangle\,.$$

Note that for each $1 \le k \le K$ and $1 \le x \le N$, either $\boldsymbol{R}_k(x)+\boldsymbol{R}'_k(x) \in \{\pm 2\}$ and $\boldsymbol{R}_k(x)-\boldsymbol{R}'_k(x) = 0$, or vice versa. For each $1 \le k \le K$, let $\boldsymbol{S}^+_k : [N] \to \{\pm 1\}$ be a random Boolean function distributed as follows: if $\boldsymbol{R}_k(x) + \boldsymbol{R}'_k(x) \in \{\pm 2\}$, then $\boldsymbol{S}^+_k(x) = \frac{1}{2} \cdot (\boldsymbol{R}_k(x) + \boldsymbol{R}'_k(x))$. Otherwise, when $\boldsymbol{R}_k(x) + \boldsymbol{R}'_k(x) = 0$, choose $\boldsymbol{S}^+_k(x)$ independently and uniformly at random from $\{\pm 1\}$. Define $\boldsymbol{S}^-_k(x)$ similarly. Then the next two equations follow by definition:

$$2 \cdot \mathbf{E}[\,|\psi_{\boldsymbol{S}^+_k}\rangle] = |\psi_{\boldsymbol{R}_k}\rangle + |\psi_{\boldsymbol{R}'_k}\rangle\,, \quad \text{and} \quad 2 \cdot \mathbf{E}[\,|\psi_{\boldsymbol{S}^-_k}\rangle] = |\psi_{\boldsymbol{R}_k}\rangle - |\psi_{\boldsymbol{R}'_k}\rangle\,.$$

Plugging this in above,

$$\begin{aligned}
(25) &= 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'} \max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{S}^+,\boldsymbol{S}^-}\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[\,\langle\psi_{\boldsymbol{S}^+_k}| \cdot W_{\mathcal{O}_f} \cdot |\psi_{\boldsymbol{S}^-_k}\rangle]\right| \\
&\le 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\mathop{\mathbf{E}}_{\boldsymbol{S}^+,\boldsymbol{S}^-} \max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[\,\langle\psi_{\boldsymbol{S}^+_k}| \cdot W_{\mathcal{O}_f} \cdot |\psi_{\boldsymbol{S}^-_k}\rangle]\right|.
\end{aligned}$$

The expression inside the max only depends on $\boldsymbol{R}$ and $\boldsymbol{R}'$ through $\boldsymbol{S}^+$ and $\boldsymbol{S}^-$. Hence, this is equal to

$$4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{S}^+,\boldsymbol{S}^-} \max_f \left| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[\,\langle\psi_{\boldsymbol{S}^+_k}| \cdot W_{\mathcal{O}_f} \cdot |\psi_{\boldsymbol{S}^-_k}\rangle]\right| = 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{S}^+,\boldsymbol{S}^-} \max_f \Big\{\Delta_{\mathsf{Decup}}(\boldsymbol{S}^+,\boldsymbol{S}^- \mid f)\Big\}.$$

But it can be checked that $\boldsymbol{S}^+$ and $\boldsymbol{S}^-$ are just distributed as two independent and uniformly random function families. This completes the proof. □

4.2 A spectral relaxation for the decoupled distinguishing advantage

In this section, we develop a spectral relaxation for the decoupled distinguishing advantage

$$\Delta_{\mathsf{Decup}}(R, R' \mid f) = \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [\, \langle \psi_{R_{\boldsymbol{k}}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{R'_{\boldsymbol{k}}}\rangle] \Big|.$$

As a precursor to this, we will develop a formula for expressing the vectors $V \cdot |\psi_{R_{\boldsymbol{k}}}\rangle$ and $V \cdot |\psi_{R'_{\boldsymbol{k}}}\rangle$.

Let $h : [N] \to \{\pm 1\}$ be a Boolean function, and consider the binary phase state

$$|\psi_h\rangle = \frac{1}{\sqrt{N}} \sum_{x\in[N]} h(x) \, |x\rangle \, .$$

The isometry V maps $|\psi_h\rangle$ to

$$V \cdot |\psi_h\rangle = \Big(\sum_{i=1}^{M} |i\rangle\langle v_i| \Big) \cdot |\psi_h\rangle = \sum_{i=1}^{M} \langle v_i|\psi_h\rangle \cdot |i\rangle \, .$$

Thus, the amplitude on the i-th basis element is $\langle v_i|\psi_h\rangle$. We would like to estimate the magnitude of this amplitude for a "typical" binary phase state. This is given by the following proposition.

Proposition 4.4 (Typical amplitudes)**.** *Let $\boldsymbol{h} : [N] \to \{\pm 1\}$ be a uniformly random Boolean function. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{h}} |\, \langle v_i|\psi_{\boldsymbol{h}}\rangle \,|^2 = \mathsf{wt}_{V,i},$$

where $\mathsf{wt}_{V,i}$ is the isometry weight defined in Definition 4.1.

Proof. We calculate the expectation as follows:

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{h}} |\, \langle v_i|\psi_{\boldsymbol{h}}\rangle \,|^2 &= \mathop{\mathbf{E}}_{\boldsymbol{h}} \langle \psi_{\boldsymbol{h}}|v_i\rangle \cdot \langle v_i|\psi_{\boldsymbol{h}}\rangle \\
&= \mathop{\mathbf{E}}_{\boldsymbol{h}} \Big(\sum_{x=1}^{N} \tfrac{1}{\sqrt{N}} \boldsymbol{h}(x) \cdot v_{i,x}^\dagger \Big) \cdot \Big(\sum_{y=1}^{N} v_{i,y} \cdot \tfrac{1}{\sqrt{N}} \boldsymbol{h}(y) \Big) \\
&= \frac{1}{N} \cdot \sum_{x,y=1}^{N} v_{i,x}^\dagger v_{i,y} \cdot \mathop{\mathbf{E}}_{\boldsymbol{h}} [\boldsymbol{h}(x)\boldsymbol{h}(y)] \\
&= \frac{1}{N} \cdot \sum_{x=1}^{N} |v_{i,x}|^2 = \frac{1}{N} \langle v_i|v_i\rangle \, ,
\end{aligned}$$

where the second-to-last equality used the fact that $\mathbf{E}_{\boldsymbol{h}}[\boldsymbol{h}(x)\boldsymbol{h}(y)] = 1$ if $x = y$ and 0 otherwise, because $\boldsymbol{h}$ is uniformly random. The proof concludes by applying the definition of $\mathsf{wt}_{V,i}$. □

In light of this, it is natural to define the following vector, which contains the "typical" amplitudes of $V \cdot |\psi_h\rangle$.

Definition 4.5 (The weight vector)**.** The *weight vector* is the unit vector given by

$$|\mathsf{wt}_V\rangle = \sum_{i=1}^{M} \sqrt{\mathsf{wt}_{V,i}} \, |i\rangle \, .$$

We can express $V \cdot |\psi_h\rangle$ in terms of the weight vector as

$$V \cdot |\psi_h\rangle = \sum_{i=1}^{M} \langle v_i|\psi_h\rangle \cdot |i\rangle = \sum_{i=1}^{M} \frac{\langle v_i|\psi_h\rangle}{\sqrt{\mathsf{wt}_{V,i}}} \cdot \sqrt{\mathsf{wt}_{V,i}} \cdot |i\rangle = \Big(\sum_{i=1}^{M} \frac{\langle v_i|\psi_h\rangle}{\sqrt{\mathsf{wt}_{v,i}}} \cdot |i\rangle\langle i| \Big) \cdot |\mathsf{wt}_V\rangle\,. \tag{26}$$

This motivates the following definition.

Definition 4.6 (The rescaling matrix)**.** Let $h : [N] \to \{\pm 1\}$ be a Boolean function. Then the corresponding *rescaling matrix* is the diagonal matrix given by

$$D_{V,h} = \sum_{i=1}^{M} D_{V,h,i} \cdot |i\rangle\langle i|\,, \text{ where } D_{V,h,i} = \frac{\langle v_i|\psi_h\rangle}{\sqrt{\mathsf{wt}_{V,i}}}.$$

By construction, we have that $V \cdot |\psi_h\rangle = D_{V,h} \cdot |\mathsf{wt}_V\rangle$.

Remark 4.7. Loosely speaking, the size of the rescaling matrix indicates how close the amplitudes of $V \cdot |\psi_h\rangle$ are to their "typical" values. If each diagonal entry of $D_{V,h}$ is close to 1 in magnitude, then the amplitudes of $V \cdot |\psi_h\rangle$'s are roughly typical; otherwise, at least one of $V \cdot |\psi_h\rangle$'s amplitudes is atypically large or small.

We can therefore express the decoupled distinguishing advantage as

$$\begin{aligned}
\Delta_{\mathsf{Decup}}(R, R' \mid f) &= \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [\, \langle \psi_{R_{\boldsymbol{k}}}| \cdot V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \cdot |\psi_{R'_{\boldsymbol{k}}}\rangle] \Big| \\
&= \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [\, \langle \mathsf{wt}_V| \cdot D^\dagger_{V,R_{\boldsymbol{k}}} \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot D_{V,R'_{\boldsymbol{k}}} \cdot |\mathsf{wt}_V\rangle] \Big|.
\end{aligned}$$

Now we observe that $\mathcal{O}$ and $D_{V,R_{\boldsymbol{k}}}$ are both diagonal matrices, and hence they both commute (and similarly for $D_{V,R'_{\boldsymbol{k}}}$). As a result, this is equal to

$$\begin{aligned}
\Delta_{\mathsf{Decup}}(R, R' \mid f) &= \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [\, \langle \mathsf{wt}_V| \cdot \mathcal{O}_f \cdot D^\dagger_{V,R_{\boldsymbol{k}}} \cdot \Pi \cdot D_{V,R'_{\boldsymbol{k}}} \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle] \Big| \\
&= \Big| \langle \mathsf{wt}_V| \cdot \mathcal{O}_f \cdot \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [D^\dagger_{V,R_{\boldsymbol{k}}} \cdot \Pi \cdot D_{V,R'_{\boldsymbol{k}}}] \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle \Big|.
\end{aligned}$$

Note that for any function f, $\mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$ is a unit vector. We can therefore upper-bound this expression by relaxing $\mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$ to be an arbitrary unit vector maximizing this expression. This gives the spectral relaxation.

Definition 4.8 (Spectral relaxation)**.** Let $R, R' : [K] \times [N] \to \{\pm 1\}$ be two function families, and let A denote an adversary. The *spectral relaxation of the decoupled distinguishing advantage* is given by

$$\Delta^{\mathsf{Spectral}}_{\mathsf{Decup}}(R, R') = \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [D^\dagger_{V,R_{\boldsymbol{k}}} \cdot \Pi \cdot D_{V,R'_{\boldsymbol{k}}}] \Big\|_{\mathrm{op}}.$$

From the above discussion, the following lemma is immediate.

Lemma 4.9. *Let $R, R' : [K] \times [N] \to \{\pm 1\}$ be two function families. Then*

$$\max_f \Big\{ \Delta_{\mathsf{Decup}}(R, R' \mid f) \Big\} \le \Delta^{\mathsf{Spectral}}_{\mathsf{Decup}}(R, R').$$

4.3 Expectation of the spectral relaxation with one parameter held fixed

The spectral relaxation is the operator norm of a matrix which is bilinear in both R and R'. Keeping R fixed, we can consider a uniformly random $\boldsymbol{R}' : [K] \times [N] \to \{\pm 1\}$, and doing so makes this a random matrix whose entries are linear combinations of random $\{\pm 1\}$ variables. The key technical result we will use to study such matrices is the following, stated in [Tro15, Theorem 4.1.1].

Theorem 4.10 (Concentration for matrix Rademacher series)**.** *Let $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$ be n independent, uniformly distributed $\{\pm 1\}$ random variables. Let $\boldsymbol{Z}$ be a $d_1 \times d_2$ complex matrix whose entries are linear combinations of the $\boldsymbol{x}_k$'s, i.e.*

$$\boldsymbol{Z}_{i,j} = c_{i,j,1} \cdot \boldsymbol{x}_1 + \cdots + c_{i,j,n} \cdot \boldsymbol{x}_n,$$

where each $c_{i,j,k}$ is a fixed, complex number. Let $v(\boldsymbol{Z})$ be the matrix variance statistic of $\boldsymbol{Z}$, i.e.

$$v(\boldsymbol{Z}) = \max\{\| \mathbf{E}[\boldsymbol{Z} \cdot \boldsymbol{Z}^\dagger]\|_{\mathrm{op}}, \| \mathbf{E}[\boldsymbol{Z}^\dagger \cdot \boldsymbol{Z}]\|_{\mathrm{op}}\}.$$

Then

$$\mathbf{E}[\|\boldsymbol{Z}\|_{\mathrm{op}}] \leq \sqrt{2\ln(d_1 + d_2)} \cdot \sqrt{v(\boldsymbol{Z})}.$$

Furthermore, for all $t \geq 0$,

$$\mathbf{Pr}\left[\|\boldsymbol{Z}\|_{\mathrm{op}} \geq t\right] \leq (d_1 + d_2) \cdot \exp\left(-\frac{t^2}{2 \cdot v(\boldsymbol{Z})}\right).$$

We now use this to upper bound the expectation of the spectral relaxation when one of the parameters is held fixed. It states that this expectation can be bounded in terms of a quantity called the *width* of the function family R. Roughly speaking, the width is a measure of the "size" of the diagonal rescaling matrices D_{V,R_k}, over all $1 \leq k \leq K$. As discussed in Remark 4.7, when R is a "typical" function family, we expect that these rescaling matrices should have small (i.e. close to 1) entries on the diagonal, in which case the width of R will be small. For atypical function families, on the other hand, the width might be large, but we expect such families to be extremely rare.

Lemma 4.11 (Expectation of the spectral relaxation with one parameter held fixed)**.** *Let $R : [K] \times [N] \to \{\pm 1\}$ be a fixed function family. Define the* width *of R to be the quantity*

$$\mathsf{width}(R) \coloneqq \max_{1 \leq i \leq M} \left\{ \frac{1}{K} \sum_{k=1}^{K} \frac{|\langle v_i | \psi_{R_k} \rangle|^2}{\mathsf{wt}_{V,i}} \right\}.$$

In addition, let $\boldsymbol{R}' : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Then

$$\mathop{\mathbf{E}}_{\boldsymbol{R}'}[\Delta^{\mathsf{Spectral}}_{\mathsf{Decup}}(R, \boldsymbol{R}')] \leq \sqrt{\frac{2\ln(2M) \cdot \mathsf{width}(R)}{K}}.$$

Proof. For the reader's convenience, we will recall the definition of the diagonal rescaling matrix corresponding to a Boolean function $h : [N] \to \{\pm 1\}$:

$$D_{V,h} = \sum_{i=1}^{M} \frac{\langle v_i | \psi_h \rangle}{\sqrt{\mathsf{wt}_{V,i}}} \cdot |i\rangle\langle i| \,.$$

Note that each entry of $D_{V,h}$ is a linear combination of the Boolean values $h(1), \ldots, h(N)$. In addition, note that

$$\begin{aligned} D_{V,h} \cdot D_{V,h}^\dagger &= \Big(\sum_{i=1}^{M} \frac{\langle v_i | \psi_h \rangle}{\sqrt{\mathsf{wt}_{V,i}}} \cdot |i\rangle\langle i| \Big) \cdot \Big(\sum_{i=1}^{M} \frac{\langle \psi_h | v_i \rangle}{\sqrt{\mathsf{wt}_{V,i}}} \cdot |i\rangle\langle i| \Big) \\ &= \sum_{i=1}^{M} \frac{|\langle v_i | \psi_h \rangle|^2}{\mathsf{wt}_{V,i}} \cdot |i\rangle\langle i| \\ &= \Big(\sum_{i=1}^{M} \frac{\langle \psi_h | v_i \rangle}{\sqrt{\mathsf{wt}_{V,i}}} \cdot |i\rangle\langle i| \Big) \cdot \Big(\sum_{i=1}^{M} \frac{\langle v_i | \psi_h \rangle}{\sqrt{\mathsf{wt}_{V,i}}} \cdot |i\rangle\langle i| \Big) = D_{V,h}^\dagger \cdot D_{V,h}. \end{aligned} \tag{27}$$

Our goal is to compute

$$\mathop{\mathbf{E}}_{\boldsymbol{R}'}[\Delta_{\mathsf{Decup}}^{\mathsf{Spectral}}(R, \boldsymbol{R}')] = \mathop{\mathbf{E}}_{\boldsymbol{R}'} \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [D_{V,R_{\boldsymbol{k}}}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{R}'_{\boldsymbol{k}}}] \Big\|_{\mathrm{op}}.$$

To this end, define the matrix

$$\boldsymbol{Z} \coloneqq \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [D_{V,R_{\boldsymbol{k}}}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{R}'_{\boldsymbol{k}}}] = \frac{1}{K} \cdot \sum_{k=1}^{K} D_{V,R_k}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{R}'_k}.$$

For each $1 \le k \le K$, $D_{V,\boldsymbol{R}'_k}$ is a matrix whose entries are linear combinations of the $\boldsymbol{R}'_k(x)$'s. As a result, the entries of $\boldsymbol{Z}$ are linear combinations of the $K \cdot N$ many $\{\pm 1\}$-valued random variables in $\boldsymbol{R}'$. Hence, we can apply Theorem 4.10 to bound $\mathbf{E}_{\boldsymbol{R}'}[\|\boldsymbol{Z}\|_{\mathrm{op}}]$. To do so, we must first compute the matrix variance statistic of $\boldsymbol{Z}$. To begin,

$$\begin{aligned} \mathop{\mathbf{E}}_{\boldsymbol{R}'}[\boldsymbol{Z} \cdot \boldsymbol{Z}^\dagger] &= \mathop{\mathbf{E}}_{\boldsymbol{R}'} \Big[\Big(\frac{1}{K} \cdot \sum_{k=1}^{K} D_{V,R_k}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{R}'_k} \Big) \cdot \Big(\frac{1}{K} \cdot \sum_{k'=1}^{K} D_{V,\boldsymbol{R}'_{k'}}^\dagger \cdot \Pi \cdot D_{V,R_{k'}} \Big) \Big] \\ &= \frac{1}{K^2} \cdot \sum_{k,k'=1}^{K} D_{V,R_k}^\dagger \cdot \Pi \cdot \mathop{\mathbf{E}}_{\boldsymbol{R}'} [D_{V,\boldsymbol{R}'_k} \cdot D_{V,\boldsymbol{R}'_{k'}}^\dagger] \cdot \Pi \cdot D_{V,R_{k'}}. \end{aligned} \tag{28}$$

Now, if $k \neq k'$, then $\boldsymbol{R}'_k$ and $\boldsymbol{R}'_{k'}$ are distributed independently from each other. As a result, for any fixed matrix C,

$$\mathop{\mathbf{E}}_{\boldsymbol{R}'}[D_{V,\boldsymbol{R}'_k} \cdot C \cdot D_{V,\boldsymbol{R}'_{k'}}^\dagger] = \mathop{\mathbf{E}}_{\boldsymbol{R}'_k}[D_{V,\boldsymbol{R}'_k}] \cdot C \cdot \mathop{\mathbf{E}}_{\boldsymbol{R}'_{k'}}[D_{V,\boldsymbol{R}'_{k'}}^\dagger] = 0, \tag{29}$$

because $D_{V,\boldsymbol{R}'_k}$ and $D_{V,\boldsymbol{R}'_{k'}}$ are mean-zero. (For Equation (28) above we only need the $C = \mathsf{Id}_{M \times M}$ case, but we will apply it below using a different matrix C.) On the other hand, if $k = k'$, then by Proposition 4.4,

$$\begin{aligned} \mathop{\mathbf{E}}_{\boldsymbol{R}'}[D_{V,\boldsymbol{R}'_k} \cdot D_{V,\boldsymbol{R}'_k}^\dagger] &= \sum_{i=1}^{M} \mathop{\mathbf{E}}_{\boldsymbol{R}} \Big[\frac{|\langle v_i | \psi_{\boldsymbol{R}_k} \rangle|^2}{\mathsf{wt}_{V,i}} \Big] \cdot |i\rangle\langle i| && \text{(by Equation (27))} \\ &= \sum_{i=1}^{M} |i\rangle\langle i| = \mathsf{Id}_{M \times M}. \end{aligned} \tag{30}$$

Combining these two facts, we have that

$$(28) = \frac{1}{K^2} \cdot \sum_{k=1}^{K} D_{V,R_k}^\dagger \cdot \Pi^2 \cdot D_{V,R_k} \preceq \frac{1}{K^2} \cdot \sum_{k=1}^{K} D_{V,R_k}^\dagger \cdot \mathsf{Id}_{M\times M} \cdot D_{V,R_k} = \frac{1}{K^2} \cdot \sum_{k=1}^{K} D_{V,R_k}^\dagger \cdot D_{V,R_k}.$$

Finally, we bound this by

$$\begin{aligned}
\frac{1}{K^2} \cdot \sum_{k=1}^{K} D_{V,R_k}^\dagger \cdot D_{V,R_k} &= \frac{1}{K^2} \cdot \sum_{k=1}^{K} \sum_{i=1}^{M} \frac{|\langle v_i|\psi_{R_k}\rangle|^2}{\mathsf{wt}_{V,i}} \cdot |i\rangle\langle i| && \text{(by Equation (27))} \\
&= \frac{1}{K} \cdot \sum_{i=1}^{M} \Big(\frac{1}{K} \cdot \sum_{k=1}^{K} \frac{|\langle v_i|\psi_{R_k}\rangle|^2}{\mathsf{wt}_{V,i}}\Big) \cdot |i\rangle\langle i| \\
&\preceq \frac{1}{K} \cdot \sum_{i=1}^{M} \mathsf{width}(R) \cdot |i\rangle\langle i| = \Big(\frac{\mathsf{width}(R)}{K}\Big) \cdot \mathsf{Id}_{M\times M}. && (31)
\end{aligned}$$

So far, we have shown that

$$\|\mathop{\mathbf{E}}_{\boldsymbol{R}'}[\boldsymbol{Z}\cdot\boldsymbol{Z}^\dagger]\|_{\mathrm{op}} \le \Big\|\Big(\frac{\mathsf{width}(R)}{K}\Big)\cdot\mathsf{Id}_{M\times M}\Big\|_{\mathrm{op}} \le \frac{\mathsf{width}(R)}{K}.$$

Thus far, we have only computed the first term in the matrix variance statistic of $\boldsymbol{Z}$. Now we move on to the second term. Fortunately, we can reuse many of the steps involved in computing the first term to compute the second term:

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{R}'}[\boldsymbol{Z}^\dagger\cdot\boldsymbol{Z}] &= \mathop{\mathbf{E}}_{\boldsymbol{R}'}\Big[\Big(\frac{1}{K}\cdot\sum_{k=1}^{K} D_{V,\boldsymbol{R}'_k}^\dagger\cdot\Pi\cdot D_{V,R_k}\Big)\cdot\Big(\frac{1}{K}\cdot\sum_{k'=1}^{K} D_{V,R_{k'}}^\dagger\cdot\Pi\cdot D_{V,\boldsymbol{R}'_{k'}}\Big)\Big] \\
&= \frac{1}{K^2}\cdot\sum_{k,k'=1}^{K}\mathop{\mathbf{E}}_{\boldsymbol{R}'}\Big[D_{V,\boldsymbol{R}'_k}^\dagger\cdot\Pi\cdot D_{V,R_k}\cdot D_{V,R_{k'}}^\dagger\cdot\Pi\cdot D_{V,\boldsymbol{R}'_{k'}}\Big] \\
&= \frac{1}{K^2}\cdot\sum_{k=1}^{K}\mathop{\mathbf{E}}_{\boldsymbol{R}'}\Big[D_{V,\boldsymbol{R}'_k}^\dagger\cdot\Pi\cdot D_{V,R_k}\cdot D_{V,R_k}^\dagger\cdot\Pi\cdot D_{V,\boldsymbol{R}'_k}\Big]. && \text{(by Equations (27) and (29))}
\end{aligned}$$

Now, let $\boldsymbol{h}:[N]\to\{\pm1\}$ be a uniformly random Boolean function. Then $\boldsymbol{h}$ has the same distribution as $\boldsymbol{R}'_k$. for each $1\le k\le K$. As a result, this is equal to

$$\begin{aligned}
&\frac{1}{K^2}\cdot\sum_{k=1}^{K}\mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[D_{V,\boldsymbol{h}}^\dagger\cdot\Pi\cdot D_{V,R_k}\cdot D_{V,R_k}^\dagger\cdot\Pi\cdot D_{V,\boldsymbol{h}}\Big] \\
&= \mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[D_{V,\boldsymbol{h}}^\dagger\cdot\Pi\cdot\Big(\frac{1}{K^2}\cdot\sum_{k=1}^{K} D_{V,R_k}\cdot D_{V,R_k}^\dagger\Big)\cdot\Pi\cdot D_{V,\boldsymbol{h}}\Big] \\
&\preceq \mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[D_{V,\boldsymbol{h}}^\dagger\cdot\Pi\cdot\Big(\Big(\frac{\mathsf{width}(R)}{K}\Big)\cdot\mathsf{Id}_{M\times M}\Big)\cdot\Pi\cdot D_{V,\boldsymbol{h}}\Big] && \text{(by Equation (31))} \\
&= \Big(\frac{\mathsf{width}(R)}{K}\Big)\cdot\mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[D_{V,\boldsymbol{h}}^\dagger\cdot\Pi^2\cdot D_{V,\boldsymbol{h}}\Big] \\
&\preceq \Big(\frac{\mathsf{width}(R)}{K}\Big)\cdot\mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[D_{V,\boldsymbol{h}}^\dagger\cdot\mathsf{Id}_{M\times M}\cdot D_{V,\boldsymbol{h}}\Big] \\
&= \Big(\frac{\mathsf{width}(R)}{K}\Big)\cdot\mathop{\mathbf{E}}_{\boldsymbol{h}}\Big[D_{V,\boldsymbol{h}}^\dagger\cdot D_{V,\boldsymbol{h}}\Big] = \Big(\frac{\mathsf{width}(R)}{K}\Big)\cdot\mathsf{Id}_{M\times M}. && \text{(by Equations (27) and (30))}
\end{aligned}$$

In total, this shows that

$$\| \mathop{\mathbf{E}}_{\boldsymbol{R}'}[\boldsymbol{Z}^\dagger \cdot \boldsymbol{Z}]\|_{\mathrm{op}} \leq \left\| \left(\frac{\mathsf{width}(R)}{K} \right) \cdot \mathsf{Id}_{M \times M} \right\|_{\mathrm{op}} \leq \frac{\mathsf{width}(R)}{K}.$$

As a result, the matrix variance statistic of $\boldsymbol{Z}$ is

$$v(\boldsymbol{Z}) = \max\{\| \mathbf{E}[\boldsymbol{Z} \cdot \boldsymbol{Z}^\dagger]\|_{\mathrm{op}}, \| \mathbf{E}[\boldsymbol{Z}^\dagger \cdot \boldsymbol{Z}]\|_{\mathrm{op}}\} \leq \frac{\mathsf{width}(R)}{K}.$$

Now we apply Theorem 4.10. It states that

$$\mathop{\mathbf{E}}_{\boldsymbol{R}'}[\|\boldsymbol{Z}\|_{\mathrm{op}}] \leq \sqrt{2\ln(2M)} \cdot \sqrt{v(\boldsymbol{Z})} \leq \sqrt{2\ln(2M)} \cdot \sqrt{\frac{\mathsf{width}(R)}{K}}.$$

This completes the proof. □

4.4 A bound on the width of a random state family

In the previous section, we showed that the expectation of the spectral relaxation, when one of the input state families $\boldsymbol{R}' : [K] \times [N] \to \{\pm 1\}$ is randomized, can be bounded by a parameter of the other input family R referred to as its *width*. In this section, we show how to bound the expected width of a uniformly random family of binary phase states $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$. For intuition, recall that $\mathsf{width}(\boldsymbol{R})$ is defined to be the quantity

$$\max_{1 \leq i \leq M} \left\{ \frac{1}{K} \sum_{k=1}^{K} \frac{|\langle v_i | \psi_{\boldsymbol{R}_k} \rangle|^2}{\mathsf{wt}_{V,i}} \right\}. \tag{32}$$

Let us fix a value $1 \leq i \leq M$ and consider the i-th average being maximized over. By Proposition 4.4, for each $1 \leq k \leq K$, the k-th term in the average has expectation exactly equal to 1, and indeed we will show that this term is close to 1 with high probability. As the i-th average is an average over K such terms, we expect that it should be extremely close to 1 with an extremely high probability, a probability so high that we can then union bound over all $1 \leq i \leq M$ and show that $\mathsf{width}(\boldsymbol{R})$ itself is close to 1 with high probability. From this, we will be able to conclude that the expectation is close to 1 as well.

To start, let us focus on the k-th term in the i-th average. It is the absolute value squared of the following quantity:

$$\frac{\langle v_i | \psi_{\boldsymbol{R}_k} \rangle}{\sqrt{\mathsf{wt}_{V,i}}} = \frac{1}{\sqrt{\mathsf{wt}_{V,i}}} \cdot \sum_{x=1}^{N} \left(v_{i,x} \cdot \frac{1}{\sqrt{N}} \cdot \boldsymbol{R}_k(x) \right) = \sum_{x=1}^{N} \left(\frac{v_{i,x}}{\sqrt{\langle v_i | v_i \rangle}} \right) \cdot \boldsymbol{R}_k(x).$$

This is just a complex-weighted linear combination of random $\{\pm 1\}$ variables. In addition, the sum of the squared weights is given by

$$\sum_{x=1}^{N} \left| \frac{v_{i,x}}{\sqrt{\langle v_i | v_i \rangle}} \right|^2 = \frac{1}{\langle v_i | v_i \rangle} \cdot \sum_{x=1}^{N} |v_{i,x}|^2 = \frac{1}{\langle v_i | v_i \rangle} \cdot \langle v_i | v_i \rangle = 1.$$

We would like to show that weighted sums of this form are highly concentrated. In particular, we will show that they posses a particular concentration property known as being *sub-exponential*.

Definition 4.12 (Sub-exponential random variables)**.** A random variable $\boldsymbol{X}$ is *sub-exponential with parameter* $\gamma > 0$ if

$$\mathbf{Pr}\left[|\boldsymbol{X}| > t\right] \leq 2 \cdot \exp\left(-\frac{t}{\gamma}\right), \quad \text{for all } t \geq 0.$$

This is shown in the next lemma, whose proof we defer to Section 4.6.

Lemma 4.13 (Each term in the width is sub-exponential)**.** *There exists a constant* $\gamma \geq 1$ *such that the following is true. Let* $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m$ *be independent and uniform* ± 1 *random variables, and let* $a_1, \ldots, a_m$ *be complex numbers such that* $|a_1|^2 + \cdots + |a_m|^2 = 1$*. Define* $\boldsymbol{S} = a_1 \cdot \boldsymbol{b}_1 + \cdots + a_m \cdot \boldsymbol{b}_m$*. Then the random variable* $|\boldsymbol{S}|^2 - 1$ *is mean-zero and sub-exponential with parameter* γ*.*

Now that we have shown our random variables are well-concentrated, we would like to that averages of them, as occur in the formula for the width (Equation (32) above), are extremely well-concentrated. This can be shown using Bernstein's inequality for averages of independent sub-exponential random variables, which is stated in [Ver18, Corollary 2.8.3].

Theorem 4.14 (Bernstein's inequality)**.** *There exists a constant* $c > 0$ *such that the following is true. Let* $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_m$ *be independent, mean-zero, sub-exponential random variables, each with sub-exponential parameter at most* γ*. Then we have*

$$\mathbf{Pr}\left[\left|\frac{1}{m}\sum_{i=1}^{m} \boldsymbol{X}_i\right| \geq t\right] \leq 2 \cdot \exp\left(-c \cdot \min\left\{\frac{t^2}{\gamma^2}, \frac{t}{\gamma}\right\} \cdot m\right), \quad \textit{for all } t \geq 0.$$

Now we combine these ingredients to show the following tail bound on the width.

Lemma 4.15 (Tail bound on the width)**.** *There exists a constant* $c > 0$ *such that the following is true. Let* $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ *be a uniformly random function. Then for all* $t \geq 0$*,*

$$\Pr_{\boldsymbol{R}}[\mathsf{width}(\boldsymbol{R}) \geq 1 + t] \leq 2M \cdot \exp\left(-c \cdot \min\{t^2, t\} \cdot K\right).$$

Proof of Lemma 4.15. For each $1 \leq i \leq M$ and $1 \leq k \leq K$, let us define the random variables

$$\mathsf{width}_{i,k}(\boldsymbol{R}) := \frac{|\langle v_i | \psi_{\boldsymbol{R}_k} \rangle|^2}{\mathsf{wt}_{V,i}} \quad \text{and} \quad \mathsf{width}_i(\boldsymbol{R}) := \frac{1}{K} \cdot \sum_{k=1}^{K} \mathsf{width}_{i,k}(\boldsymbol{R}).$$

Lemma 4.13 states that there is a constant $\gamma \geq 1$ such that $(\mathsf{width}_{i,k}(\boldsymbol{R}) - 1)$ is sub-exponential with parameter γ, for all $1 \leq i \leq M$ and $1 \leq k \leq K$. Now, fix a value $1 \leq i \leq M$. Since each $\mathsf{width}_{i,k}(\boldsymbol{R})$ only depends on $\boldsymbol{R}_k$, the random variables $(\mathsf{width}_{i,k}(\boldsymbol{R}) - 1)$ are independent across all $1 \leq k \leq K$. As a result, Bernstein's inequality states that there exists a constant $c > 0$ such that for all $t \geq 0$,

$$\begin{aligned}
\Pr_{\boldsymbol{R}}[|\mathsf{width}_i(\boldsymbol{R}) - 1| \geq t] &= \Pr_{\boldsymbol{R}}\left[\left|\frac{1}{K} \cdot \sum_{k=1}^{K} (\mathsf{width}_{i,k}(\boldsymbol{R}) - 1)\right| \geq t\right] \\
&\leq 2 \cdot \exp\left(-c \cdot \min\left\{\frac{t^2}{\gamma^2}, \frac{t}{\gamma}\right\} \cdot K\right) \\
&\leq 2 \cdot \exp\left(-c \cdot \min\left\{\frac{t^2}{\gamma^2}, \frac{t}{\gamma^2}\right\} \cdot K\right) && \text{(because } \gamma \geq 1\text{)} \\
&= 2 \cdot \exp\left(-\left(\frac{c}{\gamma^2}\right) \cdot \min\{t^2, t\} \cdot K\right).
\end{aligned}$$

Now, since the width is defined as $\mathsf{width}(\boldsymbol{R}) = \max_{1 \le i \le M}\{\mathsf{width}_i(\boldsymbol{R})\}$, we have that

$$\begin{aligned}
\mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\mathsf{width}(\boldsymbol{R}) \ge 1 + t] &= \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\exists_{1 \le i \le M}\{\mathsf{width}_i(\boldsymbol{R}) \ge 1 + t\}] \\
&\le \sum_{i=1}^{M} \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\mathsf{width}_i(\boldsymbol{R}) \ge 1 + t] && \text{(by the union bound)} \\
&\le \sum_{i=1}^{M} \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[|\mathsf{width}_i(\boldsymbol{R}) - 1| \ge t] \\
&\le \sum_{i=1}^{M} 2 \cdot \exp\left(-\left(\frac{c}{\gamma^2}\right) \cdot \min\{t^2, t\} \cdot K\right) \\
&= 2M \cdot \exp\left(-\left(\frac{c}{\gamma^2}\right) \cdot \min\{t^2, t\} \cdot K\right).
\end{aligned}$$

This completes the proof, by taking the constant "c" in the lemma statement to be c/γ^2. □

Finally, we use our tail bound to derive an expectation bound on the width. Our proof will allow us to prove a bound of $1 + o(1)$ for a wide range of parameters M and K, as our initial intuition suggested. However, to get a bound which applies to the widest relevant range of parameters, we will prove a slightly weaker $O(1)$ bound, which is still sufficient for our applications.

Lemma 4.16 (Expectation bound on the width)**.** *There exists a constant $C \ge 1$ such that the following is true. Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Suppose that $M \le e^K$. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{R}}[\mathsf{width}(\boldsymbol{R})] \le C.$$

Proof. Fix some $\alpha \ge 1$, to be determined later. Then

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{R}}[\mathsf{width}(\boldsymbol{R})] &= \int_0^\infty \mathbf{Pr}[\mathsf{width}(\boldsymbol{R}) \ge t]dt \\
&= \int_0^{1+\alpha} \mathbf{Pr}[\mathsf{width}(\boldsymbol{R}) \ge t]dt + \int_{1+\alpha}^\infty \mathbf{Pr}[\mathsf{width}(\boldsymbol{R}) \ge t]dt \\
&\le 1 + \alpha + \int_{1+\alpha}^\infty \mathbf{Pr}[\mathsf{width}(\boldsymbol{R}) \ge t]dt \\
&= 1 + \alpha + \int_{\alpha}^\infty \mathbf{Pr}[\mathsf{width}(\boldsymbol{R}) \ge 1 + t]dt \\
&\le 1 + \alpha + \int_{\alpha}^\infty 2M \cdot \exp\left(-cK \cdot \min\{t^2, t\}\right)dt && \text{(by Lemma 4.15)} \\
&= 1 + \alpha + \int_{\alpha}^\infty 2M \cdot \exp\left(-cK \cdot t\right)dt. && \text{(because } \alpha \ge 1\text{)}
\end{aligned}$$

We can compute the integral exactly:

$$2M \cdot \int_\alpha^\infty \exp\left(-cK \cdot t\right)dt = -\frac{2M}{cK} \cdot \exp(-cK \cdot t)\Big|_\alpha^\infty = \frac{2M}{cK} \cdot \exp(-cK \cdot \alpha).$$

In total, this gives us a bound of

$$\mathop{\mathbf{E}}_{\boldsymbol{R}}[\mathsf{width}(\boldsymbol{R})] \le 1 + \alpha + \frac{2M}{cK} \cdot \exp(-cK \cdot \alpha) \le 1 + \alpha + \frac{2M}{c} \cdot \exp(-cK \cdot \alpha).$$

Now we select α to be $\alpha = \max\{1, c^{-1}\}$. Then we get a bound of

$$1 + \max\{1, c^{-1}\} + \frac{2M}{c} \cdot \exp(-cK \cdot \max\{1, c^{-1}\}) \leq 1 + \max\{1, c^{-1}\} + \frac{2M}{c} \cdot \exp(-K).$$

When $M \leq e^K$, this is at most

$$1 + \max\{1, c^{-1}\} + \frac{2}{c},$$

which is a constant. Picking this for the "C" in the lemma statement completes the proof. □

4.5 The one-query lower bound

Now we complete the proof of the one-query lower bound. We begin by proving a bound on the expected value of the distinguishing advantage.

Theorem 4.17 (Expectation bound for the distinguishing advantage)**.** *There exists a constant $C > 0$ such that the following is true. Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})] \leq C \cdot \sqrt{\frac{\ln(M)}{K}}.$$

Proof of Theorem 4.17. We will consider two regimes of parameters: $M \leq e^K$ and $M > e^K$. Let us first consider the case of $M \leq e^K$. Let $\boldsymbol{R}, \boldsymbol{R}' : [K] \times [N] \to \{\pm 1\}$ be two independent and uniformly random function families. Then

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})] &\leq 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}\left[\max_f \left\{\Delta_{\mathsf{Decup}}(\boldsymbol{R}, \boldsymbol{R}' \mid f)\right\}\right] && \text{(by Lemma 4.3)} \\
&\leq 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{R}'}[\Delta^{\mathsf{Spectral}}_{\mathsf{Decup}}(\boldsymbol{R}, \boldsymbol{R}')] && \text{(by Lemma 4.9)} \\
&\leq 4 \cdot \mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\sqrt{\frac{2\ln(2M) \cdot \mathsf{width}(\boldsymbol{R})}{K}}\right] && \text{(by Lemma 4.11)} \\
&\leq 4 \cdot \sqrt{\mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\frac{2\ln(2M) \cdot \mathsf{width}(\boldsymbol{R})}{K}\right]} && \text{(by Jensen's inequality)} \\
&\leq 4 \cdot \sqrt{\frac{2\ln(2M) \cdot C}{K}} && \text{(by Lemma 4.16, for some constant } C \geq 1) \\
&\leq 4 \cdot \sqrt{\frac{2 \cdot (2\ln(M)) \cdot C}{K}} && \text{(because } M \geq 2) \\
&= 8\sqrt{C} \cdot \sqrt{\frac{\ln(M)}{K}}.
\end{aligned}$$

Picking the "C" in the theorem statement to be $8\sqrt{C}$, this completes the $M \leq e^K$ case. As for the $M > e^K$ case, we note that because the distinguishing advantage is a difference of two probabilities, it is always at most 1. Hence,

$$\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R})] \leq 1 \leq 8\sqrt{C} \leq 8\sqrt{C} \cdot \sqrt{\frac{\ln(M)}{K}}.$$

The first inequality is because $C \geq 1$, and the second inequality is because $M > e^K$. This completes the $M > e^K$ case, and therefore completes the proof. □

Combining this with Lemma 3.18, we have our main technical result.

Theorem 4.18 (Main theorem)**.** *There exists a constants $C_1, C_2 > 0$ such that the following is true. Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Then*

$$\Pr_{\boldsymbol{R}}\left[\Delta_A(\boldsymbol{R}) \geq C_1 \cdot \sqrt{\frac{\ln(M)}{K}} + \varepsilon\right] \leq 4 \cdot \exp\bigl(-C_2 \cdot \varepsilon^2 KN\bigr).$$

In particular, this implies Theorem 1.4.

4.6 Technical lemma: sub-exponential random variables

Now we prove Lemma 4.13. For convenience, we restate it here.

Lemma 4.19 (Lemma 4.13 restated)**.** *There exists a constant $\gamma \geq 1$ such that the following is true. Let $\boldsymbol{b}_1, \dots, \boldsymbol{b}_m$ be independent and uniform ± 1 random variables, and let $a_1, \dots, a_m$ be complex numbers such that $|a_1|^2 + \cdots + |a_m|^2 = 1$. Define $\boldsymbol{S} = a_1 \cdot \boldsymbol{b}_1 + \cdots + a_m \cdot \boldsymbol{b}_m$. Then the random variable $|\boldsymbol{S}|^2 - 1$ is mean-zero and sub-exponential with parameter γ.*

To compute the mean of $|\boldsymbol{S}|^2 - 1$, we will use the following proposition.

Proposition 4.20. *Let $\boldsymbol{b}_1, \dots, \boldsymbol{b}_m$ be independent and uniform ± 1 random variables, and let $a_1, \dots, a_m$ be complex numbers. Then $\boldsymbol{S} = a_1 \cdot \boldsymbol{b}_1 + \cdots + a_m \cdot \boldsymbol{b}_m$ satisfies*

$$\mathbf{E}[|\boldsymbol{S}|^2] = \sum_{i=1}^m |a_i|^2.$$

Proof. We calculate

$$\begin{aligned}\mathbf{E}[|\boldsymbol{S}|^2] = \mathbf{E}[\boldsymbol{S}^\dagger \cdot \boldsymbol{S}] &= \mathbf{E}[\Big(\sum_{i=1}^m a_i \cdot \boldsymbol{b}_i\Big)^\dagger \cdot \Big(\sum_{j=1}^m a_j \cdot \boldsymbol{b}_i\Big)] \\ &= \mathbf{E}\Big[\sum_{i,j=1}^m a_i^\dagger a_j \cdot \boldsymbol{b}_i \boldsymbol{b}_j\Big] = \sum_{i,j=1}^m a_i^\dagger a_j \cdot \mathbf{E}[\boldsymbol{b}_i \boldsymbol{b}_j] = \sum_{i=1}^m |a_i|^2,\end{aligned}$$

where the final step used $\mathbf{E}[\boldsymbol{b}_i \boldsymbol{b}_j] = 1$ if $i = j$ and 0 if $i \neq j$. This completes the proof. □

To show concentration for $|\boldsymbol{S}|^2 - 1$, we use the following tail bound, a version of Hoeffding's inequality for complex-weighted random sums.

Theorem 4.21 (Sub-Gaussian concentration for sums of complex random variables)**.** *Let $\boldsymbol{b}_1, \dots, \boldsymbol{b}_m$ be independent and uniform ± 1 random variables, and let $a_1, \dots, a_m$ be complex numbers. Then $\boldsymbol{S} = a_1 \cdot \boldsymbol{b}_1 + \cdots + a_m \cdot \boldsymbol{b}_m$ satisfies*

$$\mathbf{Pr}[|\boldsymbol{S}| \geq t] \leq 2 \cdot \exp\left(-\frac{t^2}{2 \cdot \sum_{i=1}^m |a_i|^2}\right).$$

As it turns out, this can be proved as a (very) special case of the *matrix concentration* tail bound stated in Theorem 4.10.

Proof of Theorem 4.21. We invoke Theorem 4.10 by treating $\boldsymbol{S}$ as a 1×1 complex-valued matrix. In particular, define the 1×1 matrix

$$\widehat{\boldsymbol{S}} \coloneqq (\boldsymbol{S}) = (a_1 \cdot \boldsymbol{b}_1 + \cdots + a_m \cdot \boldsymbol{b}_m).$$

Then its "matrix" variance parameter is

$$v(\widehat{\boldsymbol{S}}) = \max\{\|\mathbf{E}[\widehat{\boldsymbol{S}} \cdot \widehat{\boldsymbol{S}}^\dagger]\|_{\text{op}}, \|\mathbf{E}[\widehat{\boldsymbol{S}}^\dagger \cdot \widehat{\boldsymbol{S}}]\|_{\text{op}}\} = \|\mathbf{E}\left(|\boldsymbol{S}|^2\right)\|_{\text{op}} = \mathbf{E}[|\boldsymbol{S}|^2] = \sum_{i=1}^{m} |a_i|^2,$$

where the final step used Proposition 4.20. Then the tail bound of Theorem 4.10 implies that for all $t \geq 0$,

$$\mathbf{Pr}\left[|\boldsymbol{S}| \geq t\right] = \mathbf{Pr}\left[\|\widehat{\boldsymbol{S}}\|_{\text{op}} \geq t\right] \leq 2 \cdot \exp\left(-\frac{t^2}{2 \cdot \sum_{i \in [m]} |a_i|^2}\right).$$

This completes the proof. □

The following is an immediate corollary of Theorem 4.21.

Corollary 4.22. *Let $\boldsymbol{S}$ be as in Lemma 4.19. Then $|\boldsymbol{S}|^2$ is a sub-exponential random variable with parameter $\gamma = 2$.*

Proof. By Theorem 4.21,

$$\mathbf{Pr}[|\boldsymbol{S}| \geq t] \leq 2 \cdot \exp\left(-\frac{t^2}{2 \cdot \sum_{i=1}^{m} |a_i|^2}\right) = 2 \cdot \exp(-t^2/2).$$

Hence,

$$\mathbf{Pr}[|\boldsymbol{S}|^2 \geq t] \leq 2 \cdot \exp(-t/2).$$

This means that $|\boldsymbol{S}|^2$ is a sub-exponential random variable with parameter $\gamma = 2$. □

Now we want to show that $|\boldsymbol{S}|^2 - 1$ is also sub-exponential, taking advantage of the fact that $\mathbf{E}[|\boldsymbol{S}|^2] = 1$. To do so, we will use standard facts about sub-exponential random variables from [Ver18, Section 2.7]. In particular, we will rely on an alternative method of parameterizing sub-exponential random variables in terms of their *moment generation functions (MGFs)*.

Definition 4.23 (Sub-exponential norm)**.** Given a real random variable $\boldsymbol{X}$, the MGF of $|\boldsymbol{X}|$ is *bounded at point $\kappa > 0$* if

$$\mathbf{E}[\exp(|\boldsymbol{X}|/\kappa)] \leq 2.$$

The smallest κ for which this equation is holds is given by the *sub-exponential norm of $\boldsymbol{X}$*, denoted $\|\boldsymbol{X}\|_{\psi_1}$, and is defined formally as follows:

$$\|\boldsymbol{X}\|_{\psi_1} = \inf\{t > 0 : \mathbf{E}[\exp(|\boldsymbol{X}|/\kappa)] \leq 2\}.$$

We require two facts about this method of parameterizing sub-exponential random variables. The first is stated in [Ver18, Proposition 2.7.1] and the second is stated in [Ver18, Exercise 2.7.10].

Proposition 4.24 (Approximate equivalence of the two parameterizations)**.** *There is an absolute constant $C_1 > 0$ such that the following is true. If the MGF of $|\boldsymbol{X}|$ is bounded at point κ, then $\boldsymbol{X}$ is sub-exponential with parameter γ, for some $\gamma \leq C_1 \cdot \kappa$. Likewise, if $\boldsymbol{X}$ is sub-exponential with parameter γ, then the MGF of $|\boldsymbol{X}|$ is bounded at point κ, for some $\kappa \leq C_1 \cdot \gamma$.*

Proposition 4.25 (Centering)**.** *There is an absolute constant $C_2 > 0$ such that the following is true. If $\boldsymbol{X}$ is a sub-exponential random variable, then so is $\boldsymbol{X} - \mathbf{E}[\boldsymbol{X}]$, and it satisfies*

$$\|\boldsymbol{X} - \mathbf{E}[\boldsymbol{X}]\|_{\psi_1} \leq C_2 \cdot \|\boldsymbol{X}\|_{\psi_1}.$$

Now we prove Lemma 4.19.

Proof of Lemma 4.19. First, Proposition 4.20 states that

$$\mathbf{E}[|\boldsymbol{S}|^2] = \sum_{i=1}^{m} |a_i|^2 = 1.$$

Hence, $|\boldsymbol{S}|^2 - 1$ is mean-zero. Next, Corollary 4.22 states that $|\boldsymbol{S}|^2$ is a sub-exponential random variable with parameter $\gamma_1 = 2$. Proposition 4.24 then implies that the MGF of $||\boldsymbol{S}|^2|$ is bounded at point κ_1, for some

$$\kappa_1 \leq C_1 \cdot \gamma_1 = 2 \cdot C_1.$$

By definition of the sub-exponential norm, this immediately implies that $\||\boldsymbol{S}|^2\|_{\psi_1} \leq 2 \cdot C_1$. Proposition 4.25 then implies that

$$\||\boldsymbol{S}|^2 - 1\|_{\psi_1} = \||\boldsymbol{S}|^2 - \mathbf{E}[|\boldsymbol{S}|^2]\|_{\psi_1} \leq C_2 \cdot \||\boldsymbol{S}|^2\|_{\psi_1} \leq C_2 \cdot 2 \cdot C_1.$$

Now $|\boldsymbol{S}|^2 - 1$ is a non-constant random variable, and in particular it is nonzero with finite probability. In addition, it only obtains a discrete set of values. Hence, the infimum over $\{t > 0\}$ in the definition of the sub-exponential norm $\||\boldsymbol{S}|^2 - 1\|_{\psi_1}$ is achieved at a nonzero minimizing value $\kappa_2 > 0$; in other words, if we set

$$\kappa_2 = \||\boldsymbol{S}|^2 - 1\|_{\psi_1} \leq 2 \cdot C_1 \cdot C_2,$$

then the MGF of $||\boldsymbol{S}|^2 - 1|$ is bounded at point κ_2. Applying Proposition 4.24 again, this implies that $|\boldsymbol{S}|^2 - 1$ is sub-exponential with parameter γ_2, for some

$$\gamma_2 \leq C_1 \cdot \kappa_2 \leq 2 \cdot C_1^2 \cdot C_2.$$

Now, we note that if a random variable $\boldsymbol{X}$ is sub-exponential with parameter $a > 0$, then it is also sub-exponential with parameter b, for any $b \geq a$. This is because for all $t > 0$,

$$\mathbf{Pr}\left[|\boldsymbol{X}| > t\right] \leq 2 \cdot \exp\left(-\frac{t}{a}\right) \leq 2 \cdot \exp\left(-\frac{t}{b}\right).$$

Hence, because $|\boldsymbol{S}|^2 - 1$ is sub-exponential for parameter $\gamma_2 \leq 2 \cdot C_1^2 \cdot C_2$, it is also sub-exponential for parameter $\gamma = \max\{1, 2 \cdot C_1^2 \cdot C_2\}$. This is a constant which is greater than or equal to 1, which completes the proof. □

5 Pseudorandom states relative to a random oracle

In this section, we use Theorem 4.18 to derive Theorem 1.2, our lower bound for breaking pseudorandom state families. We begin with a definition of (single-copy) pseudorandom states in the plain model, for reference.

Definition 5.1 (Pseudorandom state family)**.** Let $n : \mathbb{N} \to \mathbb{N}$ be a function and $\{|\psi_{\lambda,k}\rangle\}_{k \in \{0,1\}^\lambda}$ be a family of $n(\lambda)$-qubit quantum states for each $\lambda \in \mathbb{N}$. Then the state family ensemble

$$\{\{|\psi_{\lambda,k}\rangle\}_{k \in \{0,1\}^\lambda}\}_{\lambda \in \mathbb{N}}$$

is a *pseudorandom state (PRS) family* if it has the following properties.

- **Efficient constructability:** there is a polynomial-time quantum algorithm that on input $(1^\lambda, k)$, for $k \in \{0,1\}^\lambda$, outputs $|\psi_{\lambda,k}\rangle$.
- **Stretch:** $n(\lambda) \geq \lambda + 1$, for all λ.
- **Pseudorandomness:** for all algorithms A described by polynomial-size quantum circuit families, we have that

$$\left| \Pr_{\boldsymbol{k}\sim\{0,1\}^\lambda}\left[A(|\psi_{\lambda,\boldsymbol{k}}\rangle) \text{ outputs "0"}\right] - \Pr_{|\boldsymbol{\psi}\rangle}\left[A(|\boldsymbol{\psi}\rangle) \text{ outputs "0"}\right]\right| = \mathsf{negl}(\lambda),$$

  where $|\boldsymbol{\psi}\rangle$ is drawn from the Haar distribution on $n(\lambda)$-qubit states.

Our Oracle and Adversary Model. In this paper, we consider pseudorandom state families defined relative to an oracle $R : \{0,1\}^* \times \{0,1\}^* \to \{\pm 1\}$. In that case, the efficient constructability property requires that there is a quantum polynomial-time oracle algorithm that on input $(1^\lambda, k)$, for $k \in \{0,1\}^\lambda$, outputs $|\psi_{\lambda,k}\rangle$, given oracle access to R.

In addition, the pseudorandomness property should require that the PRS family be secure against all algorithms A^f, where $A^{(\cdot)}$ is an oracle algorithm described by a polynomial-size oracle circuit family and $f = f_R$ is an arbitrary R-dependent oracle; equivalently, $A^{(\cdot)}$ is computable by a family of quantum circuits output by a polynomial-time Turing machine with the help of polynomial-size non-uniform advice.

The main result of this section (Theorem 5.2 below) proves that relative to a random oracle, there are PRS families secure against all one-query attacks. Explicitly, the adversary model we consider is as follows:

- For a given function $R : \{0,1\}^* \times \{0,1\}^* \to \{\pm 1\}$, the adversary is described by an R-dependent Turing machine and R-dependent collection of advice strings $(z_\lambda)_{\lambda\in\mathbb{N}}$.
- On input z_λ, the Turing machine outputs the description of a one-query oracle circuit $A^{(\cdot)} := A^{(\cdot)}_{R,z_\lambda}$.
- On input the state $|\psi\rangle$, the adversary executes the oracle circuit $A^{f_R}(|\psi\rangle)$ for a function $f_R : \{0,1\}^* \to \{\pm 1\}$ that may depend on R.

Theorem 5.2 (Theorem 1.2 formalized)**.** *Let $n(\lambda)$ be any efficiently computable polynomial function in λ such that $n(\lambda) \geq \lambda + 1$ for all λ. Then with probability 1 over the choice of a random oracle $\boldsymbol{R} : \{0,1\}^* \times \{0,1\}^* \to \{\pm 1\}$, the following is true relative to $\boldsymbol{R}$. There exists a PRS family consisting of $n(\lambda)$-qubit quantum states that is secure against all polynomial-time quantum algorithms A^f that have polynomial-size non-uniform classical advice and make one query to an arbitrary Boolean function $f : \{0,1\}^* \to \{\pm 1\}$.*

Proof. For each $\lambda \in \mathbb{N}$, we define the function family $\boldsymbol{R}^\lambda : \{0,1\}^\lambda \times \{0,1\}^{n(\lambda)} \to \{\pm 1\}$ by setting

$$\boldsymbol{R}^\lambda_k(x) := \boldsymbol{R}(k, x),$$

for each $k \in \{0,1\}^\lambda$ and $x \in \{0,1\}^{n(\lambda)}$. Then the candidate PRS family is the state family ensemble which contains the family of $n(\lambda)$-qubit quantum states

$$\{\, |\psi_{\boldsymbol{R}^\lambda_k}\rangle\}_{k\in\{0,1\}^\lambda},$$

for each security parameter $\lambda \in \mathbb{N}$. By construction, the state $|\psi_{\boldsymbol{R}_k^\lambda}\rangle$ can be generated in time $\mathrm{poly}(\lambda)$ given a single oracle call to $\boldsymbol{R}$. Thus, all that remains is to establish security.

Security *nearly* follows from Theorem 4.18, except that the order of quantifiers is wrong: in Theorem 4.18, the oracle circuit $A^{(\cdot)}$ is not allowed to depend on $\boldsymbol{R}$, although the function f it queries is. However, in this setting, $A^{(\cdot)}$ *is* allowed to depend on $\boldsymbol{R}$. We handle this by a standard quantifier-switching argument using the Borel-Cantelli lemma [BG81, IR89], which applies even in the case of A with bounded non-uniformity.

The argument is as follows. We abuse notation and let $A(\cdot)$ denote a polynomial-time Turing machine that on input z_λ outputs a one-query oracle circuit $A_{z_\lambda}^{(\cdot)}(\cdot)$. The adversary runs $A_{z_\lambda}^{f}$ on input state $|\psi\rangle$ using an arbitrary $\boldsymbol{R}$-dependent oracle $f = f_{\boldsymbol{R}}$. Here, $z = \{z_\lambda\}_\lambda$ is a collection of advice strings in which z_λ has length $\mathrm{poly}(\lambda)$. Because $A(\cdot)$ runs in polynomial time, the query length of $A_{z_\lambda}^{f}$ is bounded by some $p(\lambda) = \mathrm{poly}(\lambda)$. As a result, by Theorem 4.18 (setting $\varepsilon = \frac{1}{\sqrt{K}}$), we know that for every security parameter $\lambda \in \mathbb{N}$,

$$\Pr_{\boldsymbol{R}^\lambda}\left[\Delta_{A_{z_\lambda}}(\boldsymbol{R}^\lambda) \geq \frac{(1 + C_1 \cdot \sqrt{p(\lambda)})}{\sqrt{K}}\right] \leq 4 \cdot \exp(-C_2 \cdot N) = 4 \cdot \exp(-C_2 \cdot 2^{n(\lambda)}) \leq 4 \cdot \exp(-C_2 \cdot 2^\lambda),$$

where the last inequality uses the fact that $n(\lambda) \geq \lambda + 1$. We may then union bound over the $2^{p(\lambda)}$ possible advice strings z_λ and conclude that

$$\Pr_{\boldsymbol{R}^\lambda}\left[\exists z : \Delta_{A_{z_\lambda}}(\boldsymbol{R}^\lambda) \geq \frac{(1 + C_1 \cdot \sqrt{p(\lambda)})}{\sqrt{K}}\right] \leq 2^{p(\lambda)} \cdot 4 \cdot \exp(-C_2 \cdot 2^\lambda) \leq 4 \cdot \exp(-c \cdot 2^\lambda),$$

for a universal constant $c > 0$ and all sufficiently large λ.

Let $\mathcal{E}_\lambda$ denote the above event. Then, we know that the summation

$$\sum_{\lambda \in \mathbb{N}} \Pr_{\boldsymbol{R}:\{0,1\}^* \times \{0,1\}^* \to \{\pm 1\}}[\mathcal{E}_\lambda] < \infty$$

converges to a real number. Therefore, by the Borel-Cantelli lemma,

$$\Pr_{\boldsymbol{R}:\{0,1\}^* \times \{0,1\}^* \to \{\pm 1\}}\left[\mathcal{E}_\lambda \text{ occurs for infinitely many } \lambda\right] = 0.$$

Therefore, for all sufficiently large $\lambda \in \mathbb{N}$,

$$\Delta_{A_{z_\lambda}}(\boldsymbol{R}^\lambda) \leq \frac{\mathrm{poly}(\lambda)}{\sqrt{K}}, \tag{33}$$

no matter what advice $z = \{z_\lambda\}_\lambda$ the algorithm is given. Finally, we observe that the probability space above is uncountable. Therefore, we may union bound over all countably many polynomial-time Turing machines $A(\cdot)$ and conclude that Equation (33) holds for all $A(\cdot)$ and all sufficiently large $\lambda \in \mathbb{N}$. This shows that the PRS family satisfies the claimed pseudorandomness property, concluding the proof. □

References

[Aar16] Scott Aaronson. The complexity of quantum states and transformations: from quantum money to black holes. Technical report, arXiv:1607.05256, 2016. 1, 1.3.2

[Aar21] Scott Aaronson. Open problems related to quantum query complexity, comment #36, 2021. https://scottaaronson.blog/?p=5837. 3.3

[AK07] Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. In *Proceedings of the 22nd Annual IEEE Conference on Computational Complexity*, pages 115–128, 2007. 1, 1, 1.3.1, 1.3.1, 1.4, 3.3, B.2

[AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *Proceedings of the 42nd Annual International Cryptology Conference*, pages 208–236, 2022. 1.1

[BCKM21] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In *Proceedings of the 41st Annual International Cryptology Conference*, pages 467–496, 2021. 1.1

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In *Proceedings of the 14th Innovations in Theoretical Computer Science*, pages 24:1–24:21, 2023. 1.1

[BEM+23] John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the Uhlmann transformation problem. *Technical report, arXiv:2306.13073*, 2023. 1.2, 1.3.3, 3.3

[BG81] Charles Bennett and John Gill. Relative to a random oracle a, $\mathrm{P}^a \neq \mathrm{NP}^a \neq \mathrm{co}-\mathrm{NP}^a$ with probability 1. *SIAM Journal on Computing*, 10(1):96–113, 1981. 5

[BG22] Anne Broadbent and Alex Grilo. QMA-hardness of consistency of local density matrices with applications to quantum zero-knowledge. *SIAM Journal on Computing*, 51(4):1400–1450, 2022. 1.1

[BV97] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM journal on computing*, 26(5):1411–1473, 1997. 1, 2.1

[CGLQ20] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*, pages 673–684, 2020. 2.2

[CLQ20] Kai-Min Chung, Tai-Ning Liao, and Luowen Qian. Lower bounds for function inversion with quantum advice. In *Proceedings of the 1st Conference on Information-Theoretic Cryptography*, 2020. 2.2

[DGLM23] Hugo Delavenne, François Le Gall, Yupan Liu, and Masayuki Miyamoto. Quantum Merlin-Arthur proof systems for synthesizing quantum states. Technical report, arXiv: 2303.01877, 2023. 1.3.3

[GJMZ23] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1579–1588, 2023. 1.1

[GLSV21] Alex Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in MiniQCrypt. In *Proceedings of the 40th Annual International Cryptology Conference*, pages 531–561, 2021. 1.1

[HXY19] Minki Hhan, Keita Xagawa, and Takashi Yamakawa. Quantum random oracle model with auxiliary input. In *Proceedings of the 25th Annual International Conference on the Theory and Application of Cryptology and Information Security*, pages 584–614, 2019. 2.2

[INN+22] Sandy Irani, Anand Natarajan, Chinmay Nirkhe, Sujit Rao, and Henry Yuen. Quantum search-to-decision reductions and the state synthesis problem. In *Proceedings of the 37th Annual IEEE Conference on Computational Complexity*, pages 1–19, 2022. 1, 1.3.2, 1.3.3

[IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 44–61, 1989. 5

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Proceedings of the 38th Annual International Cryptology Conference*, pages 126–152, 2018. 1.1, 1

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in Algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, 2023. 1.1

[Kre21] William Kretschmer. Quantum pseudorandomness and classical complexity. In *Proceedings of the 16th Conference on the Theory of Quantum Computation, Communication and Cryptography*, 2021. 1

[Kre23] William Kretschmer. Does quantum cryptography imply classical lower bounds? Talk at the Simons Institute, 2023. 1.3.3

[Liu23] Qipeng Liu. Non-uniformity and quantum advice in the quantum random oracle model. In *Proceedings of the 42nd Annual International Cryptology Conference*, pages 117–143, 2023. 2.2

[LMR14] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014. 2

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Proceedings of the 42nd Annual International Cryptology Conference*, pages 269–295, 2022. 1.1, 1

[MY23] Tony Metger and Henry Yuen. stateQIP= statePSPACE. In *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science*, 2023. 1.3.3

[NC10] Michael Nielsen and Isaac Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010. 3.5

[Pet12] Peter Petersen. *Linear algebra*. Springer, 2012. 3.5

[Ros22] Gregory Rosenthal. Query and depth upper bounds for quantum unitaries via Grover search. Technical report, arXiv:2111.07992, 2022. 1, 1.3.1, 3.3

[Ros23a] Gregory Rosenthal. Efficient quantum state synthesis with one query. Technical report, arXiv:2306.01723, 2023. 1, 1.2, 1.3.2, 1.3.3, 3

[Ros23b] Gregory Rosenthal. *Quantum State and Unitary Complexity*. PhD thesis, University of Toronto, 2023. 1, 1.3.2, 3.5, 3.6

[RY22] Gregory Rosenthal and Henry Yuen. Interactive proofs for synthesizing quantum states and unitaries. In *Proceedings of the 13th Innovations in Theoretical Computer Science*, 2022. 1.3.3, 3.3

[Sha49] Claude Shannon. The synthesis of two-terminal switching circuits. *The Bell System Technical Journal*, 28(1):59–98, 1949. 1

[Tro12] Joel Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12:389–434, 2012. 2.2, A.3

[Tro15] Joel Tropp. *An introduction to matrix concentration inequalities*, volume 8. 2015. 4.3

[Ver11] Roman Vershynin. A simple decoupling inequality in probability theory. Found at https://www.math.uci.edu/~rvershyn/papers/decoupling-simple.pdf, 2011. 2b, 4.1

[Ver18] Roman Vershynin. *High-dimensional probability: an introduction with applications in data science*. Cambridge University Press, 2018. 2.1, 3.4, 4.4, 4.6, 4.6

[vH17] Ramon van Handel. Structured random matrices. *Convexity and concentration*, pages 107–156, 2017. 2b, 4.1

[Wat18] John Watrous. *The theory of quantum information*. Cambridge University Press, 2018. 3.3

[Yan22] Jun Yan. General properties of quantum bit commitments. In *Proceedings of the 28th Annual International Conference on the Theory and Application of Cryptology and Information Security*, pages 628–657, 2022. 1.1

[Yue22a] Henry Yuen. Lecture 6 from COMS E6998: Frontiers of quantum complexity and cryptography. Found at https://www.henryyuen.net/spring2022/lec6-statesynthesis.pdf and https://www.henryyuen.net/spring2022/lec6-unitarysynthesis.pdf, 2022. 1, 1

[Yue22b] Henry Yuen. Lecture 7 from COMS E6998: Frontiers of quantum complexity and cryptography. Found at https://www.henryyuen.net/spring2022/lec7-quantumprograms.pdf, 2022. 1, 2

A A matrix Chernoff proof of the one-query lower bound

In this section, we will give an alternative proof of the one-query lower bound using a variant of the matrix Chernoff bound known as the matrix Hoeffding bound. Although this proof strategy ultimately results in worse bounds than the proof presented in Section 4, we have included it because we believe its techniques will be more familiar to a computer science audience. This section is organized as follows: in Appendix A.1 we will develop a natural spectral relaxation of the distinguishing advantage $\Delta_A(R)$. Next, we will perform a slight modification to this spectral relaxation in Appendix A.2 to handle function families R which produce outlier values. Finally, in Appendix A.3 we will use these ingredients to complete the proof of the one-query lower bound.

We will largely follow the same notation as the proof in Section 4, which can be found in Section 3.1 as well as in Section 4.2. For convenience, we repeat several important definitions and results here.

Definition A.1 (Isometry weights, Definition 4.1 restated)**.** The *isometry weights* are the numbers

$$\mathsf{wt}_{V,i} := \tfrac{1}{N} \cdot \langle v_i | v_i \rangle \,,$$

for $1 \le i \le M$. Note that these sum to one and therefore form a probability distribution.

Proposition A.2 (Typical amplitudes, Proposition 4.4 restated)**.** *Let $\boldsymbol{h} : [N] \to \{\pm 1\}$ be a uniformly random Boolean function. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{h}} |\, \langle v_i | \psi_{\boldsymbol{h}} \rangle \,|^2 = \mathsf{wt}_{V,i}.$$

Definition A.3 (The weight vector, Definition 4.5 restated)**.** The *weight vector* is the unit vector given by

$$|\mathsf{wt}_V\rangle = \sum_{i=1}^{M} \sqrt{\mathsf{wt}_{V,i}}\, |i\rangle \,.$$

Definition A.4 (The rescaling matrix, Definition 4.6 restated)**.** Let $h : [N] \to \{\pm 1\}$ be a Boolean function. Then the corresponding *rescaling matrix* is the diagonal matrix given by

$$D_{V,h} = \sum_{i=1}^{M} D_{V,h,i} \cdot |i\rangle\langle i| \,, \text{ where } D_{V,h,i} = \frac{\langle v_i | \psi_h \rangle}{\sqrt{\mathsf{wt}_{V,i}}}.$$

By construction, we have that $V \cdot |\psi_h\rangle = D_{V,h} \cdot |\mathsf{wt}_V\rangle$.

Theorem A.5 (Sub-Gaussian concentration for sums of complex random variables, Theorem 4.21 restated)**.** *Let $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m$ be independent and uniform ± 1 random variables, and let $a_1, \ldots, a_m$ be complex numbers. Then $\boldsymbol{S} = a_1 \cdot \boldsymbol{b}_1 + \cdots + a_m \cdot \boldsymbol{b}_m$ satisfies*

$$\mathbf{Pr}[|\boldsymbol{S}| \ge t] \le 2 \cdot \exp\left(-\frac{t^2}{2 \cdot \sum_{i=1}^{m} |a_i|^2} \right).$$

Now we move to the proof.

A.1 A spectral relaxation for the distinguishing advantage

Let $h : [N] \to \{\pm 1\}$ be a Boolean function, and consider the adversary's execution on the binary phase state $|\psi_h\rangle$. First, it applies the isometry, resulting in the state

$$V \cdot |\psi_h\rangle = D_{V,h} \cdot |\mathsf{wt}_V\rangle \,.$$

Next, it applies an oracle $\mathcal{O}_f$. This will produce the state

$$\mathcal{O}_f \cdot D_{V,h} \cdot |\mathsf{wt}_V\rangle \,.$$

Now we observe that $\mathcal{O}_f$ and $D_{V,h}$ are both diagonal matrices, and hence they both commute. As a result,

$$\mathcal{O}_f \cdot D_{V,h} \cdot |\mathsf{wt}_V\rangle = D_{V,h} \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle \,.$$

Note that $\mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$ is always a unit vector, and that it is independent of h.

Finally, the adversary performs the measurement $\{\Pi, I - \Pi\}$ and accepts if it observes the first outcome. We can therefore calculate the acceptance probability of the adversary A with oracle access to a function $f : [M] \to \{\pm 1\}$ as:

$$\begin{aligned} p_A(h \mid f) &= \langle \psi_h | \, V^\dagger \cdot \mathcal{O}_f \cdot \Pi \cdot \mathcal{O}_f \cdot V \, |\psi_h\rangle \\ &= \langle \mathsf{wt}_V | \, \mathcal{O}_f \cdot D_{V,h}^\dagger \cdot \Pi \cdot D_{V,h} \cdot \mathcal{O}_f \, |\mathsf{wt}_V\rangle \, . \end{aligned}$$

Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. By the above calculation and the definition of the distinguishing advantage $\Delta_A(R)$ allows us to conclude that

$$\begin{aligned} \Delta_A(R) &= \max_{f:[M]\to\{\pm 1\}} \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} [p_A(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}} [p_A(\boldsymbol{h} \mid f)] \Big| \\ &= \max_f \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} \langle \mathsf{wt}_V | \, \mathcal{O}_f \cdot D_{V,R_{\boldsymbol{k}}}^\dagger \cdot \Pi \cdot D_{V,R_{\boldsymbol{k}}} \cdot \mathcal{O}_f \, |\mathsf{wt}_V\rangle - \mathop{\mathbf{E}}_{\boldsymbol{h}} \langle \mathsf{wt}_V | \, \mathcal{O}_f \cdot D_{V,\boldsymbol{h}}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{h}} \cdot \mathcal{O}_f \, |\mathsf{wt}_V\rangle \Big| \\ &= \max_f \Big| \langle \mathsf{wt}_V | \, \mathcal{O}_f \cdot \Big(\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D_{V,R_{\boldsymbol{k}}}^\dagger \cdot \Pi \cdot R_{V,R_{\boldsymbol{k}}} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D_{V,\boldsymbol{h}}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{h}} \Big) \cdot \mathcal{O}_f \, |\mathsf{wt}_V\rangle \Big| . \end{aligned}$$

Recall that $\mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$ is a unit vector which depends on V and f. We can therefore upper-bound this expression by relaxing $\mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$ to be an arbitrary unit vector maximizing this expression. This gives the spectral relaxation.

Definition A.6 (Spectral relaxation)**.** Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. The *spectral relaxation* of the distinguishing probability on R is given by

$$\Delta_A^{\mathsf{Spectral}}(R) = \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D_{V,R_{\boldsymbol{k}}}^\dagger \cdot \Pi \cdot D_{V,R_{\boldsymbol{k}}} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D_{V,\boldsymbol{h}}^\dagger \cdot \Pi \cdot D_{V,\boldsymbol{h}} \Big\|_{\mathrm{op}} .$$

From the above discussion, the following lemma is immediate.

Lemma A.7. *Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. Then $\Delta_A(R) \leq \Delta_A^{\mathsf{Spectral}}(R)$.*

In the worst case, this relaxation can be quite poor. The following is an example in which the relaxation is equal to $\sqrt{N} - 1$, even though the distinguishing value $\Delta_A(R)$ can never be more than one.

Example A.8 (A large relaxation value)**.** For this example, we will view the space $\mathbb{C}^N$ as corresponding to n qubits, so that the standard basis contains the vector $|x\rangle$ for each $x \in \{0,1\}^n$. With this viewpoint, a binary phase state is specified by a Boolean function $h : \{0,1\}^n \to \{\pm 1\}$ and is given by

$$|\psi_h\rangle = \frac{1}{\sqrt{N}} \sum_{x \in \{0,1\}^n} h(x) \, |x\rangle \, .$$

Suppose the adversary's strategy does not expand the Hilbert space (so that $M = N$). In addition, suppose that the isometry V is just the n-qubit Hadamard transform $V = H^{\otimes n}$, and that the measurement Π is just the n-qubit identity matrix $\Pi = I_{N \times N}$. In this case, the rows of V are just the binary phase states $|\psi_{\chi_\alpha}\rangle$, where $\chi_\alpha : \{0,1\}^n \to \{\pm 1\}$ is the Boolean function $\chi_\alpha(x) = (-1)^{\langle \alpha, x \rangle}$; in other words,

$$V = \sum_{\alpha \in \{0,1\}^n} |\alpha\rangle\langle\psi_{\chi_\alpha}| \, .$$

As a result, the weight $\mathsf{wt}_{V,\alpha} = 1/N$ for all $\alpha \in \{0,1\}^n$, and so the rescaling matrix is given by

$$R_{V,h} = \sum_{\alpha \in \{0,1\}^n} \sqrt{N} \cdot \langle \psi_{\chi_\alpha} | \psi_h \rangle \cdot |\alpha\rangle\langle\alpha| \,.$$

Now we compute the two terms in the spectral relaxation. The second term is independent of the function family R, and so we compute it first:

$$\mathop{\mathbf{E}}_{\boldsymbol{h}} D^\dagger_{V,\boldsymbol{h}} \cdot \Pi \cdot D_{V,\boldsymbol{h}} = \mathop{\mathbf{E}}_{\boldsymbol{h}} D^\dagger_{V,\boldsymbol{h}} \cdot D_{V,\boldsymbol{h}} = \mathop{\mathbf{E}}_{\boldsymbol{h}} \sum_{\alpha \in \{0,1\}^n} N \cdot |\, \langle \psi_{\chi_\alpha} | \psi_{\boldsymbol{h}} \rangle \,|^2 \cdot |\alpha\rangle\langle\alpha| = I_{N\times N},$$

where the last equality used the fact that $\mathbf{E}_{\boldsymbol{h}} \, \langle \psi_{\chi_\alpha} | \psi_{\boldsymbol{h}} \rangle \,|^2 = \mathsf{wt}_{V,\alpha} = N$ due to Proposition A.2. As for the first term, consider a worst-case function family R in which every R_k is equal to the same parity $R_k = \chi_\alpha$, for some $\alpha \in \{0,1\}^n$. Then for every $1 \le k \le K$, the rescaling matrix is given by $D_{V,R_k} = \sqrt{N} \cdot |\alpha\rangle\langle\alpha|$. As a result,

$$\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D^\dagger_{V,R_{\boldsymbol{k}}} \cdot \Pi \cdot D_{V,R_{\boldsymbol{k}}} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D^\dagger_{V,\boldsymbol{h}} \cdot \Pi \cdot D_{V,\boldsymbol{h}} = \sqrt{N} \cdot |\alpha\rangle\langle\alpha| - I_{N\times N}.$$

The operator norm of this matrix is $\sqrt{N} - 1$, and so $\Delta_A^{\mathsf{Spectral}}(R) = \sqrt{N} - 1$.

The reason that this example has such a large relaxation value is that the rescaling matrices D_{V,R_k} all have an extremely large diagonal entry and therefore an extremely large operator norm. We would like to rule out examples like this by only considering function families R in which the rescaling matrices have operator norms which are not too much larger than 1. This motivates the following definition.

Definition A.9 (B-bounded function families)**.** A function $h : [N] \to \{\pm 1\}$ is *B-bounded* if $|D_{V,h,i}| \le B$ for all $1 \le i \le M$. In addition, a function family $R : [K] \times [N] \to \{\pm 1\}$ is *B-bounded* if R_k is B-bounded for all $1 \le k \le K$.

Example A.8 showed that there exist worst-case function families R which are not B-bounded for small values of B. However, the next lemma shows that an average-case function family will in fact be B-bounded with extremely high probability.

Lemma A.10 (Random function families are bounded)**.** *Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Then*

$$\Pr_{\boldsymbol{R}}[\boldsymbol{R} \text{ is not } B\text{-bounded}] \le 4KM \cdot e^{-B^2/4}.$$

This lemma follows as a simple corollary of the following lemma by applying it to each function $\boldsymbol{R}_k$ separately and then union bounding over all $1 \le k \le K$.

Lemma A.11 (Random functions are bounded)**.** *Let $\boldsymbol{h} : [N] \to \{\pm 1\}$ be a uniformly random Boolean function. Then*

$$\Pr_{\boldsymbol{h}}[\boldsymbol{h} \text{ is not } B\text{-bounded}] \le 2M \cdot e^{-B^2/2}.$$

The key technical tool in the proof of this lemma is Theorem A.5

Proof of Lemma A.10. Fix a $1 \le i \le M$, and let us consider $D_{V,\boldsymbol{h},i}$. By definition,

$$D_{V,\boldsymbol{h},i} = \frac{1}{\sqrt{\mathsf{wt}_{V,i}}} \cdot \langle v_i | \psi_{\boldsymbol{h}} \rangle = \frac{1}{\sqrt{\mathsf{wt}_{V,i}}} \cdot \sum_{x=1}^{N} v_{i,x} \cdot \tfrac{1}{\sqrt{N}} \boldsymbol{h}(x) = \sum_{x=1}^{N} a_x \cdot \boldsymbol{h}(x).$$

where

$$a_x = \frac{1}{\sqrt{\mathsf{wt}_{V,i} \cdot N}} \cdot v_{i,x}.$$

Note that

$$\sum_{x=1}^{N} |a_x|^2 = \sum_{x=1}^{N} \frac{1}{\mathsf{wt}_{v,i} \cdot N} \cdot |v_{i,x}|^2 = \frac{1}{\mathsf{wt}_{v,i} \cdot N} \cdot \langle v_i | v_i \rangle = \frac{1}{\mathsf{wt}_{v,i}} \cdot \mathsf{wt}_{v,i} = 1.$$

As a result, Theorem A.5 says that

$$\mathbf{Pr}[|D_{V,\boldsymbol{h},i}| \geq B] \leq 2 \cdot \exp\left(-\frac{B^2}{2 \cdot \sum_{x=1}^{N} |a_x|^2}\right) = 2 \cdot \exp\left(-\frac{B^2}{2}\right).$$

Union bounding over all $1 \leq i \leq M$, we have that

$$\mathop{\mathbf{Pr}}_{\boldsymbol{h}}[\boldsymbol{h} \text{ is not } B\text{-bounded}] \leq \sum_{i=1}^{M} \mathbf{Pr}[|D_{V,\boldsymbol{h},i}| \geq B] \leq \sum_{i=1}^{M} 2 \cdot e^{-B^2/2} = 2M \cdot e^{-B^2/2}.$$

This completes the proof. □

A.2 Truncating the spectral relaxation

Although Lemma A.10 shows that the overwhelming majority of function families R are B-bounded, it will still be convenient to modify the spectral relaxation slightly so that the rare bad events do not lead to extremely large values, as in Example A.8. We will handle this by truncation.

Definition A.12 (B-truncation)**.** Let $\mathrm{trunc}_B : \mathbb{C} \to \mathbb{C}$ be the function which, on input $t \in \mathbb{C}$, acts as follows:

$$\mathrm{trunc}_B(t) = \begin{cases} t & \text{if } |t| \leq B \\ B \cdot t/|t| & \text{if } |t| > B. \end{cases}$$

By design, $|\mathrm{trunc}_B(t)| \leq B$ for all $t \in \mathbb{C}$. Now we use this to define a truncated version of the rescaling matrix.

Definition A.13 (B-truncated rescaling matrix)**.** Let $h : [N] \to \{\pm 1\}$ be a Boolean function. Then the *B-truncated rescaling matrix* is the diagonal matrix given by

$$D_{V,h,B} = \sum_{i=1}^{M} \mathrm{trunc}_B(D_{V,h,i}) \cdot |i\rangle\langle i|.$$

With this in hand, we can define truncated analogues of the distinguishing advantage and the spectral relaxation.

Definition A.14 (B-truncated advantage and spectral relaxation)**.** Let $R : [K] \times [N] \to \{\pm 1\}$ be a function family. Let $h : [N] \to \{\pm 1\}$ be a Boolean function. Given an oracle $\mathcal{O}$, the *B-truncated acceptance probability* is defined as

$$p_{A,B}(h \mid f) = \langle \mathsf{wt}_V | \cdot \mathcal{O}_f \cdot D_{V,h,B}^{\dagger} \cdot \Pi \cdot D_{V,h,B} \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle$$

In addition, the *B-truncated distinguishing advantage* and *B-truncated spectral relaxation* are defined as follows.

$$\begin{aligned}\Delta_{A,B}(R) &= \max_f \Big| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} p_{A,B}(R_{\boldsymbol{k}} \mid f) - \mathop{\mathbf{E}}_{\boldsymbol{h}} p_{A,B}(\boldsymbol{h} \mid f) \Big| \\ &= \max_f \Big| \langle \mathsf{wt}_V | \cdot \mathcal{O}_f \cdot \Big(\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D^\dagger_{V,R_{\boldsymbol{k}},B} \cdot \Pi \cdot D_{V,R_{\boldsymbol{k}},B} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D^\dagger_{V,\boldsymbol{h},B} \cdot \Pi \cdot D_{V,\boldsymbol{h},B} \Big) \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle \Big|, \\ \Delta^{\mathsf{Spectral}}_{A,B}(R) &= \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D^\dagger_{V,R_{\boldsymbol{k}},B} \cdot \Pi \cdot D_{V,R_{\boldsymbol{k}},B} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D^\dagger_{V,\boldsymbol{h},B} \cdot \Pi \cdot D_{V,\boldsymbol{h},B} \Big\|_{\mathrm{op}}.\end{aligned}$$

Note that the B-truncated spectral relaxation remains a spectral relaxation of the B-truncated distinguishing advantage, in that $\Delta_{A,B}(R) \le \Delta^{\mathsf{Spectral}}_{A,B}(R)$.

As we have seen, a random function family is B-bounded with overwhelming probability. This suggests that the B-truncated analogue of the distinguishing advantage should not be too far from the regular distinguishing advantage. This is shown in the next lemma.

Lemma A.15 (Truncating doesn't change the distinguishing advantage by much)**.** *Let $R : [K] \times [N] \to \{\pm 1\}$ be a B-bounded function family. Then*

$$\Delta_A(R) \le \Delta_{A,B}(R) + 4M \cdot e^{-B^2/2}.$$

Before proving this, we will need to establish the following technical lemma.

Lemma A.16. *For any function family $R : [K] \times [N] \to \{\pm 1\}$, any function $h : [N] \to \{\pm 1\}$, and any function f,*

$$|p_A(h \mid f) - p_{A,B}(h \mid f)| \le 2.$$

Proof. The first step of the proof is a simple triangle inequality:

$$\begin{aligned}|p_A(h \mid f) - p_{A,B}(h \mid f)| &\le |p_A(h \mid f)| + |p_{A,B}(h \mid f)| \\ &\le 1 + |p_{A,B}(h \mid f)| \\ &= 1 + |\langle \mathsf{wt}_V | \cdot \mathcal{O}_f \cdot D^\dagger_{V,h,B} \cdot \Pi \cdot D_{V,h,B} \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle|,\end{aligned}$$

where the second inequality is because $p_A(h \mid f)$ is an acceptance probability and therefore at most 1. As for the second term, we note that

$$D_{V,h,B} \cdot \mathcal{O}_f \cdot |\mathsf{wt}_V\rangle = \mathcal{O}_f \cdot D_{V,h,B} \cdot |\mathsf{wt}_V\rangle$$

because $D_{V,h,B}$ and $\mathcal{O}_f$ are diagonal matrices. Expanding $D_{V,h,B} \cdot |\mathsf{wt}_V\rangle$,

$$\begin{aligned}D_{V,h,B} \cdot |\mathsf{wt}_V\rangle &= D_{V,h,B} \cdot \Big(\sum_{i=1}^M \sqrt{\mathsf{wt}_{V,i}}\, |i\rangle \Big) \\ &= \sum_{i=1}^M \mathrm{trunc}_B(D_{V,h,i}) \cdot \sqrt{\mathsf{wt}_{V,i}}\, |i\rangle \\ &= \sum_{i=1}^M \mathrm{trunc}_{B\cdot\sqrt{\mathsf{wt}_{V,i}}}(D_{V,h,i} \cdot \sqrt{\mathsf{wt}_{V,i}}) \cdot |i\rangle \\ &= \sum_{i=1}^M \mathrm{trunc}_{B\cdot\sqrt{\mathsf{wt}_{V,i}}}(\langle v_i | \psi_h \rangle) \cdot |i\rangle\,.\end{aligned}$$

We note that

$$|\mathrm{trunc}_{B\cdot\sqrt{\mathsf{wt}_{V,i}}}(\langle v_i|\psi_h\rangle)| \le |\langle v_i|\psi_h\rangle|,$$

which is the amplitude on $|i\rangle$ in the state $V\cdot|\psi_h\rangle$ due to Equation (26). As a result, $D_{V,h,B}\cdot|\mathsf{wt}_V\rangle$ is a subnormalized vector, and therefore so is $\mathcal{O}_f\cdot D_{V,h,B}\cdot|\mathsf{wt}_V\rangle$ because $\mathcal{O}_f$ is a unitary matrix. Putting everything together, this tells us that

$$|\langle \mathsf{wt}_V|\cdot\mathcal{O}_f\cdot D^\dagger_{V,h,B}\cdot\Pi\cdot D_{V,h,B}\cdot\mathcal{O}_f\cdot|\mathsf{wt}_V\rangle| \le 1,$$

because $0\preceq\Pi\preceq\mathsf{Id}$. Thus, the sum of the two terms is at most 2. □

Now we use this to prove Lemma A.15.

Proof of Lemma A.15. By definition,

$$\begin{aligned}\Delta_A(R) &= \max_f\left|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} p_A(R_{\boldsymbol{k}}\mid f) - \mathop{\mathbf{E}}_{\boldsymbol{h}} p_A(\boldsymbol{h}\mid f)\right|\\ &= \max_f\left|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} p_{A,B}(R_{\boldsymbol{k}}\mid f) - \mathop{\mathbf{E}}_{\boldsymbol{h}} p_A(\boldsymbol{h}\mid f)\right|,\end{aligned}$$

where the second equality holds because R is B-bounded, so $p_A(R_{\boldsymbol{k}}\mid f) = p_{A,B}(R_{\boldsymbol{k}}\mid f)$ for every $1\le k\le K$. By the triangle inequality, this can be upper-bounded by

$$\begin{aligned}&\max_f\left|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} p_{A,B}(R_{\boldsymbol{k}}\mid f) - \mathop{\mathbf{E}}_{\boldsymbol{h}} p_{A,B}(\boldsymbol{h}\mid f)\right| + \max_f\left|\mathop{\mathbf{E}}_{\boldsymbol{h}} p_{A,B}(\boldsymbol{h}\mid f) - \mathop{\mathbf{E}}_{\boldsymbol{h}} p_A(\boldsymbol{h}\mid f)\right|\\ &= \Delta_{A,B}(R) + \max_f\left|\mathop{\mathbf{E}}_{\boldsymbol{h}} p_{A,B}(\boldsymbol{h}\mid f) - \mathop{\mathbf{E}}_{\boldsymbol{h}} p_A(\boldsymbol{h}\mid f)\right|\\ &\le \Delta_{A,B}(R) + \max_f \mathop{\mathbf{E}}_{\boldsymbol{h}}\left|p_{A,B}(\boldsymbol{h}\mid f) - p_A(\boldsymbol{h}\mid f)\right|.\end{aligned}$$

Let us first focus on the second term. If $\boldsymbol{h}$ is B-bounded, then $p_{A,B}(\boldsymbol{h}\mid f) = p_A(\boldsymbol{h}\mid f)$, and the term inside the expectation is zero. Otherwise the term inside the expectation is at most 2, by Lemma A.16. As a result, for every function f, the expectation is at most

$$2\cdot\Pr_{\boldsymbol{h}}[\boldsymbol{h}\text{ is not }B\text{-bounded}] \le 4M\cdot e^{-B^2/2},$$

by Lemma A.11. Putting everything together,

$$\Delta_A(R) \le \Delta_{A,B}(R) + 4M\cdot e^{-B^2/2}.$$

This completes the proof. □

A.3 The one-query lower bound

Now we complete the proof of the one-query lower bound. The quantitative bound we prove is as follows.

Theorem A.17. *Let $\boldsymbol{R}:[K]\times[N]\to\{\pm1\}$ be a uniformly random function family. Then*

$$\Pr_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \ge 1/K^{1/4} + 4M\cdot e^{-K^{1/8}/2}] \le 6KM\cdot e^{-K^{1/8}/32}.$$

We note that although this bound is quantitatively weaker than Theorem 4.18, it still gives a strong lower bound. For our typical settings of K and M, it states that $\Delta_A(\boldsymbol{R})$ is roughly bounded by $1/K^{1/4}$ with all but a negligible probability. The key technical result we use to prove this is the following variant of the matrix Chernoff bound, stated in [Tro12, Theorem 1.3].

Theorem A.18 (Matrix Hoeffding)**.** *Let $\{\boldsymbol{Z}_k\}$ be a set of independent, Hermitian, random matrices with dimension D. Let $\{C_k\}$ be a set of fixed Hermitian matrices. Assume that for all k, $\mathbf{E}\,\boldsymbol{Z}_k = 0$ and $\boldsymbol{Z}_k^2 \preceq C_k^2$. Set*

$$\sigma^2 = \Big\| \sum_k C_k^2 \Big\|.$$

Then

$$\mathbf{Pr}\left[\lambda_{\max}\Big(\sum_k \boldsymbol{Z}_k\Big) \geq t\right] \leq D \cdot e^{-t^2/8\sigma^2}.$$

By applying matrix Hoeffding to both $(\sum_k \boldsymbol{Z}_k)$ and $-(\sum_k \boldsymbol{Z}_k)$, we can derive the following concentration bound for the operator norm.

Corollary A.19 (Matrix Hoeffding for operator norm)**.** *Under the same assumptions as Theorem A.18, we have that*

$$\mathbf{Pr}\left[\Big\| \sum_k \boldsymbol{Z}_k \Big\|_{\mathrm{op}} \geq t\right] \leq 2D \cdot e^{-t^2/8\sigma^2}.$$

Proof of Theorem A.17. Let $B \geq 0$ be a nonnegative real number to be determined later. By Lemma A.15, we have that

$$\Delta_A(R) \leq \Delta_{A,B}(R) + 4M \cdot e^{-B^2/2}.$$

for any B-bounded function family R. This means that if $\Delta_A(R) \geq \varepsilon + 4M \cdot e^{-B^2/2}$ for some number ε to be determined later, it must either be the case that $\Delta_{A,B}(R) \geq \varepsilon$ or that R is not B-bounded. Hence, by the union bound, if $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is a random function family,

$$\begin{aligned}\mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \geq \varepsilon + 4M \cdot e^{-B^2/2}] &\leq \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,B}(\boldsymbol{R}) \geq \varepsilon] + \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\boldsymbol{R} \text{ is not } B\text{-bounded}] \\ &\leq \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A,B}^{\mathsf{Spectral}}(\boldsymbol{R}) \geq \varepsilon] + 4KM \cdot e^{-B^2/2},\end{aligned}$$

where the second inequality is due to Lemma A.10 and the fact that $\Delta_{A,B}(\boldsymbol{R}) \leq \Delta_{A,B}^{\mathsf{Spectral}}(\boldsymbol{R})$. We will now focus on bounding the first term. By definition of the B-bounded spectral relaxation,

$$\Delta_{A,B}^{\mathsf{Spectral}}(\boldsymbol{R}) = \Big\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D_{V,\boldsymbol{R}_{\boldsymbol{k}},B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{R}_{\boldsymbol{k}},B} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D_{V,\boldsymbol{h},B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{h},B} \Big\|_{\mathrm{op}}.$$

To analyze this, we note that for each $1 \leq k \leq K$, $\boldsymbol{R}_k$ is distributed as a uniformly random function, and so $\boldsymbol{R}_k$ has the same distribution as $\boldsymbol{h}$. Hence, if we keep k fixed and randomize over $\boldsymbol{R}$,

$$\mathop{\mathbf{E}}_{\boldsymbol{R}} D_{V,\boldsymbol{R}_k,B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B} = \mathop{\mathbf{E}}_{\boldsymbol{h}} D_{V,\boldsymbol{h},B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{h},B}.$$

This means that the random matrix

$$X_{\boldsymbol{R}_k} \coloneqq D_{V,\boldsymbol{R}_k,B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D_{V,\boldsymbol{h},B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{h},B}$$

has $\mathbf{E}_{\boldsymbol{R}}[X_{\boldsymbol{R}_k}] = 0$. In terms of these matrices, our goal is to bound

$$\| \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]} D_{V,\boldsymbol{R}_{\boldsymbol{k}},B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{R}_{\boldsymbol{k}},B} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D_{V,\boldsymbol{h},B}^{\dagger} \cdot \Pi \cdot D_{V,\boldsymbol{h},B} \| = \| \mathop{\mathbf{E}}_{\boldsymbol{k}} X_{\boldsymbol{R}_{\boldsymbol{k}}} \|.$$

Note the following properties of $X_{\boldsymbol{R}_k}$:

1. For each k, $X_{\boldsymbol{R}_k}$ only depends on $\boldsymbol{R}_k$. Hence, the random variables $X_{\boldsymbol{R}_k}$, over all $1 \le k \le K$, are independent and identically distributed.

2. $X_{\boldsymbol{R}_k}$ is an $M \times M$ matrix.

3. To bound the operator norm of $X_{\boldsymbol{R}_k}$, we begin with the bound

$$\begin{aligned}\|X_{\boldsymbol{R}_k}\| &= \|D^\dagger_{V,\boldsymbol{R}_k,B} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B} - \mathop{\mathbf{E}}_{\boldsymbol{h}} D^\dagger_{V,\boldsymbol{h},B} \cdot \Pi \cdot D_{V,\boldsymbol{h},B}\| \\ &\le \|D^\dagger_{V,\boldsymbol{R}_k,B} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B}\| + \mathop{\mathbf{E}}_{\boldsymbol{h}} \|D^\dagger_{V,\boldsymbol{h},B} \cdot \Pi \cdot D_{V,\boldsymbol{h},B}\|.\end{aligned}$$

   Now, let us bound the operator norm of $D^\dagger_{V,\boldsymbol{R}_k,B} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B}$. Let $|v\rangle$ be any unit vector. Then because $D_{V,\boldsymbol{R}_k,B}$ is a diagonal matrix whose diagonal entries have magnitude at most B, $D_{V,\boldsymbol{R}_k,B} \cdot |v\rangle$ has norm at most B. Hence,

$$\langle v| \cdot D^\dagger_{V,\boldsymbol{R}_k,B} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B} \cdot |v\rangle \le B^2.$$

   As a result, $D^\dagger_{V,\boldsymbol{R}_k,B} \cdot \Pi \cdot D_{V,\boldsymbol{R}_k,B}$ has spectral norm at most B^2; a similar argument will show that $D^\dagger_{V,\boldsymbol{h},B} \cdot \Pi \cdot D_{V,\boldsymbol{h},B}$ has spectral norm at most B^2 as well. Thus, the spectral norm of $X_{\boldsymbol{R}_k}$ is at most $2B^2$, and so, $X^2_{\boldsymbol{R}_k} \le 4B^4 \cdot \mathsf{Id}$ always.

Now we are in a good place to apply matrix Hoeffding. In our setting, the matrices $X_{\boldsymbol{R}_1}, \dots, X_{\boldsymbol{R}_K}$ are independent, Hermitian, and have dimension M. Furthermore, we know that $X^2_{\boldsymbol{R}_k} \le 4B^4 \cdot \mathsf{Id}$ always. Hence, our value of σ^2 is

$$\sigma^2 = \Big\| \sum_{k=1}^K 4B^4 \mathsf{Id} \Big\| = 4KB^4.$$

Now, our goal is to bound

$$\begin{aligned}\Pr_{\boldsymbol{R}}[\Delta^{\mathsf{Spectral}}_{A,B}(\boldsymbol{R}) \ge \varepsilon] &= \Pr_{\boldsymbol{R}}\Big[\Big\| \mathop{\mathbf{E}}_{\boldsymbol{k}\in[K]} X_{\boldsymbol{R}_{\boldsymbol{k}}} \Big\|_{\mathrm{op}} \ge \varepsilon\Big] \\ &= \Pr_{\boldsymbol{R}}\Big[\Big\| \frac{1}{K} \sum_{k=1}^K X_{\boldsymbol{R}_k} \Big\|_{\mathrm{op}} \ge \varepsilon\Big] = \Pr_{\boldsymbol{R}}\Big[\Big\| \sum_{k=1}^K X_{\boldsymbol{R}_k} \Big\|_{\mathrm{op}} \ge \varepsilon K\Big].\end{aligned}$$

This we can apply Corollary A.19 to, which tells us that

$$\Pr_{\boldsymbol{R}}[\Delta^{\mathsf{Spectral}}_{A,B}(\boldsymbol{R}) \ge \varepsilon] \le 2M \cdot e^{-\varepsilon^2 K^2/8(4KB^4)} = 2M \cdot e^{-\varepsilon^2 K/(32B^4)}.$$

Putting everything together, we have

$$\Pr_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \ge \varepsilon + 4M \cdot e^{-B^2/2}] \le 2M \cdot e^{-\varepsilon^2 K/(32B^4)} + 4KM \cdot e^{-B^2/2}.$$

Now we select our constants to be $\varepsilon = 1/K^{1/4}$ and $B = K^{1/16}$. Then this states that

$$\Pr_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \ge 1/K^{1/4} + 4M \cdot e^{-K^{1/8}/2}] \le 2M \cdot e^{-K^{1/4}/32} + 4KM \cdot e^{-K^{1/8}/2} \le 6KM \cdot e^{-K^{1/8}/32}.$$

This completes the proof. □

B On the power of counting arguments

In this section, we will consider the power of *counting arguments* to show lower bounds for the Oracle State Distinguishing Game. Counting arguments apply to the case when the adversary cannot compute too many unitaries, which we formalize as follows.

Definition B.1 (Small oracle circuits)**.** An oracle circuit $A^{(\cdot)}$ is *S-small* if the number of distinct unitaries A^f, ranging over all oracles f, is at most S.

Although most oracle circuits are not "small" enough to be useful, there are a few interesting families of small oracle circuits, which we state below. As always, we will write $N := 2^n$ for the size of the oracle circuit's input register, and $M := 2^m$ for the total size of the oracle circuit's registers.

Example B.2 (Aaronson-Kuperberg adversaries)**.** As described in Section 1.3, Aaronson and Kuperberg [AK07] considered oracle circuits $A^{(\cdot)}$ in which for every oracle f, A^f exactly computes some unitary transformation on its first n qubits. They showed that any such oracle circuit is 4^N-small [AK07, Theorem 6.7].

Example B.3 (Adversaries with no ancillas but multiple oracles)**.** Consider an oracle circuit with n input qubits and no ancilla qubits which makes t queries. For this example only, we will depart from our usual notation and allow the queries to be made to t potentially different functions $f_1, \ldots, f_t$. Then there are at most 2^N choices for each function f_i, and so this oracle circuit is $(2^N)^t = 2^{Nt}$-small. If $t = \mathrm{poly}(n)$, then this is $2^{N \cdot \mathrm{poly}(n)}$-small.

Example B.4 (Small-ancilla adversaries)**.** Let $A^{(\cdot)}$ be an oracle circuit which makes multiple queries to a single function f and uses n input qubits and a ancilla qubits, for a total of $m = n + a$ qubits. Then $A^{(\cdot)}$ is 2^M-small.

Note that the bound in Example B.4 subsumes Example B.3. This is because by Remark 3.5, an adversary which makes t queries to different functions $f_1, \ldots, f_t$ can be simulated by an oracle circuit $A^{(\cdot)}$ which uses $\lceil \log_2 t \rceil$ additional ancilla qubits and queries a single function f.

Now we state our main bound, which rules out adversaries for the Oracle State Distinguishing Game which are "small".

Theorem B.5 (Counting bound)**.** *There is a universal constant $c > 0$ such that the following is true. Consider the Oracle State Distinguishing Game played with a uniformly random function family $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$. Let $A^{(\cdot)}$ be an adversary which is S-small, for $S = \exp\big(c \cdot \varepsilon^2 KN\big)$. Then*

$$\Pr_{\boldsymbol{R}}\Big[\Delta_A(\boldsymbol{R}) \geq \varepsilon\Big] \leq 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big).$$

In the context of our examples, this rules out Aaronson-Kuperberg adversaries, so long as $K = \Omega(1/\varepsilon^2)$. This also rules out small-ancilla adversaries. For example, if adversary uses $a = \frac{1}{2} \cdot \log_2(K)$ ancilla qubits, then its total size is $M = N \cdot \sqrt{K}$, and so it is $2^{N\sqrt{K}}$-small, which is small enough (assuming reasonable settings of parameters) for Theorem B.5 to apply. On the other hand, Theorem B.5 cannot rule out general adversaries which use $a = \log_2(K)$ ancilla qubits or more. For example, a natural adversary might intend to perform the query $(k, x) \mapsto R(k, x)$ in superposition, and to do so it needs $\log_2(K) + n$ qubits, putting it in the range where Theorem B.5 no longer applies. This shows the limitation of this style of counting argument: it becomes ineffective once the adversary has even a small number of ancilla qubits.

Now we prove Theorem B.5. We will do so using a standard "concentration and union bound" approach: we prove a tail inequality on the probability that A^f results in a good attack for a fixed f, and the we union bound over all f.

Lemma B.6 (Success probability of no-query adversaries)**.** *There is a universal constant $c > 0$ such that the following is true. Let $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ be a uniformly random function family. Let $A^{(\cdot)}$ be an adversary that does not make any queries. Then*

$$\mathop{\mathbf{Pr}}_{\boldsymbol{R}}\Big[\Delta_A(\boldsymbol{R}) \geq \varepsilon\Big] \leq 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big).$$

Proof. We prove this using tools from Section 3.4. Since $A^{()}$ makes no queries, let us fix an arbitrary f and note that $\Delta_A(R) = \Delta_A(R \mid f)$ for all function families R. Then

$$\Delta_A(R \mid f) = \Big|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_A(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_A(\boldsymbol{h} \mid f)]\Big| = \Big|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_{A,0}(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_{A,0}(\boldsymbol{h} \mid f)]\Big|, \tag{34}$$

where $p_{A,0}(\cdot \mid f)$ is the extension of $p_A(\cdot \mid f)$ to bounded functions $\underline{R} : [K] \times [N] \to [-1, 1]$ from Definition 3.21. But the function $\underline{R} \mapsto \mathbf{E}_{\boldsymbol{k}\sim[K]}[p_{A,0}(\underline{R}_{\boldsymbol{k}} \mid f)]$ is convex and $(2/\sqrt{KN})$-Lipschitz by Lemma 3.25. Hence, if $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$ is a uniformly random function family, Talagrand's inequality (Theorem 3.20) implies that

$$\mathop{\mathbf{Pr}}_{\boldsymbol{R}}\Big[\Big|\mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[p_{A,0}(\boldsymbol{R}_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}}[p_{A,0}(\boldsymbol{h} \mid f)]\Big| \geq \varepsilon\Big] \leq 2 \cdot \exp\left(-\frac{c \cdot \varepsilon^2 KN}{4}\right)$$

for some absolute constant $c > 0$. But by Equation (34) the left-hand side is $\mathbf{Pr}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \geq \varepsilon]$, and so this completes the proof, with the "c" in the lemma statement being equal to $c/4$. □

Deriving Theorem B.5 from Lemma B.6 is relatively straightforward.

Proof of Theorem B.5. Let $A^{(\cdot)}$ be an S-small adversary, for a value of S to be determined later. Then there exist S functions $f_1, \ldots, f_S$ such that the set of unitaries $A^{f_1}, \ldots, A^{f_S}$ contains every unitary computable by $A^{(\cdot)}$. Fix a $1 \leq i \leq S$. Then by hard-coding the function f_i into $A^{(\cdot)}$, we can view A^{f_i} as an oracle circuit that does not make any queries. Thus, Lemma B.6 says that

$$\mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A^{f_i}}(\boldsymbol{R}) \geq \varepsilon] \leq 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big).$$

As a result, we can upper-bound the maximum distinguishing probability by

$$\begin{aligned}
\mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R}) \geq \varepsilon] &= \mathop{\mathbf{Pr}}_{\boldsymbol{R}}\Big[\max_f\{\Delta_A(\boldsymbol{R} \mid f)\} \geq \varepsilon\Big] \\
&= \mathop{\mathbf{Pr}}_{\boldsymbol{R}}\Big[\max_i\{\Delta_A(\boldsymbol{R} \mid f_i)\} \geq \varepsilon\Big] \\
&\leq \sum_{i=1}^{S} \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_A(\boldsymbol{R} \mid f_i) \geq \varepsilon] && \text{(by the union bound)} \\
&= \sum_{i=1}^{S} \mathop{\mathbf{Pr}}_{\boldsymbol{R}}[\Delta_{A^{f_i}}(\boldsymbol{R}) \geq \varepsilon] \\
&\leq \sum_{i=1}^{S} 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big) = S \cdot 2 \cdot \exp\big(-c \cdot \varepsilon^2 KN\big).
\end{aligned}$$

Now, let us choose S to be $S = \exp\big(c/2 \cdot \varepsilon^2 KN\big)$. Then this upper bound on the maximum distinguishing probability equals $S \cdot 2 \cdot \exp\big(-c/2 \cdot \varepsilon^2 KN\big)$. This concludes the proof, with the constant "c" in the statement of the proof equal to $c/2$. □

C A one-query attack with advantage $\Omega(1/\sqrt{K})$

In this section, we give a one-query adversary for the Oracle State Distinguishing Game achieving advantage $\Omega(1/\sqrt{K})$ using only one ancilla qubit. This demonstrates that the dependence on K in our main theorem (Theorem 4.18) is tight. The adversary is given as follows.

Definition C.1 (Hadamard adversary)**.** On input an n-qubit state $|\psi\rangle$, the *Hadamard adversary* $A^{(\cdot)}_{\mathrm{Had}}$ acts as follows.

1. Apply the n-qubit Hadamard $H^{\otimes n}$ to $|\psi\rangle$.
2. Measure $H^{\otimes n} \cdot |\psi\rangle$ in the standard basis. Let $\boldsymbol{y} \in \{0,1\}^n$ be the measurement outcome.
3. Query a bit flip oracle $f : \{0,1\}^n \to \{0,1\}$ on $\boldsymbol{y}$. Let $\boldsymbol{b}' = f(\boldsymbol{y}) \in \{0,1\}$ be the result.
4. Output $\boldsymbol{b}'$.

The Hadamard's adversary's one ancilla qubit is used to store the outcome of the query to the bit flip oracle. Note that this can be simulated by an adversary which makes a query to a phase oracle instead, as discussed following the statement of Definition 3.3. We also remark that because the Hadamard adversary applies an n-qubit Hadamard, it will be more convenient to think of the adversary's state space as consisting of n-qubits, rather than being a single space of overall dimension $N := 2^n$. As a result, using the correspondence between $\{0,1\}^n$ and $[N]$ mentioned in Notation 3.1, we will prefer to format our function families as $R : [K] \times \{0,1\}^n \to \{\pm 1\}$, with the k-th binary phase state being

$$|\psi_{R_k}\rangle = \frac{1}{\sqrt{N}} \cdot \sum_{x \in \{0,1\}^n} R_k(x) \cdot |x\rangle \,.$$

Our main goal is to prove the following bounds on the Hadamard adversary's distinguishing probability.

Theorem C.2 (Distinguishing advantage of the Hadamard adversary)**.** *There exists a constant $c > 0$ such that the following is true. Let $K, N \geq c$, and let $\boldsymbol{R} : [K] \times \{0,1\}^n \to \{\pm 1\}$ be a uniformly random function family. Then*

$$\mathop{\mathbf{E}}_{\boldsymbol{R}}[\Delta_{A_{\mathrm{Had}}}(\boldsymbol{R})] \geq \Omega\Big(\frac{1}{\sqrt{K}}\Big).$$

When the Oracle State Distinguishing Game is played with some function family $R : [K] \times [N] \to \{\pm 1\}$, with probability $\frac{1}{2}$ the Hadamard adversary is given the state $|\psi_{R_{\boldsymbol{k}}}\rangle$ for $\boldsymbol{k}$ chosen uniformly at random. We will write $\mathcal{M}_R$ for the probability distribution on the measurement outcome $\boldsymbol{y}$ in this case. In other words,

$$\mathcal{M}_R(y) := \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} |\, \langle y| \cdot H^{\otimes n} \cdot |\psi_{R_{\boldsymbol{k}}}\rangle \,|^2.$$

With the remaining $\frac{1}{2}$ probability, the Hadamard adversary is given a uniformly random phase state; equivalently, it is given the maximally mixed state Id_N/N. In this case, the measurement outcome $\boldsymbol{y}$ is distributed as a uniformly random string in $\{0,1\}^n$. We will write $\mathcal{U}_N$ for this uniform probability distribution, i.e. $\mathcal{U}_N(y) := 1/N$.

The Hadamard adversary measures a $\boldsymbol{y}$ which is sampled either from $\mathcal{M}_R$ or $\mathcal{U}_N$, and it feeds $\boldsymbol{y}$ into the function f, which can be thought of as a statistical test to distinguish these two distributions. The following lemma characterizes the Hadamard adversary's distinguishing advantage in terms of the *total variation distance* $\mathrm{d_{TV}}(\cdot,\cdot)$ between these two distributions.

Lemma C.3 (Distinguishing advantage equals TV distance)**.**

$$\Delta_{A_{\mathrm{Had}}}(R) = \mathrm{d}_{\mathrm{TV}}(\mathcal{M}_R, \mathcal{U}_N) = \frac{1}{2} \cdot \sum_{y \in \{0,1\}^n} |\mathcal{M}_R(y) - \mathcal{U}_N(y)| \,.$$

Proof. Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. By definition,

$$\begin{aligned}
\Delta_{A_{\mathrm{Had}}}(R \mid f) &= \Big| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} [p_{A_{\mathrm{Had}}}(R_{\boldsymbol{k}} \mid f)] - \mathop{\mathbf{E}}_{\boldsymbol{h}} [p_{A_{\mathrm{Had}}}(\boldsymbol{h} \mid f)] \Big| \\
&= \Big| \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \mathbf{Pr}[A^f_{\mathrm{Had}} \text{ outputs “0” on } |\psi_{R_{\boldsymbol{k}}}\rangle] - \mathop{\mathbf{E}}_{\boldsymbol{h}} \mathbf{Pr}[A^f_{\mathrm{Had}} \text{ outputs “0” on } |\psi_{\boldsymbol{h}}\rangle] \Big| \\
&= \Big| \sum_{y : f(y) = 0} \mathcal{M}_R(y) - \sum_{y : f(y) = 0} \mathcal{U}_N(y) \Big| .
\end{aligned}$$

The maximum distinguishing advantage is then computed by optimizing this expression over all f, but that is exactly the definition of the total variation distance. □

Our goal is to calculate the expectation $\mathbf{E}_{\boldsymbol{R}}[\Delta_{A_{\mathrm{Had}}}(\boldsymbol{R})]$. The following lemma gives an alternative expression for this expectation in terms of a new random variable.

Lemma C.4. *Given a function family $R : [K] \times \{0,1\}^n \to \{\pm 1\}$, define the quantity*

$$X_R \coloneqq \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big(\sum_{x \in \{0,1\}^n} R_{\boldsymbol{k}}(x) \Big)^2 .$$

Then for a uniformly random function family $\boldsymbol{R} : [K] \times [N] \to \{\pm 1\}$,

$$\mathop{\mathbf{E}}_{\boldsymbol{R}} [\Delta_{A_{\mathrm{Had}}}(\boldsymbol{R})] = \frac{1}{2} \cdot \mathop{\mathbf{E}}_{\boldsymbol{R}} |X_{\boldsymbol{R}} - 1| .$$

Proof. By Lemma C.3,

$$\begin{aligned}
\mathop{\mathbf{E}}_{\boldsymbol{R}} [\Delta_{A_{\mathrm{Had}}}(\boldsymbol{R})] = \frac{1}{2} \cdot \sum_{y \in \{0,1\}^n} \mathop{\mathbf{E}}_{\boldsymbol{R}} |\mathcal{M}_{\boldsymbol{R}}(y) - \mathcal{U}_N(y)| &= \frac{1}{2} \cdot \sum_{y \in \{0,1\}^n} \mathop{\mathbf{E}}_{\boldsymbol{R}} \Big| \mathcal{M}_{\boldsymbol{R}}(y) - \frac{1}{N} \Big| \\
&= \frac{1}{2} \cdot \frac{1}{N} \cdot \sum_{y \in \{0,1\}^n} \mathop{\mathbf{E}}_{\boldsymbol{R}} \Big| N \cdot \mathcal{M}_{\boldsymbol{R}}(y) - 1 \Big| . \qquad (35)
\end{aligned}$$

Now, fix a $y \in \{0,1\}^n$. Consider the random variable

$$N \cdot \mathcal{M}_{\boldsymbol{R}}(y) = N \cdot \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \big| \langle y | \cdot H^{\otimes n} \cdot |\psi_{\boldsymbol{R}_{\boldsymbol{k}}}\rangle \big|^2 = \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k} \sim [K]} \Big(\sum_{x \in \{0,1\}^n} \boldsymbol{R}_{\boldsymbol{k}}(x) \cdot (-1)^{x \cdot y} \Big)^2 .$$

For each value of $\boldsymbol{k}$, the corresponding term is distributed as the square of the sum of N independent and uniformly random $\{\pm 1\}$ numbers, and the terms are independent across different values of $\boldsymbol{k}$. Hence, this random variable is distributed identically to $X_{\boldsymbol{R}}$. In addition, these K random variables are independent. As a result, by linearity of expectation,

$$(35) = \frac{1}{2} \cdot \frac{1}{N} \cdot \sum_{y \in \{0,1\}^n} \mathop{\mathbf{E}}_{\boldsymbol{R}} |X_{\boldsymbol{R}} - 1| = \frac{1}{2} \cdot \mathop{\mathbf{E}}_{\boldsymbol{R}} |X_{\boldsymbol{R}} - 1| .$$

This completes the proof. □

Now we study the distribution of the random variable $X_{\boldsymbol{R}}$. To begin, we compute its mean and variance.

Lemma C.5. *Let $\boldsymbol{R} : [K] \times \{0,1\}^n \to \{\pm 1\}$ be a uniformly random function family. Then the random variable $X_{\boldsymbol{R}}$ has expectation 1 and variance $2/K \cdot (N-1)/N$.*

Proof. First, we compute the mean:

$$\begin{aligned}\mathop{\mathbf{E}}_{\boldsymbol{R}}[X_{\boldsymbol{R}}] &= \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{k}}\left[\Big(\sum_{x\in\{0,1\}^n} \boldsymbol{R}_{\boldsymbol{k}}(x)\Big)^2\right] = \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{R},\boldsymbol{k}}\left[\sum_{x,y\in\{0,1\}^n} \boldsymbol{R}_{\boldsymbol{k}}(x)\cdot \boldsymbol{R}(\boldsymbol{k},y)\right] \\ &= \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}\left[\sum_{x,y\in\{0,1\}^n} \mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\boldsymbol{R}_{\boldsymbol{k}}(x)\cdot \boldsymbol{R}_{\boldsymbol{k}}(y)\right]\right] = \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}[N] = 1.\end{aligned}$$

Next, we compute the variance. To begin, note that for each R,

$$X_R - \mathop{\mathbf{E}}_{\boldsymbol{R}}[X_{\boldsymbol{R}}] = X_R - 1 = \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}\Big(\sum_{x\in\{0,1\}^n} R_{\boldsymbol{k}}(x)\Big)^2 - 1 = \frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}\Big(\sum_{x\neq y} R_{\boldsymbol{k}}(x)\cdot R_{\boldsymbol{k}}(y)\Big).$$

Thus, the variance is given by

$$\begin{aligned}\mathop{\mathbf{E}}_{\boldsymbol{R}}[(X_{\boldsymbol{R}} - \mathop{\mathbf{E}}_{\boldsymbol{R}}[X_{\boldsymbol{R}}])^2] &= \mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\Big(\frac{1}{N} \cdot \mathop{\mathbf{E}}_{\boldsymbol{k}\sim[K]}\Big(\sum_{x\neq y} \boldsymbol{R}_{\boldsymbol{k}}(x)\cdot \boldsymbol{R}_{\boldsymbol{k}}(y)\Big)\Big)^2\right] \\ &= \frac{1}{N^2} \cdot \mathop{\mathbf{E}}_{\boldsymbol{R}} \mathop{\mathbf{E}}_{\boldsymbol{k},\boldsymbol{k}'\sim[K]}\left[\sum_{x\neq y}\sum_{z\neq w} \boldsymbol{R}_{\boldsymbol{k}}(x)\cdot \boldsymbol{R}_{\boldsymbol{k}}(y)\cdot \boldsymbol{R}_{\boldsymbol{k}'}(z)\cdot \boldsymbol{R}_{\boldsymbol{k}'}(w)\right] \\ &= \frac{1}{N^2} \cdot \sum_{x\neq y}\sum_{z\neq w} \mathop{\mathbf{E}}_{\boldsymbol{k},\boldsymbol{k}'\sim[K]} \mathop{\mathbf{E}}_{\boldsymbol{R}}\left[\boldsymbol{R}_{\boldsymbol{k}}(x)\cdot \boldsymbol{R}_{\boldsymbol{k}}(y)\cdot \boldsymbol{R}_{\boldsymbol{k}'}(z)\cdot \boldsymbol{R}_{\boldsymbol{k}'}(w)\right]. \qquad (36)\end{aligned}$$

The expectation over $\boldsymbol{R}$ is zero if $\boldsymbol{k} \neq \boldsymbol{k}'$. On the other hand, if $\boldsymbol{k} = \boldsymbol{k}'$, then the expectation is 1 if $\{x,y\} = \{z,w\}$ and 0 otherwise. As a result,

$$(36) = \frac{1}{N^2}\cdot\sum_{x\neq y}\sum_{z\neq w}\frac{1}{K}\cdot\mathbf{1}[\{x,y\}=\{z,w\}] = \frac{1}{N^2}\cdot\frac{1}{K}\cdot\sum_{x\neq y}\sum_{z\neq w}\mathbf{1}[\{x,y\}=\{z,w\}] = \frac{1}{N^2}\cdot\frac{1}{K}\cdot 2N(N-1).$$

This completes the proof. □

From Lemma C.5, we roughly expect that $X_{\boldsymbol{R}}$ tends to be around the values $1 \pm \sqrt{2/K}$. If this were true, then the expectation $\mathbf{E}_{\boldsymbol{R}}|X_{\boldsymbol{R}} - 1|$ we are trying to compute would be roughly $\sqrt{2/K}$, and we would be done. However, it could be that the $X_{\boldsymbol{R}}$'s variance being roughly $2/K$ could be due to some small probability events where $X_{\boldsymbol{R}}$ is very far from 1, whereas with high probability $X_{\boldsymbol{R}}$ is much closer to 1 than $\sqrt{2/K}$. To rule this out, we prove the following concentration bound for $X_{\boldsymbol{R}}$.

Lemma C.6. *There exists a constant $c > 0$ such that the following is true. Let $\boldsymbol{R} : [K]\times\{0,1\}^n \to \{\pm 1\}$ be a uniformly random function family. Then for all $t > 0$,*

$$\Pr_{\boldsymbol{R}}[|X_{\boldsymbol{R}} - 1| \geq t] \leq 2\cdot\exp\big(-c\cdot K\cdot\min\{t^2, t\}\big).$$

Proof. For each $1 \le k \le K$, define

$$X_{\boldsymbol{R},k} \coloneqq \frac{1}{N} \cdot \Big(\sum_{x \in \{0,1\}^n} \boldsymbol{R}_k(x) \Big)^2 = \Big(\sum_{x \in \{0,1\}^n} \frac{1}{\sqrt{N}} \cdot \boldsymbol{R}_k(x) \Big)^2,$$

Then Lemma 4.13 implies that there exists a constant $\gamma \ge 1$ such that for each $1 \le k \le K$, the random variable $|X_{\boldsymbol{R},k}|^2 - 1$ is sub-exponential with parameter γ. Since $X_{\boldsymbol{R}} = \mathbf{E}_{\boldsymbol{k} \sim [K]}[X_{\boldsymbol{R},\boldsymbol{k}}]$, Bernstein's inequality (Theorem 4.14) states that

$$\Pr_{\boldsymbol{R}}[|X_{\boldsymbol{R}} - 1| \ge t] \le 2 \cdot \exp\left(-c \cdot \min\left\{ \frac{t^2}{\gamma^2}, \frac{t}{\gamma} \right\} \cdot K \right) \le 2 \cdot \exp\left(-c \cdot \min\left\{ \frac{t^2}{\gamma^2}, \frac{t}{\gamma^2} \right\} \cdot K \right),$$

because $\gamma \ge 1$. This completes the proof, by setting the "c" in the lemma statement to c/γ^2. □

Now we prove our bound on the expected distinguishing advantage of the Hadamard tester.

Proof of Theorem C.2. By Lemma C.4, it suffices to show that $\mathbf{E}_{\boldsymbol{R}} |X_{\boldsymbol{R}} - 1| \ge \Omega(1/\sqrt{K})$. We will show this by deriving the following weak anti-concentration result: there exists a constant $\varepsilon > 0$ such that

$$\Pr_{\boldsymbol{R}}\left[|X_{\boldsymbol{R}} - 1| \ge \frac{\varepsilon}{\sqrt{K}} \right] \ge \varepsilon. \tag{37}$$

Assuming this is true, then our main result can be shown as follows:

$$\mathop{\mathbf{E}}_{\boldsymbol{R}} |X_{\boldsymbol{R}} - 1| \ge \frac{\varepsilon}{\sqrt{K}} \cdot \Pr\left[|X_{\boldsymbol{R}} - 1| \ge \frac{\varepsilon}{\sqrt{K}} \right] \ge \frac{\varepsilon^2}{\sqrt{K}}.$$

Now we prove Equation (37). The proof is by contradiction: for sake of contradiction, let us assume that it is false. Then one can obtain an upper bound for the variance of $X_{\boldsymbol{R}}$ as follows:

$$\begin{aligned}
&\mathop{\mathbf{E}}_{\boldsymbol{R}}\Big[|X_{\boldsymbol{R}} - 1|^2 \Big] \\
&= \int_0^\infty \Pr\Big[|X_{\boldsymbol{R}} - 1|^2 \ge t \Big] dt \\
&= \int_0^{\frac{\varepsilon^2}{K}} \Pr\Big[|X_{\boldsymbol{R}} - 1|^2 \ge t \Big] dt + \int_{\frac{\varepsilon^2}{K}}^{\frac{1}{\varepsilon \cdot K}} \Pr\Big[|X_{\boldsymbol{R}} - 1|^2 \ge t \Big] dt + \int_{\frac{1}{\varepsilon \cdot K}}^\infty \Pr\Big[|X_{\boldsymbol{R}} - 1|^2 \ge t \Big] dt. \qquad (38)
\end{aligned}$$

We can upper-bound the first term by ε^2/K since probabilities are always at most one. As for the second term, since we are assuming that Equation (37) is false, we have that

$$\int_{\frac{\varepsilon^2}{K}}^{\frac{1}{\varepsilon \cdot K}} \Pr\Big[|X_{\boldsymbol{R}} - 1|^2 \ge t \Big] dt \le \int_{\frac{\varepsilon^2}{K}}^{\frac{1}{\varepsilon \cdot K}} \Pr\Big[|X_{\boldsymbol{R}} - 1|^2 \ge \frac{\varepsilon^2}{K} \Big] dt < \int_{\frac{\varepsilon^2}{K}}^{\frac{1}{\varepsilon \cdot K}} \varepsilon \cdot dt < \varepsilon \cdot \frac{1}{\varepsilon \cdot K} = \frac{1}{K}.$$

Finally, we can bound the third term using Lemma C.6. In total, we have that

$$\begin{aligned}
(38) &< \frac{\varepsilon^2}{K} + \frac{1}{K} + \int_{\frac{1}{\varepsilon \cdot K}}^\infty 2 \cdot \exp\Big(-c \cdot K \cdot \min(t, \sqrt{t}) \Big) dt \\
&\le \frac{\varepsilon^2}{K} + \frac{1}{K} + \int_{\frac{1}{\varepsilon \cdot K}}^\infty 2 \cdot \exp(-c \cdot K \cdot t) dt + \int_1^\infty 2 \cdot \exp(-c \cdot K \cdot t) \cdot 2t \cdot dt \\
&= \frac{\varepsilon^2}{K} + \frac{1}{K} + \frac{2}{c \cdot K} \cdot \exp\Big(-\frac{c}{\varepsilon} \Big) + 4 \cdot \Big(\frac{1}{c \cdot K} + \frac{1}{(c \cdot K)^2} \Big) \cdot \exp(-c \cdot K).
\end{aligned}$$

For ε a sufficiently small constant and K a sufficiently large constant, this is at most $\frac{3}{2} \cdot \frac{1}{K}$. But we already calculated that the variance of $X_{\boldsymbol{R}}$ is $\frac{N-1}{N} \cdot \frac{2}{K}$, in Lemma C.5. Hence, we have a contradiction for sufficiently large N, completing the proof. $\square$