

Shrinkage under Random Projections, and Cubic Formula Lower Bounds for \mathbf{AC}^0 *

Yuval Filmus[†] Or Meir[‡] Avishay Tal[§]

December 29, 2020

Abstract

Håstad showed that any De Morgan formula (composed of AND, OR and NOT gates) shrinks by a factor of $O(p^2)$ under a random restriction that leaves each variable alive independently with probability p [SICOMP, 1998]. Using this result, he gave an $\tilde{\Omega}(n^3)$ formula size lower bound for the Andreev function, which, up to lower order improvements, remains the state-of-the-art lower bound for any explicit function.

In this work, we extend the shrinkage result of Håstad to hold under a far wider family of random restrictions and their generalization — random projections. Based on our shrinkage results, we obtain an $\tilde{\Omega}(n^3)$ formula size lower bound for an explicit function computed in \mathbf{AC}^0 . This improves upon the best known formula size lower bounds for \mathbf{AC}^0 , that were only quadratic prior to our work. In addition, we prove that the KRW conjecture [Karchmer et al., Computational Complexity 5(3/4), 1995] holds for inner functions for which the unweighted quantum adversary bound is tight. In particular, this holds for inner functions with a tight Khrapchenko bound.

Our random projections are tailor-made to the function’s structure so that the function maintains structure even under projection — using such projections is necessary, as standard random restrictions simplify \mathbf{AC}^0 circuits. In contrast, we show that any De Morgan formula shrinks by a quadratic factor under our random projections, allowing us to prove the cubic lower bound.

Our proof techniques build on the proof of Håstad for the simpler case of balanced formulas. This allows for a significantly simpler proof at the cost of slightly worse parameters. As such, when specialized to the case of p -random restrictions, our proof can be used as an exposition of Håstad’s result.

1 Introduction

1.1 Background

Is there an efficient computational task that cannot be perfectly parallelized? Equivalently, is $\mathbf{P} \not\subseteq \mathbf{NC}^1$? The answer is still unknown. The question can be rephrased as follows: is there a

*An extended abstract of this work has been accepted to ITCS 2021.

[†]Technion — Israel Institute of Technology, Haifa 3200003, Israel. yuvalfi@cs.technion.ac.il. Taub Fellow — supported by the Taub Foundations. The research was funded by ISF grant 1337/16.

[‡]Department of Computer Science, University of Haifa, Haifa 3498838, Israel. ormeir@cs.haifa.ac.il. Partially supported by the Israel Science Foundation (grant No. 1445/16).

[§]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, United States. atal@berkeley.edu

function in \mathbf{P} that does not have a (De Morgan) formula of polynomial size?

The history of formula lower bounds for functions in \mathbf{P} goes back to the 1960s, with the seminal result of Subbotovskaya [Sub61] that introduced the technique of random restrictions. Subbotovskaya showed that the Parity function on n variables requires formulas of size at least $\Omega(n^{1.5})$. Khrapchenko [Khr72], using a different proof technique, showed that in fact the Parity function on n variables requires formulas of size $\Theta(n^2)$. Later, Andreev [And87] came up with a new explicit function (now known as the Andreev function) for which he was able to obtain an $\Omega(n^{2.5})$ size lower bound. This lower bound was subsequently improved by [IN93, PZ93, Hås98, Tal14] to $n^{3-o(1)}$.

The line of work initiated by Subbotovskaya and Andreev relies on the *shrinkage* of formulas under p -random restrictions. A p -random restriction is a randomly chosen partial assignment to the inputs of a function. Set a parameter $p \in (0, 1)$. We fix each variable independently with probability $1 - p$ to a uniformly random bit, and we keep the variable alive with probability p . Under such a restriction, formulas shrink (in expectation) by a factor more significant than p . Subbotovskaya showed that De Morgan formulas shrink to at most $p^{1.5}$ times their original size, whereas subsequent works of [PZ93, IN93] improved the bound to $p^{1.55}$ and $p^{1.63}$, respectively. Finally, Håstad [Hås98] showed that the shrinkage exponent of De Morgan formulas is 2, or in other words, that De Morgan formulas shrink by a factor of $p^{2-o(1)}$ under p -random restrictions. Tal [Tal14] improved the shrinkage factor to $O(p^2)$ — obtaining a tight result, as exhibited by the Parity function.

In a nutshell, shrinkage results are useful to proving lower bounds as long as the explicit function being analyzed maintains structure under such restrictions and does not trivialize. For example, the Parity function does not become constant as long as at least one variable remains alive. Thus any formula F that computes Parity must be of at least quadratic size, or else the formula F under restriction, keeping each variable alive with probability $100/n$, would likely become a constant function, whereas Parity would not. Andreev’s idea is similar, though he manages to construct a function such that under a random restriction keeping only $\Theta(\log n)$ of the variables, the formula size should be at least $\tilde{\Omega}(n)$ (in expectation). This ultimately gives the nearly cubic lower bound.

The KRW Conjecture. Despite much effort, proving $\mathbf{P} \not\subseteq \mathbf{NC}^1$, and even just breaking the cubic barrier in formula lower bounds, have remained a challenge for more than two decades. An approach to solve the \mathbf{P} versus \mathbf{NC}^1 problem was suggested by Karchmer, Raz and Wigderson [KRW95]. They conjectured that when composing two Boolean functions, f and g , the formula size of the resulting function, $f \diamond g$, is (roughly) the product of the formula sizes of f and g .¹ We will refer to this conjecture as the “KRW conjecture”. Under the KRW conjecture (and even under weaker variants of it), [KRW95] constructed a function in \mathbf{P} with no polynomial-size formulas. It remains a major open challenge to settle the KRW conjecture.

A few special cases of the KRW conjecture are known to be true. The conjecture holds when either f or g is the AND or the OR function. Håstad’s result [Hås98] and its improvement [Tal14] show that the conjecture holds when the inner function g is the Parity function and the outer function f is any function. This gives an alternative explanation to the $n^{3-o(1)}$ lower bound for the Andreev function. Indeed, the Andreev function is at least as hard as the composition of a maximally-hard function f on $\log n$ bits and $g = \text{Parity}_{n/\log n}$, where the formula size of f is $\tilde{\Omega}(n)$ and the formula size of $\text{Parity}_{n/\log n}$ is $\Theta(n^2/\log^2 n)$. Since the KRW conjecture holds for this special

¹More precisely, the original KRW conjecture [KRW95] concerns depth complexity rather than formula complexity. The variant of the conjecture for formula complexity, which is discussed above, was posed in [GMWW17].

case, the formula size of the Andreev function is at least $\tilde{\Omega}(n^3)$. In other words, the state-of-the-art formula size lower bounds for explicit functions follow from a special case of the KRW conjecture — the case in which g is the Parity function. Moreover, this special case follows from the shrinkage of De Morgan formulas under p -random restrictions.

Bottom-Up versus Top-Down Techniques. Whereas random restrictions are a “bottom-up” proof technique [HJP95], a different line of work suggested a “top-down” approach using the language of communication complexity. The connection between formula size and communication complexity was introduced in the seminal work of Karchmer and Wigderson [KW90]. They defined for any Boolean function f a two-party communication problem KW_f : Alice gets an input x such that $f(x) = 1$, and Bob gets an input y such that $f(y) = 0$. Their goal is to identify a coordinate i on which $x_i \neq y_i$, while minimizing their communication. It turns out that there is a one-to-one correspondence between any protocol tree solving KW_f and any formula computing the function f . Since protocols naturally traverse the tree from root to leaf, proving lower bounds on their size or depth is done usually in a top-down fashion. This framework has proven to be very useful in proving formula lower bounds in the monotone setting (see, e.g., [KW90, GH92, RW92, KRW95, RM99, GP18, PR17]) and in studying the KRW conjecture (see, e.g., [KRW95, EIRS01, HW93, GMWW17, DM18, KM18, Mei20, dRMN⁺20, MS20]). Moreover, a recent work by Dinur and Meir [DM18] was able to reprove Håstad’s cubic lower bound using the framework of Karchmer and Wigderson. As Dinur and Meir’s proof showed that top-down techniques can replicate Håstad’s cubic lower bound, a natural question (which motivated this project) arose:

Are top-down techniques superior to bottom-up techniques?

Towards that, we focused on a candidate problem: prove a cubic lower bound for an explicit function in \mathbf{AC}^0 .² Based on the work of Dinur and Meir [DM18], we suspected that such a lower bound could be achieved using top-down techniques. We were also *certain* that the problem cannot be solved using the random restriction technique. Indeed, in order to prove a lower bound on a function f using random restrictions, one should argue that f remains hard under a random restriction, however, it is well-known that functions in \mathbf{AC}^0 trivialize under p -random restrictions [Ajt83, FSS84, Yao85, Hås86]. Based on this intuition, surely random restrictions cannot show that a function in \mathbf{AC}^0 requires cubic size. Our intuition turned out to be false.

1.2 Our results

In this work, we construct an explicit function in \mathbf{AC}^0 which requires De Morgan formulas of size $n^{3-o(1)}$. Surprisingly, our proof is conducted via the bottom-up technique of random projections, which is a generalization of random restrictions (more details below).

Theorem 1.1. *There exists a family of Boolean functions $h_n : \{0, 1\}^n \rightarrow \{0, 1\}$ for $n \in \mathbb{N}$ such that*

1. h_n can be computed by uniform depth-4 unbounded fan-in formulas of size $O(n^3)$.
2. The formula size of h_n is at least $n^{3-o(1)}$.

²Recall that \mathbf{AC}^0 is the class of functions computed by constant depth polynomial size circuits composed of AND and OR gates of unbounded fan-in, with variables or their negation at the leaves.

Prior to our work, the best formula size lower bounds on an explicit function in \mathbf{AC}^0 were only quadratic [Nec66, CKK12, Juk12, BM12].

Our hard function is a variant of the Andreev function. More specifically, recall that the Andreev function is based on the composition $f \diamond g$, where f is a maximally-hard function and g is the Parity function. Since Parity is not in \mathbf{AC}^0 , we cannot take g to be the Parity function in our construction. Instead, our hard function is obtained by replacing the Parity function with the Surjectivity function of [BM12].

As in the case of the Andreev function, we establish the hardness of our function by proving an appropriate special case of the KRW conjecture. To this end, we introduce a generalization of the complexity measure of Khrapchenko [Khr72], called the *min-entropy Khrapchenko bound*. We prove the KRW conjecture for the special case in which the outer function f is any function, and g is a function whose formula complexity is bounded tightly by the min-entropy Khrapchenko bound. We then obtain Theorem 1.1 by applying this version of the KRW conjecture to the case where g is the Surjectivity function. We note that our KRW result also implies the known lower bounds in the cases where g is the Parity function [Hås98] and the Majority function [GTN19].

Our KRW result in fact applies more generally, to functions g whose formula complexity is bounded tightly by the “soft-adversary method”, denoted $\text{Adv}_s(g)$, which is a generalization of Ambainis’ unweighted adversary method [Amb02] (see Section 6.2).

Our proof of the special case of the KRW conjecture follows the methodology of Håstad [Hås93], who proved the special case in which g is Parity on m variables. Håstad proved that De Morgan formulas shrink by a factor of (roughly) p^2 under p -random restrictions. Choosing $p = 1/m$ shrinks a formula for $f \diamond g$ by a factor of roughly m^2 , which coincides with the formula complexity of g . On the other hand, on average each copy of g simplifies to a single input variable, and so $f \diamond g$ simplifies to f . This shows that $L(f \diamond g) \gtrsim L(f) \cdot L(g)$.

Our main technical contribution is a new shrinkage theorem that works in a far wider range of scenarios than just p -random restrictions. Given a function g with soft-adversary bound $\text{Adv}_s(g)$, we construct a random projection³ which, on the one hand, shrinks De Morgan formulas by a factor of $\text{Adv}_s(g)$, and on the other hand, simplifies $f \diamond g$ to f . We thus show that $L(f \diamond g) \gtrsim L(f) \cdot \text{Adv}_s(g)$, and in particular, if $\text{Adv}_s(g) \approx L(g)$, then $L(f \diamond g) \gtrsim L(f) \cdot L(g)$, just as in Håstad’s proof. Our random projections are tailored specifically to the structure of the function $f \diamond g$, ensuring that $f \diamond g$ simplifies to f under projection. This enables us to overcome the aforementioned difficulty. In contrast, p -random restrictions that do not respect the structure of $f \diamond g$ would likely result in a restricted function that is much simpler than f and in fact would be a constant function with high probability.

Our shrinkage theorem applies more generally to two types of random projections, which we call *fixing projections* and *hiding projections*. Fixing projections are random projections in which fixing the value of a variable results in a projection which is much more probable. Hiding projections are random projections in which fixing the value of a variable hides which coordinates it appeared on. We note that our shrinkage theorem for fixing projections captures Håstad’s result for p -random restrictions as a special case.

The proof of our shrinkage theorem is based on Håstad’s proof [Hås98], but also simplifies it. In particular, we take the simpler argument that Håstad uses for the special case of completely balanced trees, and adapt it to the general case. As such, our proof avoids a complicated case

³A projection is a mapping from the set of the variables $\{x_1, \dots, x_n\}$ to the set $\{y_1, \dots, y_m, \overline{y_1}, \dots, \overline{y_m}, 0, 1\}$, where y_1, \dots, y_m are formal variables.

analysis, at the cost of slightly worse bounds. Using our bounds, it is nevertheless easy to obtain the $n^{3-o(1)}$ lower bound for the Andreev function. Therefore, one can see the specialization of our shrinkage result to p -random restrictions as an exposition of Håstad’s cubic lower bound.

An example: our techniques when specialized to $f \diamond \text{Majority}_m$. To illustrate our choice of random projections, we present its instantiation to the special case of $f \diamond g$, where $f: \{0, 1\}^k \rightarrow \{0, 1\}$ is non-constant and $g = \text{Majority}_m$ for some odd integer m . In this case, the input variables to $f \diamond g$ are composed of k disjoint blocks, B_1, \dots, B_k , each containing m variables. We use the random projection that for each block $B_i = \{x_{m(i-1)+1}, \dots, x_{mi}\}$, picks one variable in the block B_i uniformly at random, projects this variable to the new variable y_i , and fixes the rest of the variables in the block in a balanced way so that the number of zeros and ones in the block is equal (i.e., we have exactly $(m-1)/2$ zeros and $(m-1)/2$ ones). It is not hard to see that under this choice, $f \diamond g$ simplifies to f . On the other hand, we show that this choice of random projections shrinks the formula complexity by a factor of $\approx 1/m^2$. Combining the two together, we get that $L(f \diamond \text{Majority}_m) \gtrsim L(f) \cdot m^2$. Note that in this distribution of random projections, the different coordinates are not independent of one another, and this feature allows us to maintain structure.

1.3 Related work

Our technique of using tailor-made random projections was inspired by the celebrated result of Rossman, Servedio, and Tan [RST15, HRST17] that proved an average-case depth hierarchy. In fact, the idea to use tailor-made random restrictions goes back to Håstad’s thesis [Hås87, Chapter 6.2]. Similar to our case, in [Hås87, RST15, HRST17], p -random restrictions are too crude to separate depth d from depth $d+1$ circuits. Given a circuit C of depth $d+1$, the main challenge is to construct a distribution of random restrictions or projections (tailored to the circuit C) that on the one hand maintains structure for C , but on the other hand simplify any depth d circuit C' .

Paper outline

The paper starts with brief preliminaries in Section 2. We prove our shrinkage theorem for fixing projections in Section 3, and our shrinkage theorem for hiding projections in Section 4. In Section 5 we provide a brief interlude on concatenation of projections. Khrapchenko’s method, the quantum adversary bound and their relation to hiding projections are discussed in Section 6. Finally, Section 7 contains a proof of Theorem 1.1, as a corollary of a more general result which is a special case of the KRW conjecture. In the same section we also rederive the cubic lower bound on Andreev’s function, and the cubic lower bound on the Majority-based variant considered in [GTN19].

2 Preliminaries

Throughout the paper, we use bold letters to denote random variables. For any $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$. Given a bit $\sigma \in \{0, 1\}$, we denote its negation by $\bar{\sigma}$. We assume familiarity with the basic definitions of communication complexity (see, e.g., [KN97]). All logarithms in this paper are base 2.

Definition 2.1. A (*De Morgan*) formula (*with bounded fan-in*) is a binary tree, whose leaves are labeled with literals from the set $\{x_1, \bar{x}_1, \dots, x_n, \bar{x}_n\}$, and whose internal vertices are labeled as

AND (\wedge) or OR (\vee) gates. The *size* of a formula ϕ , denoted $\text{size}(\phi)$, is the number of leaves in the tree. The *depth* of the formula is the depth of the tree. A *formula with unbounded fan-in* is defined similarly, but every internal vertex in the tree can have any number of children. Unless stated explicitly otherwise, whenever we say “formula” we refer to a formula with bounded fan-in.

Definition 2.2. A formula ϕ computes a Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ in the natural way. The *formula complexity* of a Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, denoted $L(f)$, is the size of the smallest formula that computes f . The *depth complexity* of f , denoted $D(f)$, is the smallest depth of a formula that computes f . For convenience, we define the size and depth of the constant function to be zero.

A basic property of formula complexity is that it is subadditive:

Fact 2.3. *For every two functions $f_1, f_2: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $L(f_1 \wedge f_2) \leq L(f_1) + L(f_2)$ and $L(f_1 \vee f_2) \leq L(f_1) + L(f_2)$.*

The following theorem shows that every small formula can be “balanced” to obtain a shallow formula.

Theorem 2.4 (Formula balancing, [BB94], following [Spi71, Bre74]). *For every $\alpha > 0$, the following holds: For every formula ϕ of size s , there exists an equivalent formula ϕ' of depth at most $O(2^{\frac{1}{\alpha}} \cdot \log s)$ and size at most $s^{1+\alpha}$.*

Notation 2.5. With a slight abuse of notation, we will often identify a formula ϕ with the function it computes. In particular, the notation $L(\phi)$ denotes *the formula complexity of the function* computed by ϕ , and *not the size of ϕ* (which is denoted by $\text{size}(\phi)$).

Notation 2.6. Given a Boolean variable z , we denote by z^0 and z^1 the literals z and \bar{z} , respectively. In other words, $z^b = z \oplus b$.

Notation 2.7. Given a literal ℓ , we define $\text{var}(\ell)$ to be the underlying variable, that is, $\text{var}(z) = \text{var}(\bar{z}) = z$.

Notation 2.8. Let Π be a deterministic communication protocol that takes inputs from $\mathcal{A} \times \mathcal{B}$, and recall that the leaves of the protocol induce a partition of $\mathcal{A} \times \mathcal{B}$ to combinatorial rectangles. For every leaf ℓ of Π , we denote by $\mathcal{A}_\ell \times \mathcal{B}_\ell$ the combinatorial rectangle that is associated with ℓ .

We use the framework of Karchmer–Wigderson relations [KW90], which relates the complexity of f to the complexity of a related communication problem KW_f .

Definition 2.9 ([KW90]). Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. The *Karchmer–Wigderson relation* of f , denoted KW_f , is the following communication problem: The inputs of Alice and Bob are strings $a \in f^{-1}(1)$ and $b \in f^{-1}(0)$, respectively, and their goal is to find a coordinate $i \in [n]$ such that $a_i \neq b_i$. Note that such a coordinate must exist since $f^{-1}(1) \cap f^{-1}(0) = \emptyset$ and hence $a \neq b$.

Theorem 2.10 ([KW90], see also [Raz90]). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. The communication complexity of KW_f is equal to $D(f)$, and the minimal number of leaves in a protocol that solves KW_f is $L(f)$.*

We use the following two standard inequalities.

Fact 2.11 (the AM-GM inequality). *For every two non-negative real numbers x, y it holds that $\sqrt{x \cdot y} \leq \frac{x+y}{2}$.*

Fact 2.12 (special case of Cauchy-Schwarz inequality). *For every t non-negative real numbers x_1, \dots, x_t it holds that $\sqrt{x_1} + \dots + \sqrt{x_t} \leq \sqrt{t} \cdot \sqrt{x_1 + \dots + x_t}$.*

Proof. It holds that $\sqrt{x_1} + \dots + \sqrt{x_t} \leq \sqrt{1^2 + \dots + 1^2} \cdot \sqrt{(\sqrt{x_1})^2 + \dots + (\sqrt{x_t})^2} = \sqrt{t} \cdot \sqrt{x_1 + \dots + x_t}$, as required. \square

3 Shrinkage theorem for fixing projections

In this section we prove our main result on the shrinkage of De Morgan formulas under fixing projections, which we define below. We start by defining projections and the relevant notation.

Definition 3.1. Let x_1, \dots, x_n and y_1, \dots, y_m be Boolean variables. A *projection* π from x_1, \dots, x_n to y_1, \dots, y_m is a function from the set $\{x_1, \dots, x_n\}$ to the set $\{0, 1, y_1, \bar{y}_1, \dots, y_m, \bar{y}_m\}$. Given such a projection π and a Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ over the variables x_1, \dots, x_n , we denote by $f|_\pi: \{0, 1\}^m \rightarrow \{0, 1\}$ the function obtained from f by substituting each input variable x_i with $\pi(x_i)$ in the natural way. Unless stated explicitly otherwise, all projections in this section are from x_1, \dots, x_n to y_1, \dots, y_m , and all functions from $\{0, 1\}^n$ to $\{0, 1\}$ are over the variables x_1, \dots, x_n . A *random projection* is a distribution over projections.

Notation 3.2. Let π be a projection. For every $j \in [m]$ and bit $\sigma \in \{0, 1\}$, we denote by $\pi_{y_j \leftarrow \sigma}$ the projection that is obtained from π by substituting y_j with σ .

Notation 3.3. With a slight abuse of notation, if a projection π maps all the variables x_1, \dots, x_n to constants in $\{0, 1\}$, we will sometimes treat it as a binary string in $\{0, 1\}^n$.

We use a new notion of random projections, which we call *q-fixing projections*. Intuitively, a q -fixing projection is a random projection in which for every variable x_i , the probability that π maps a variable x_i to a literal is not much larger than the probability that π fixes that literal to a constant, regardless of the values that π assigns to the other variables. This property is essentially the minimal property that is required in order to carry out the argument of Håstad [Hås98]. Formally, we define q -fixing projections as follows.

Definition 3.4. Let $0 \leq q_0, q_1 \leq 1$. We say that a random projection π is a (q_0, q_1) -fixing projection if for every projection π , every bit $\sigma \in \{0, 1\}$, and every variable x_i , it holds that

$$\Pr[\pi(x_i) \notin \{0, 1\} \text{ and } \pi_{\text{var}(\pi(x_i)) \leftarrow \sigma} = \pi] \leq q_\sigma \cdot \Pr[\pi = \pi]. \quad (1)$$

For shorthand, we say that π is a q -fixing projection, for $q = \sqrt{q_0 q_1}$.

If needed, one can consider without loss of generality only variables x_i such that $\pi(x_i) \in \{0, 1\}$, as otherwise Equation (1) holds trivially with the left-hand side equaling zero.

Example 3.5. In order to get intuition for the definition of fixing projections, let us examine how this definition applies to random restrictions. In our terms, a restriction is a projection from x_1, \dots, x_n to x_1, \dots, x_n that maps every variable x_i either to itself or to $\{0, 1\}$. Suppose that ρ is

any distribution over restrictions, and that ρ is some fixed restriction. In this case, the condition of being q -fixing can be rewritten as follows:

$$\Pr[\rho(x_i) = x_i \text{ and } \rho_{x_i \leftarrow \sigma} = \rho] \leq q_\sigma \cdot \Pr[\rho = \rho].$$

Denote by ρ', ρ' the restrictions obtained from ρ, ρ by truncating x_i (i.e., $\rho' = \rho|_{\{x_1, \dots, x_n\} - \{x_i\}}$). Using this notation, we can rewrite the foregoing equation as

$$\Pr[\rho(x_i) = x_i \text{ and } \rho' = \rho' \text{ and } \rho_{x_i \leftarrow \sigma}(x_i) = \rho(x_i)] \leq q_\sigma \cdot \Pr[\rho(x_i) = \rho(x_i) \text{ and } \rho' = \rho'].$$

Now, observe that it is always the case $\rho_{x_i \leftarrow \sigma}(x_i) = \sigma$, and therefore the probability on the left-hand side is non-zero only if $\rho(x_i) = \sigma$. Hence, we can restrict ourselves to the latter case, and the foregoing equation can be rewritten again as

$$\Pr[\rho(x_i) = x_i \text{ and } \rho' = \rho'] \leq q_\sigma \cdot \Pr[\rho(x_i) = \sigma \text{ and } \rho' = \rho'].$$

Finally, if we divide both sides by $\Pr[\rho' = \rho']$, we obtain the following intuitive condition:

$$\Pr[\rho(x_i) = x_i \mid \rho' = \rho'] \leq q_\sigma \cdot \Pr[\rho(x_i) = \sigma \mid \rho' = \rho'].$$

This condition informally says the following: ρ is a fixing projection if the probability of leaving x_i unfixed is at most q_σ times the probability of fixing it to σ , and this holds regardless of what the restriction assigns to the other variables.

In particular, it is now easy to see that the classic random restrictions are fixing projections. Recall that a p -random restriction fixes each variable independently with probability $1 - p$ to a random bit. Due to the independence of the different variables, the foregoing condition simplifies to

$$\Pr[\rho(x_i) = x_i] \leq q_\sigma \cdot \Pr[\rho(x_i) = \sigma],$$

and it is easy to see that this condition is satisfied for $q_0 = q_1 = \frac{2p}{1-p}$.

We prove the following shrinkage theorem for q -fixing projections, which is analogous to the shrinkage theorem of [Hås98] for random restrictions in the case of balanced formulas.

Theorem 3.6 (Shrinkage under fixing projections). *Let ϕ be a formula of size s and depth d , and let π be a q -fixing projection. Then*

$$\mathbb{E}[L(\phi|_\pi)] = O(q^2 \cdot d^2 \cdot s + q \cdot \sqrt{s}).$$

Our shrinkage theorem has somewhat worse parameters compared to the theorem of [Hås98]: specifically, the factor of d^2 does not appear in [Hås98]. The reason is that the proof of [Hås98] uses a fairly-complicated case-analysis in order to avoid losing that factor, and we chose to skip this analysis in order to obtain a simpler proof. We did not check if the factor of d^2 in our result can be avoided by using a similar case-analysis. By applying formula balancing (Theorem 2.4) to our shrinkage theorem, we can obtain the following result, which is independent of the depth of the formula.

Corollary 3.7. *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a function with formula complexity s , and let π be a q -fixing projection. Then*

$$\mathbb{E}[L(f|_\pi)] = q^2 \cdot s^{1+O(\frac{1}{\sqrt{\log s}})} + q \cdot s^{1/2+O(\frac{1}{\sqrt{\log s}})}.$$

Proof. By assumption, there exists a formula ϕ of size s that computes f . We balance the formula ϕ by applying Theorem 2.4 with $\alpha = \frac{1}{\sqrt{\log s}}$, and obtain a new formula ϕ' that computes f and has size $s^{1+\frac{1}{\sqrt{\log s}}}$ and depth $O(2^{\sqrt{\log s}} \cdot \log s) = s^{O(\frac{1}{\sqrt{\log s}})}$. The required result now follows by applying Theorem 3.6 to ϕ' . \square

3.1 Proof of Theorem 3.6

In this section, we prove our main shrinkage theorem, Theorem 3.6. Our proof is based on the ideas of [Hås98], but the presentation is different. Fix a formula ϕ of size s and depth d , and let π be a q -fixing projection. We would like to upper-bound the expectation of $L(\phi|\pi)$. As in [Hås98], we start by upper-bounding the probability that the projection π shrinks a formula to size 1. Specifically, we prove the following lemma in Section 3.2.

Lemma 3.8. *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function, and let π be a q -fixing projection. Then,*

$$\Pr[L(f|\pi) = 1] \leq q \cdot \sqrt{L(f)}.$$

Next, we show that to upper-bound the expectation of $L(\phi|\pi)$, it suffices to upper-bound the probability that the projection π shrinks two formulas to size 1 simultaneously. In order to state this claim formally, we introduce some notation.

Notation 3.9. Let g be a gate of ϕ . We denote the depth of g in ϕ by $\text{depth}_\phi(g)$ (the root has depth 0), and omit ϕ if it is clear from context. If g is an internal node, we denote the sub-formulas that are rooted in its left and right children by $\text{left}(g)$ and $\text{right}(g)$, respectively.

We prove the following lemma, which says that in order to upper-bound $L(\phi|\pi)$ it suffices to upper-bound, for every internal gate g , the probability that $\text{left}(g)$ and $\text{right}(g)$ shrink to size 1 under π .

Lemma 3.10. *For every projection π it holds that*

$$L(\phi|\pi) \leq \sum_{\text{internal gate } g \text{ of } \phi} (\text{depth}(g) + 2) \cdot \mathbb{1}_{\{L(\text{left}(g)|\pi)=1 \text{ and } L(\text{right}(g)|\pi)=1\}} + \mathbb{1}_{L(\phi|\pi)=1}.$$

We would like to use Lemmas 3.8 and 3.10 to prove the shrinkage theorem. As a warm-up, let us make the simplifying assumption that for every two functions $f_1, f_2: \{0, 1\}^n \rightarrow \{0, 1\}$, the events $L(f_1|\pi) = 1$ and $L(f_2|\pi) = 1$ are independent. If this was true, we could have upper-

bounded $\mathbb{E}[L(\phi|\pi)]$ as follows:

$$\begin{aligned}
\mathbb{E}[L(\phi|\pi)] &\leq \sum_{\text{int. gate } g \text{ of } \phi} (\text{depth}(g) + 2) \cdot \mathbb{E}[1_{\{L(\text{left}(g)|\pi)=1 \text{ and } L(\text{right}(g)|\pi)=1\}}] && \text{(Lemma 3.10)} \\
&+ \mathbb{E}[1_{L(\phi|\pi)=1}] \\
&\leq (d+2) \cdot \sum_{\text{int. gate } g \text{ of } \phi} \Pr[L(\text{left}(g)|\pi) = 1 \text{ and } L(\text{right}(g)|\pi) = 1] && (\phi \text{ is of depth } d) \\
&+ \mathbb{E}[1_{L(\phi|\pi)=1}] \\
&= (d+2) \cdot \sum_{\text{int. gate } g \text{ of } \phi} \Pr[L(\text{left}(g)|\pi) = 1] \cdot \Pr[L(\text{right}(g)|\pi) = 1] && \text{(simplifying assumption)} \\
&+ \mathbb{E}[1_{L(\phi|\pi)=1}] \\
&\leq (d+2) \cdot \sum_{\text{int. gate } g \text{ of } \phi} q^2 \cdot \sqrt{L(\text{left}(g)) \cdot L(\text{right}(g))} + q \cdot \sqrt{s} && \text{(Lemma 3.8)} \\
&\leq q^2 \cdot (d+2) \cdot \sum_{\text{int. gate } g \text{ of } \phi} (L(\text{left}(g)) + L(\text{right}(g))) + q \cdot \sqrt{s} && \text{(AM-GM inequality)} \\
&\leq q^2 \cdot (d+2) \cdot \sum_{\text{int. gate } g \text{ of } \phi} (\text{size}(\text{left}(g)) + \text{size}(\text{right}(g))) + q \cdot \sqrt{s} \\
&= q^2 \cdot (d+2) \cdot \sum_{\text{int. gate } g \text{ of } \phi} \text{size}(g) + q \cdot \sqrt{s}.
\end{aligned}$$

The last sum counts every leaf ℓ of ϕ once for each internal ancestor of ℓ , so the last expression is equal to

$$\begin{aligned}
&q^2 \cdot (d+2) \cdot \sum_{\text{leaf } \ell \text{ of } \phi} \text{depth}(\ell) + q \cdot \sqrt{s} \\
&\leq q^2 \cdot (d+2) \cdot \sum_{\text{leaf } \ell \text{ of } \phi} d + q \cdot \sqrt{s} \\
&= q^2 \cdot (d+2) \cdot d \cdot s + q \cdot \sqrt{s} \\
&= O(q^2 \cdot d^2 \cdot s + q \cdot \sqrt{s}),
\end{aligned}$$

which is the bound we wanted. However, the above calculation only works under our simplifying assumption, which is false: the events $L(f_1|\pi) = 1$ and $L(f_2|\pi) = 1$ will often be dependent. In particular, in order for the foregoing calculation to work, we need to the following inequality to hold:

$$\Pr[L(f_2|\pi) = 1 \mid L(f_1|\pi) = 1] \leq q \cdot \sqrt{L(f_2)}.$$

This inequality holds under our simplifying assumption by Lemma 3.8, but may not hold in general. Nevertheless, we prove the following similar statement in Section 3.3.

Lemma 3.11. *Let π be a q -fixing projection. Let $f_1, f_2: \{0, 1\}^n \rightarrow \{0, 1\}$, let $\sigma, \tau \in \{0, 1\}$, and let y_j be a variable. Then,*

$$\Pr[L(f_2|\pi_{y_j \leftarrow \sigma}) = 1 \mid f_1|\pi = y_j^\tau] \leq q \cdot \sqrt{L(f_2)}.$$

Intuitively, Lemma 3.11 breaks the dependency between the events $L(f_1|\pi) = 1$ and $L(f_2|\pi) = 1$ by fixing in f_2 the single literal to which f_1 has shrunk. We would now like to use Lemma 3.11 to prove the theorem. To this end, we prove an appropriate variant of Lemma 3.10, which allows using the projection $\pi_{y_j \leftarrow \sigma}$ rather than π in the second function. This variant is motivated by the following “one-variable simplification rules” of [Hås98], which are easy to verify.

Fact 3.12 (one-variable simplification rules). *Let $h: \{0, 1\}^m \rightarrow \{0, 1\}$ be a function over the variables y_1, \dots, y_m , and let $\sigma \in \{0, 1\}$. We denote by $h_{y_j \leftarrow \sigma}$ the function obtained from h by setting y_j to the bit σ . Then:*

- The function $y_j^\sigma \vee h$ is equal to the function $y_j^\sigma \vee h_{y_j \leftarrow \sigma}$.
- The function $y_j^\sigma \wedge h$ is equal to the function $y_j^\sigma \wedge h_{y_j \leftarrow \bar{\sigma}}$.

In order to use the simplification rules, we define, for every internal gate g of ϕ and projection π , an event $\mathcal{E}_{g,\pi}$ as follows: if g is an OR gate, then $\mathcal{E}_{g,\pi}$ is the event that there exists some literal y_j^σ (for $\sigma \in \{0, 1\}$) such that $\text{left}(g)|_\pi = y_j^\sigma$ and $L(\text{right}(g)|_{\pi_{y_j \leftarrow \sigma}}) = 1$. If g is an AND gate, then $\mathcal{E}_{g,\pi}$ is defined similarly, except that we replace $\pi_{y_j \leftarrow \sigma}$ with $\pi_{y_j \leftarrow \bar{\sigma}}$. We have the following lemma, which is proved in Section 3.4.

Lemma 3.13. *For every projection π it holds that*

$$L(\phi|\pi) \leq \sum_{\text{internal gate } g \text{ of } \phi} (\text{depth}(g) + 2) \cdot \mathbf{1}_{\mathcal{E}_{g,\pi}} + \mathbf{1}_{L(\phi|\pi)=1}.$$

We can now use the following corollary of Lemma 3.11 to replace our simplifying assumption.

Corollary 3.14. *For every internal gate g of ϕ it holds that*

$$\Pr[\mathcal{E}_{g,\pi}] \leq q^2 \cdot \sqrt{L(\text{left}(g)) \cdot L(\text{right}(g))}.$$

Proof. Let g be an internal gate of ϕ . We prove the corollary for the case where g is an OR gate, and the proof for the case that g is an AND gate is similar. It holds that

$$\begin{aligned} \Pr[\mathcal{E}_{g,\pi}] &= \Pr \left[\exists \text{ literal } y_j^\sigma : \text{left}(g)|_\pi = y_j^\sigma \text{ and } L(\text{right}(g)|_{\pi_{y_j \leftarrow \sigma}}) = 1 \right] \\ &= \sum_{\text{literal } y_j^\sigma} \Pr \left[\text{left}(g)|_\pi = y_j^\sigma \text{ and } L(\text{right}(g)|_{\pi_{y_j \leftarrow \sigma}}) = 1 \right] \\ &= \sum_{\text{literal } y_j^\sigma} \Pr \left[L(\text{right}(g)|_{\pi_{y_j \leftarrow \sigma}}) = 1 \mid \text{left}(g)|_\pi = y_j^\sigma \right] \cdot \Pr[\text{left}(g)|_\pi = y_j^\sigma] \\ &\leq q \cdot \sqrt{L(\text{right}(g))} \cdot \sum_{\text{literal } y_j^\sigma} \Pr[\text{left}(g)|_\pi = y_j^\sigma] && \text{(Lemma 3.11)} \\ &= q \cdot \sqrt{L(\text{right}(g))} \cdot \Pr[L(\text{left}(g)|_\pi) = 1] \\ &\leq q^2 \cdot \sqrt{L(\text{left}(g)) \cdot L(\text{right}(g))}, && \text{(Lemma 3.8)} \end{aligned}$$

as required. □

The shrinkage theorem now follows using the same calculation as above, replacing Lemma 3.10 with Lemma 3.13 and the simplifying assumption with Corollary 3.14:

$$\begin{aligned}
\mathbb{E}[L(\phi|\pi)] &\leq q \cdot \sqrt{s} + (d+2) \cdot \sum_{\text{internal gate } g \text{ of } \phi} \Pr[\mathcal{E}_{g,\pi}] && \text{(Lemma 3.13)} \\
&\leq q^2 \cdot (d+2) \cdot \sum_{\text{internal gate } g \text{ of } \phi} \sqrt{L(\text{left}(g)) \cdot L(\text{right}(g))} + q \cdot \sqrt{s} && \text{(Corollary 3.14)} \\
&\leq q^2 \cdot (d+2) \cdot \sum_{\text{internal gate } g \text{ of } \phi} (L(\text{left}(g)) + L(\text{right}(g))) + q \cdot \sqrt{s} && \text{(AM-GM inequality)} \\
&= q^2 \cdot (d+2) \cdot \sum_{\text{internal gate } g \text{ of } \phi} \text{size}(g) + q \cdot \sqrt{s} \\
&\leq O(q^2 \cdot d^2 \cdot s + q \cdot \sqrt{s}).
\end{aligned}$$

In the remainder of this section, we prove Lemmas 3.8, 3.11 and 3.13.

Remark 3.15. In this paper, we do not prove Lemma 3.10, since we do not actually need it for our proof. However, this lemma can be established using the proof of Lemma 3.13, with some minor changes.

3.2 Proof of Lemma 3.8

Let $f: \{0,1\}^n \rightarrow \{0,1\}$, and let \mathcal{E} be the set of projections π such that $L(f|\pi) = 1$. We prove that the probability that $\pi \in \mathcal{E}$ is at most $q \cdot \sqrt{L(f)}$. Our proof follows closely the proof of [Hås98, Lemma 4.1].

Let Π be a protocol that solves KW_f and has $L(f)$ leaves (such a protocol exists by Theorem 2.10). Let \mathcal{A} and \mathcal{B} be the sets of projections π for which $f|\pi$ is the constants 1 and 0, respectively. We extend the protocol Π to take inputs from $\mathcal{A} \times \mathcal{B}$ as follows: when Alice and Bob are given as inputs the projections $\pi^A \in \mathcal{A}$ and $\pi^B \in \mathcal{B}$, respectively, they construct strings $a, b \in \{0,1\}^n$ from π^A, π^B by substituting 0 in all the variables y_1, \dots, y_m , and invoke Π on the inputs a and b . Observe that a and b are indeed legal inputs for Π (since $f(a) = 1$ and $f(b) = 0$). Moreover, recall that the protocol Π induces a partition of $\mathcal{A} \times \mathcal{B}$ to combinatorial rectangles, and that we denote the rectangle of the leaf ℓ by $\mathcal{A}_\ell \times \mathcal{B}_\ell$ (see Notation 2.8).

Our proof strategy is the following: We associate with every projection $\pi \in \mathcal{E}$ a leaf of Π , denoted $\text{leaf}(\pi)$. We consider the two disjoint events $\mathcal{E}^+, \mathcal{E}^-$ that correspond to the event that $f|\pi$ is a single positive literal or a single negative literal, respectively, and show that for every leaf ℓ it holds that

$$\Pr[\pi \in \mathcal{E}^+ \text{ and } \text{leaf}(\pi) = \ell] \leq q \cdot \sqrt{\Pr[\pi \in \mathcal{A}_\ell] \cdot \Pr[\pi \in \mathcal{B}_\ell]} \quad (2)$$

$$\Pr[\pi \in \mathcal{E}^- \text{ and } \text{leaf}(\pi) = \ell] \leq q \cdot \sqrt{\Pr[\pi \in \mathcal{A}_\ell] \cdot \Pr[\pi \in \mathcal{B}_\ell]}. \quad (3)$$

Together, the two inequalities imply that

$$\Pr[\pi \in \mathcal{E} \text{ and } \text{leaf}(\pi) = \ell] \leq 2q \cdot \sqrt{\Pr[\pi \in \mathcal{A}_\ell] \cdot \Pr[\pi \in \mathcal{B}_\ell]}.$$

The desired bound on $\Pr[\pi \in \mathcal{E}]$ will follow by summing the latter bound over all the leaves ℓ of Π .

We start by explaining how to associate a leaf with every projection $\pi \in \mathcal{E}^+$. Let $\pi \in \mathcal{E}^+$. Then, it must be the case that $f|\pi = y_j$ for some $j \in [m]$. We define the projections $\pi^1 = \pi_{y_j \leftarrow 1}$

and $\pi^0 = \pi_{y_j \leftarrow 0}$, and observe that $\pi^1 \in \mathcal{A}$ and $\pi^0 \in \mathcal{B}$. We now define $\text{leaf}(\pi)$ to be the leaf to which Π arrives when invoked on inputs π^1 and π^0 . Observe that the output of Π at $\text{leaf}(\pi)$ must be a variable x_i that satisfies $\pi(x_i) \in \{y_j, \bar{y}_j\}$, and thus $\pi_{\text{var}(\pi(x_i)) \leftarrow 1} = \pi^1$.

Next, fix a leaf ℓ . We prove that $\Pr[\boldsymbol{\pi} \in \mathcal{E}^+ \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] \leq q_1 \cdot \Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell]$. Let x_i be the output of the protocol Π at ℓ . Then,

$$\begin{aligned} \Pr[\boldsymbol{\pi} \in \mathcal{E}^+ \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] &\leq \Pr[\boldsymbol{\pi}^1 \in \mathcal{A}_\ell \text{ and } \boldsymbol{\pi}(x_i) \notin \{0, 1\} \text{ and } \boldsymbol{\pi}_{\text{var}(\pi(x_i)) \leftarrow 1} = \boldsymbol{\pi}^1] \\ &= \Pr[\boldsymbol{\pi}(x_i) \notin \{0, 1\} \text{ and } \boldsymbol{\pi}_{\text{var}(\pi(x_i)) \leftarrow 1} \in \mathcal{A}_\ell] \\ &= \sum_{\boldsymbol{\pi} \in \mathcal{A}_\ell} \Pr[\boldsymbol{\pi}(x_i) \notin \{0, 1\} \text{ and } \boldsymbol{\pi}_{\text{var}(\pi(x_i)) \leftarrow 1} = \boldsymbol{\pi}] \\ &\leq q_1 \cdot \sum_{\boldsymbol{\pi} \in \mathcal{A}_\ell} \Pr[\boldsymbol{\pi} = \boldsymbol{\pi}] \quad (\text{since } \boldsymbol{\pi} \text{ is } (q_0, q_1)\text{-fixing}) \\ &= q_1 \cdot \Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell]. \end{aligned}$$

Similarly, it can be proved that $\Pr[\boldsymbol{\pi} \in \mathcal{E}^+ \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] \leq q_0 \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]$. Together, the two bounds imply that

$$\Pr[\boldsymbol{\pi} \in \mathcal{E}^+ \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] \leq \sqrt{q_1 \Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot q_0 \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]} = q \sqrt{\Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]}$$

for every leaf ℓ of Π . We define $\text{leaf}(\pi)$ for projections $\pi \in \mathcal{E}^-$ in an analogous way, and then a similar argument shows that

$$\Pr[\boldsymbol{\pi} \in \mathcal{E}^- \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] \leq q \cdot \sqrt{\Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]}.$$

It follows that

$$\Pr[\boldsymbol{\pi} \in \mathcal{E} \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] \leq 2q \cdot \sqrt{\Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]}.$$

Finally, let \mathcal{L} denote the set of leaves of Π . It holds that

$$\begin{aligned} \Pr[\boldsymbol{\pi} \in \mathcal{E}] &= \sum_{\ell \in \mathcal{L}} \Pr[\boldsymbol{\pi} \in \mathcal{E} \text{ and } \text{leaf}(\boldsymbol{\pi}) = \ell] \\ &\leq 2q \cdot \sum_{\ell \in \mathcal{L}} \sqrt{\Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]} \\ &\leq 2q \cdot \sqrt{|\mathcal{L}|} \cdot \sqrt{\sum_{\ell \in \mathcal{L}} \Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]} \quad (\text{Cauchy-Schwarz – see Fact 2.12}) \\ &= 2q \cdot \sqrt{L(f)} \cdot \sqrt{\sum_{\ell \in \mathcal{L}} \Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell]}. \end{aligned}$$

We conclude the proof by showing that $\sum_{\ell \in \mathcal{L}} \Pr[\boldsymbol{\pi} \in \mathcal{A}_\ell] \cdot \Pr[\boldsymbol{\pi} \in \mathcal{B}_\ell] \leq \frac{1}{4}$. To this end, let $\boldsymbol{\pi}^A, \boldsymbol{\pi}^B$

be two independent random variables that are distributed identically to π . Then, it holds that

$$\begin{aligned}
\sum_{\ell \in \mathcal{L}} \Pr[\pi \in \mathcal{A}_\ell] \cdot \Pr[\pi \in \mathcal{B}_\ell] &= \sum_{\ell \in \mathcal{L}} \Pr[\pi^A \in \mathcal{A}_\ell] \cdot \Pr[\pi^B \in \mathcal{B}_\ell] \\
&= \sum_{\ell \in \mathcal{L}} \Pr[(\pi^A, \pi^B) \in \mathcal{A}_\ell \times \mathcal{B}_\ell] && (\pi^A, \pi^B \text{ are independent}) \\
&= \Pr[(\pi^A, \pi^B) \in \mathcal{A} \times \mathcal{B}] && (\mathcal{A}_\ell \times \mathcal{B}_\ell \text{ are a partition of } \mathcal{A} \times \mathcal{B}) \\
&= \Pr[\pi^A \in \mathcal{A}] \cdot \Pr[\pi^B \in \mathcal{B}] && (\pi^A, \pi^B \text{ are independent}) \\
&= \Pr[\pi \in \mathcal{A}] \cdot \Pr[\pi \in \mathcal{B}] \\
&\leq \Pr[\pi \in \mathcal{A}] \cdot (1 - \Pr[\pi \in \mathcal{A}]). && (\mathcal{A}, \mathcal{B} \text{ are disjoint})
\end{aligned}$$

It is not hard to check that the last expression is always at most $\frac{1}{4}$.

3.3 Proof of Lemma 3.11

Recall that π is a (q_0, q_1) -fixing random projection, and that $q = \sqrt{q_0 q_1}$. Let $f_1, f_2: \{0, 1\}^n \rightarrow \{0, 1\}$, let $\sigma, \tau \in \{0, 1\}$, and let y_j be a variable. We would like to prove that

$$\Pr\left[L(f_2|_{\pi_{y_j \leftarrow \sigma}}) = 1 \mid f_1|_{\pi} = y_j^\tau\right] \leq q \cdot \sqrt{L(f_2)}. \quad (4)$$

For simplicity, we focus on the case that $f_1|_{\pi} = y_j$, and the case that $f_1|_{\pi} = \bar{y}_j$ can be dealt with similarly. The crux of the proof is to show that the random projection $\pi_{y_j \leftarrow \sigma}$ is essentially a (q_0, q_1) -fixing projection even when conditioned on the event $f_1|_{\pi} = y_j$, and therefore Equation (4) is implied immediately by Lemma 3.8.

In order to carry out this argument, we first establish some notation. Let $I^+ = \pi^{-1}(y_j)$ and $I^- = \pi^{-1}(\bar{y}_j)$, and denote by π' the restriction of π to $\{x_1, \dots, x_n\} \setminus (I^+ \cup I^-)$. We denote by f_{2, I^+, I^-} the function over $\{x_1, \dots, x_n\} \setminus (I^+ \cup I^-)$ that is obtained from f_2 by hard-wiring σ and $\bar{\sigma}$ to the variables in I^+ and I^- , respectively. Observe that $f_2|_{\pi_{y_j \leftarrow \sigma}} = f_{2, I^+, I^-}|_{\pi'}$, so it suffices to prove that for every two disjoint sets $I^+, I^- \subseteq \{x_1, \dots, x_n\}$ it holds that

$$\Pr\left[L(f_{2, I^+, I^-}|_{\pi'}) = 1 \mid f_1|_{\pi} = y_j, I^+ = I^+, I^- = I^-\right] \leq q \cdot \sqrt{L(f_2)}. \quad (5)$$

Let $I^+, I^- \subseteq \{x_1, \dots, x_n\}$ be disjoint sets, and let \mathcal{I} be the event that $I^+ = I^+$ and $I^- = I^-$. For convenience, let $K = \{x_1, \dots, x_n\} \setminus (I^+ \cup I^-)$ and $Y = \{y_1, \dots, y_m\} \setminus \{y_j\}$, so π' is a random projection from K to Y when conditioned on \mathcal{I} . To prove Equation (5), it suffices to prove that π' is a (q_0, q_1) -fixing projection when conditioned on the events \mathcal{I} and $f_1|_{\pi} = y_j$, and then the inequality will follow from Lemma 3.8. We first prove that π' is a (q_0, q_1) -fixing projection when conditioning only on the event \mathcal{I} (and not on $f_1|_{\pi} = y_j$).

Proposition 3.16. *Conditioned on the event \mathcal{I} , the projection π' is a (q_0, q_1) -fixing projection.*

Proof. We prove that π' satisfies the definition of a fixing projection. Let π' be a projection from

K to Y , and let $x_i \in K$. Let $\sigma \in \{0, 1\}$. It holds that

$$\begin{aligned}
& \Pr \left[\pi'(x_i) \notin \{0, 1\} \text{ and } \pi'_{\text{var}(\pi'(x_i)) \leftarrow \sigma} = \pi' \mid \mathcal{I} \right] \\
&= \Pr \left[\pi(x_i) \notin \{0, 1\} \text{ and } \pi_{\text{var}(\pi(x_i)) \leftarrow \sigma} \mid K = \pi' \text{ and } \mathcal{I} \right] / \Pr[\mathcal{I}] \\
&= \sum_{\substack{\pi: \pi|_K = \pi', \pi^{-1}(y_j) = I^+, \\ \pi^{-1}(\bar{y}_j) = I^-}} \Pr \left[\pi(x_i) \notin \{0, 1\} \text{ and } \pi_{\text{var}(\pi(x_i)) \leftarrow \sigma} = \pi \right] / \Pr[\mathcal{I}] \\
&\leq \sum_{\substack{\pi: \pi|_K = \pi', \pi^{-1}(y_j) = I^+, \\ \pi^{-1}(\bar{y}_j) = I^-}} q_\sigma \cdot \Pr[\pi = \pi] / \Pr[\mathcal{I}] \quad (\text{since } \pi \text{ is } (q_0, q_1)\text{-fixing}) \\
&= q_\sigma \cdot \Pr \left[\pi|_K = \pi' \text{ and } \mathcal{I} \right] / \Pr[\mathcal{I}] \\
&= q_\sigma \cdot \Pr \left[\pi' = \pi' \mid \mathcal{I} \right],
\end{aligned}$$

as required. \square

We now prove that π' remains a (q_0, q_1) -fixing projection when conditioning on $f_1|_\pi = y_j$ in addition to \mathcal{I} . The crucial observation is that the event $f_1|_\pi = y_j$ is essentially a *filter*, defined next.

Definition 3.17. A set of projections \mathcal{E} from x_1, \dots, x_n to y_1, \dots, y_m is a *filter* if it is closed under assignment to variables, i.e., if for every $\pi \in \mathcal{E}$, every variable y_j , and every bit $\tau \in \{0, 1\}$, it holds that $\pi_{y_j \leftarrow \tau} \in \mathcal{E}$.

It turns out that the property of a projection being a (q_0, q_1) -fixing projection is preserved when conditioning on filters. Formally:

Proposition 3.18. *Let \mathcal{E} be a filter and let π^* be a (q_0, q_1) -fixing projection. Then, $\pi^*|_{\mathcal{E}}$ is a (q_0, q_1) -fixing projection.*

Proof. Let π^* be a projection, and let x_i be a variable. Let $\sigma \in \{0, 1\}$. We would like to prove that

$$\Pr \left[\pi^*(x_i) \notin \{0, 1\} \text{ and } \pi^*_{\text{var}(\pi^*(x_i)) \leftarrow \sigma} = \pi^* \mid \pi^* \in \mathcal{E} \right] \leq q_\sigma \cdot \Pr[\pi^* = \pi^* \mid \pi^* \in \mathcal{E}].$$

If $\pi^* \notin \mathcal{E}$, then both sides of the equation are equal to zero: this is obvious for the right-hand side, and holds for the left-hand side since if there is a projection $\pi^0 \in \mathcal{E}$ and a variable y_j such that $\pi^0_{y_j \leftarrow \sigma} = \pi^*$ then it must be the case that $\pi^* \in \mathcal{E}$ by the definition of a filter. Thus, we may assume that $\pi^* \in \mathcal{E}$. Now, it holds that

$$\begin{aligned}
& \Pr \left[\pi^*(x_i) \notin \{0, 1\} \text{ and } \pi^*_{\text{var}(\pi^*(x_i)) \leftarrow \sigma} = \pi^* \mid \pi^* \in \mathcal{E} \right] \\
&\leq \Pr \left[\pi^*(x_i) \notin \{0, 1\} \text{ and } \pi^*_{\text{var}(\pi^*(x_i)) \leftarrow \sigma} = \pi^* \right] / \Pr[\pi^* \in \mathcal{E}] \\
&\leq q_\sigma \cdot \Pr[\pi^* = \pi^*] / \Pr[\pi^* \in \mathcal{E}] \quad (\pi^* \text{ is } (q_0, q_1)\text{-fixing}) \\
&= q_\sigma \cdot \Pr[\pi^* = \pi^* \mid \pi^* \in \mathcal{E}], \quad (\pi^* \in \mathcal{E})
\end{aligned}$$

as required. \square

Consider the event $f_1|_{\pi} = y_j$. Viewed as a set of projections from x_1, \dots, x_n to y_1, \dots, y_m , this event is not a filter, since it is not closed under assignments to y_j . However, this event is closed under assignments to *all variables except* y_j : when $f_1|_{\pi} = y_j$, the equality continues to hold even if the variables in Y are fixed to constants. Moreover, observe that conditioned on \mathcal{I} , the event $f_1|_{\pi} = y_j$ depends only on the values that π assigns to K . Thus, we can view the event $f_1|_{\pi} = y_j$ as a set of projections from K to $\{y_1, \dots, y_m\} \setminus \{y_j\}$, and taking this view, *this event is a filter*. Since π' is a (q_0, q_1) -fixing projection from K to $\{y_1, \dots, y_m\} \setminus \{y_j\}$ when conditioned on \mathcal{I} , we conclude that it is a (q_0, q_1) -fixing projection when conditioned on both \mathcal{I} and $f_1|_{\pi} = y_j$. It follows by Lemma 3.8 that

$$\Pr[L(f_{2,I^+,I^-}|_{\pi'}) = 1 \mid f_1|_{\pi} = y_j, \mathbf{I}^+ = I^+, \mathbf{I}^- = I^-] \leq q \cdot \sqrt{L(f_{2,I^+,I^-})} \leq q \cdot \sqrt{L(f_2)},$$

as required.

3.4 Proof of Lemma 3.13

Let π be a projection. We prove that

$$L(\phi|_{\pi}) \leq \sum_{\text{internal gate } g \text{ of } \phi} (\text{depth}(g) + 2) \cdot \mathbf{1}_{\mathcal{E}_{g,\pi}} + \mathbf{1}_{L(\phi|_{\pi})=1},$$

where $\mathcal{E}_{g,\pi}$ is the event that there exists some literal y_j^{σ} such that $\text{left}(g)|_{\pi} = y_j^{\sigma}$ and

- $L(\text{right}(g)|_{\pi_{y_j \leftarrow \sigma}}) = 1$ if g is an OR gate, or
- $L(\text{right}(g)|_{\pi_{y_j \leftarrow \bar{\sigma}}}) = 1$ if g is an AND gate.

We prove this claim by induction. If ϕ consists of a single leaf, then the upper bound clearly holds. Otherwise, the root of ϕ is an internal gate. Without loss of generality, assume that the root is an OR gate. We denote the sub-formulas rooted at the left and right children of the root by ϕ_{ℓ} and ϕ_r , respectively. We consider several cases:

- If $L(\phi|_{\pi}) = 1$, then the upper bound clearly holds.
- Suppose that $L(\phi_{\ell}|_{\pi}) \geq 2$. By the subadditivity of formula complexity (Fact 2.3), it holds that

$$L(\phi|_{\pi}) \leq L(\phi_{\ell}|_{\pi}) + L(\phi_r|_{\pi}).$$

By the induction hypothesis, it holds that

$$\begin{aligned} L(\phi_{\ell}|_{\pi}) &\leq \sum_{\text{internal gate } g \text{ of } \phi_{\ell}} (\text{depth}_{\phi_{\ell}}(g) + 2) \cdot \mathbf{1}_{\mathcal{E}_{g,\pi}} + \mathbf{1}_{L(\phi_{\ell}|_{\pi})=1} \\ &= \sum_{\text{internal gate } g \text{ of } \phi_{\ell}} (\text{depth}_{\phi_{\ell}}(g) + 2) \cdot \mathbf{1}_{\mathcal{E}_{g,\pi}} && (L(\phi_{\ell}|_{\pi}) \geq 2) \\ &= \sum_{\text{internal gate } g \text{ of } \phi_{\ell}} (\text{depth}_{\phi}(g) + 1) \cdot \mathbf{1}_{\mathcal{E}_{g,\pi}}, \end{aligned}$$

where the equality holds since $\text{depth}_{\phi_\ell}(g) = \text{depth}_\phi(g) - 1$ for every gate g of ϕ_ℓ . Since $L(\phi_\ell|\pi) \geq 2$, at least one of the terms in the last sum must be non-zero, so it holds that

$$\sum_{\text{internal gate } g \text{ of } \phi_\ell} (\text{depth}_\phi(g) + 1) \cdot 1_{\mathcal{E}_{g,\pi}} \leq \sum_{\text{internal gate } g \text{ of } \phi_\ell} (\text{depth}_\phi(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} - 1.$$

Next, by the induction hypothesis it holds that

$$\begin{aligned} L(\phi_r|\pi) &\leq \sum_{\text{internal gate } g \text{ of } \phi_r} (\text{depth}_{\phi_r}(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} + 1_{L(\phi_r|\pi)=1} \\ &\leq \sum_{\text{internal gate } g \text{ of } \phi_r} (\text{depth}_\phi(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} + 1. \end{aligned}$$

By combining the two bounds, we get that

$$\begin{aligned} L(\phi|\pi) &\leq L(\phi_\ell|\pi) + L(\phi_r|\pi) \\ &\leq \sum_{\text{internal gate } g \text{ of } \phi_\ell} (\text{depth}_\phi(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} - 1 \\ &\quad + \sum_{\text{internal gate } g \text{ of } \phi_r} (\text{depth}_\phi(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} + 1 \\ &\leq \sum_{\text{internal gate } g \text{ of } \phi} (\text{depth}(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} \\ &\leq \sum_{\text{internal gate } g \text{ of } \phi} (\text{depth}(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} + 1_{L(\phi|\pi)=1}, \end{aligned}$$

as required.

- If $L(\phi_r|\pi) \geq 2$, then we use the same argument of the previous case by exchanging ϕ_ℓ and ϕ_r .
- Suppose that $L(\phi|\pi) \geq 2$, $L(\phi_\ell|\pi) \leq 1$ and $L(\phi_r|\pi) \leq 1$. Then, it must be the case that $L(\phi|\pi) = 2$ and also that $L(\phi_\ell|\pi) = 1$ and $L(\phi_r|\pi) = 1$ (or otherwise $L(\phi|\pi) = 1$). In particular, $\phi_\ell|\pi$ is equal to some literal y_j^σ . It follows that $\phi|\pi = y_j^\sigma \vee \phi_r|\pi$, and by the one-variable simplification rules (Fact 3.12), this function is equal to

$$y_j^\sigma \vee (\phi_r|\pi)_{y_j \leftarrow \sigma} = y_j^\sigma \vee \phi_r|\pi_{y_j \leftarrow \sigma}.$$

Thus, it must be the case that $L(\phi_r|\pi_{y_j \leftarrow \sigma}) = 1$ (since $L(\phi|\pi) = 2$). It follows that if we let g be the root of ϕ , then the event $\mathcal{E}_{g,\pi}$ occurs and so

$$(\text{depth}(g) + 2) \cdot 1_{\mathcal{E}_{g,\pi}} = 2 = L(\phi|\pi),$$

so the desired upper bound holds.

We proved that the upper bound holds in each of the possible cases, so the required result follows.

4 Hiding projections

Fixing projections are the most general type of random projection for which our shrinkage theorem holds. In this section we define a new type of random projections, *hiding projections*, which are sometimes easier to understand and construct, and show how to convert them into fixing projections. Intuitively, a hiding projection is a random projection π in which, when given $\pi_{y_j \leftarrow \sigma}$, it is hard to tell in which locations the variable y_j appears in π . Formally, we define q -hiding projections as follows.

Definition 4.1. Let $0 \leq q_0, q_1 \leq 1$. We say that a random projection π is a (q_0, q_1) -*hiding projection* if for every projection π , every bit $\sigma \in \{0, 1\}$, and every variables x_i, y_j , it holds that

$$\Pr[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \pi_{y_j \leftarrow \sigma} = \pi] \leq q_\sigma,$$

whenever the event conditioned on has positive probability. For shorthand, we say that π is a q -hiding projection, for $q = \sqrt{q_0 q_1}$.

To illustrate the definition, consider the following natural random restriction: given n variables x_1, \dots, x_n , the restriction chooses a set of m variables uniformly at random, and fixes all the other variables to random bits. This restriction is not captured by the notion of p -random restrictions or by fixing projections, but as we demonstrate next, it can be implemented by hiding projections. We start with the simple case of $m = 1$, and then consider the general case.

Example 4.2. In order to implement the case of $m = 1$, consider the random projection π from x_1, \dots, x_n to y that is defined as follows: the projection π chooses an index $i \in [n]$ and a bit $\tau \in \{0, 1\}$ uniformly at random, sets $\pi(x_i) = y^\tau$, and sets $\pi(x_{i'})$ to a random bit for all $i' \in [n] \setminus \{i\}$. It is clear that π is essentially equivalent to the random restriction described above for $m = 1$. We claim that π is a $\frac{1}{n}$ -hiding projection. To see it, observe that for every bit $\sigma \in \{0, 1\}$, the projection $\pi_{y \leftarrow \sigma}$ is a uniformly distributed string in $\{0, 1\}^n$, and moreover, this is true conditioned on any possible value of i . In particular, the random variable i is independent of $\pi_{y \leftarrow \sigma}$. Therefore, for every projection $\pi \in \{0, 1\}^n$ and index $i \in [n]$ it holds that

$$\Pr[\pi(x_i) \in \{y, \bar{y}\} \mid \pi_{y \leftarrow \sigma} = \pi] = \Pr[i = i \mid \pi_{y \leftarrow \sigma} = \pi] = \Pr[i = i] = \frac{1}{n},$$

so π satisfies the definition of a (q_0, q_1) -hiding projection with $q_0 = q_1 = \frac{1}{n}$.

Example 4.3. We turn to consider the case of a general $m \in \mathbb{N}$. Let π be the random projection from x_1, \dots, x_n to y_1, \dots, y_m that is defined as follows: the projection π chooses m distinct indices $i_1, \dots, i_m \in [n]$ and m bits $\tau_1, \dots, \tau_m \in \{0, 1\}$ uniformly at random, sets $\pi(x_{i_j}) = y_j^{\tau_j}$ for every $j \in [m]$, and sets all the other variables x_i to random bits. It is clear that π is essentially equivalent to the random projection described above. We show that π is a $\frac{1}{n-m+1}$ -hiding projection. To this end, we should show that for every $i \in [n]$, $j \in [m]$, and $\sigma \in \{0, 1\}$ it holds that

$$\Pr[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \pi_{y_j \leftarrow \sigma} = \pi] \leq \frac{1}{n-m+1}.$$

For simplicity of notation, we focus on the case where $j = m$. Observe that $\pi_{y_m \leftarrow \sigma}$ reveals the values of i_1, \dots, i_{m-1} , and the projection of $\pi_{y_m \leftarrow \sigma}$ on the remaining $n - m + 1$ indices is uniform over

$\{0, 1\}^{n-m+1}$. Moreover, the latter assertion remains true conditioned on any possible value of \mathbf{i}_m (which must be different from the known values of $\mathbf{i}_1, \dots, \mathbf{i}_{m-1}$). Therefore given $\mathbf{i}_1, \dots, \mathbf{i}_{m-1}$, the random variable \mathbf{i}_m (ranging over all indices other than $\mathbf{i}_1, \dots, \mathbf{i}_{m-1}$) is independent of the projection of $\pi_{y_m \leftarrow \sigma}$ on the $n - m + 1$ indices other than $\mathbf{i}_1, \dots, \mathbf{i}_{m-1}$. It follows that for every projection π in the support of $\pi_{y_m \leftarrow \sigma}$ and every index $i \in [n]$, if with $\pi(i_1) = y_1^{\tau_1}, \dots, \pi(i_{m-1}) = y_{m-1}^{\tau_{m-1}}$ then:

$$\begin{aligned} \Pr[\pi(x_i) \in \{y_m, \overline{y_m}\} \mid \pi_{y_m \leftarrow \sigma} = \pi] &= \Pr[\mathbf{i}_m = i \mid \pi_{y_m \leftarrow \sigma} = \pi] \\ &= \Pr[\mathbf{i}_m = i \mid \mathbf{i}_1 = i_1, \dots, \mathbf{i}_{m-1} = i_{m-1}] \leq \frac{1}{n - m + 1}, \end{aligned}$$

as required.

As mentioned above, hiding projections can be converted to fixing projections. This is captured by the following result.

Lemma 4.4. *Let π be a q -hiding projection from x_1, \dots, x_n to y_1, \dots, y_m . Then, there exists an $(4m^2 \cdot q)$ -fixing projection π' from x_1, \dots, x_n to y_1, \dots, y_m such that $\Pr[\pi' = \pi] \geq \frac{1}{2}$. Furthermore, the event $\pi' = \pi$ is independent of π .*

We prove Lemma 4.4 in Section 4.1. Note that the conversion from hiding to fixing projections loses a factor of m^2 in the parameter q . While this loss might be large in general, it will be of no importance in our applications. The factor m^2 can be improved to m in our actual applications, as we indicate during the proof of Lemma 4.4. By combining Lemma 4.4 with our shrinkage theorem for fixing projections, we obtain the following shrinkage theorem for hiding projections.

Theorem 4.5. *Let ϕ be a formula of size s and depth d , and let π be a q -hiding projection. Then*

$$\mathbb{E}[L(\phi|\pi)] = O(m^4 \cdot q^2 \cdot d^2 \cdot s + m^2 \cdot q \cdot \sqrt{s}).$$

Proof. Let π' be the $(m^2 \cdot q)$ -fixing projection obtained from π by Lemma 4.4. Since $\Pr[\pi' = \pi] \geq \frac{1}{2}$, it holds that

$$\mathbb{E}[L(\phi|\pi')] \geq \Pr[\pi' = \pi] \cdot \mathbb{E}[L(\phi|\pi') \mid \pi' = \pi] \geq \frac{1}{2} \cdot \mathbb{E}[L(\phi|\pi) \mid \pi' = \pi] = \frac{1}{2} \cdot \mathbb{E}[L(\phi|\pi)],$$

where the last equality holds since the event $\pi' = \pi$ is independent of π . On the other hand, by applying Theorem 3.6 to π' , we obtain that

$$\mathbb{E}[L(\phi|\pi')] = O(m^4 \cdot q^2 \cdot d^2 \cdot s + m^2 \cdot q \cdot \sqrt{s}).$$

The theorem follows by combining the two bounds. \square

Applying formula balancing, we can obtain an analog of Corollary 3.7, with an identical proof.

Corollary 4.6. *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a function with formula complexity s , and let π be a q -hiding projection. Then*

$$\mathbb{E}[L(f|\pi)] = m^4 \cdot q^2 \cdot s^{1+O(\frac{1}{\sqrt{\log s}})} + m^2 \cdot q \cdot s^{1/2+O(\frac{1}{\sqrt{\log s}})}.$$

Remark 4.7. The following example shows that the loss of a factor of at least m^2 is necessary for Theorem 4.5. Let $f = \text{Parity}_n$ be the Parity function over x_1, \dots, x_n , and note that the formula complexity of f is $s = \Theta(n^2)$. Let π be the random projection from Example 4.3, and recall that it is a q -hiding projection for $q = \frac{1}{n-m+1}$. Then, $f|_\pi$ is the Parity function over m bits, and therefore its formula complexity is $\Theta(m^2)$. It follows that when $m \leq n/2$,

$$\mathbb{E}[L(f|_\pi)] = \Theta(m^2) = \Theta(m^2 \cdot q^2 \cdot s).$$

4.1 Proof of Lemma 4.4

We use the following straightforward generalization of the property of hiding projections.

Claim 4.8. *Let π be (q_0, q_1) -hiding projection, and let \mathcal{E} be a random set of projections that is independent of π . Then, for every $\sigma \in \{0, 1\}$, it holds that*

$$\Pr[\pi(x_i) \in \{y_j, \overline{y_j}\} \mid \pi_{y_j \leftarrow \sigma} \in \mathcal{E}] \leq q_\sigma.$$

The straightforward proof of Claim 4.8 works by applying the property of hiding projections separately to each possible value of \mathcal{E} and each possible projection $\pi \in \mathcal{E}$, and can be found in Appendix A.

We turn to proving Lemma 4.4. Suppose that π is a (q_0, q_1) -hiding projection from x_1, \dots, x_n to y_1, \dots, y_m . Let ρ be a $(1 - \frac{1}{2m})$ -random restriction over y_1, \dots, y_m , i.e., ρ is the random projection from y_1, \dots, y_m to y_1, \dots, y_m that assigns each y_j independently as follows:

$$\rho(y_j) = \begin{cases} y_j & \text{with probability } 1 - \frac{1}{2m}, \\ 0 & \text{with probability } \frac{1}{4m}, \\ 1 & \text{with probability } \frac{1}{4m}. \end{cases}$$

For convenience, we define $\rho(0) = 0$, $\rho(1) = 1$, and $\rho(\overline{y_j}) = \overline{\rho(y_j)}$ for every $j \in [m]$. We now choose the random projection π' to be the composition $\rho \circ \pi$. Observe that the event $\pi' = \pi$ occurs whenever it holds that $\rho(y_j) = y_j$ for every $j \in [m]$, and therefore

$$\Pr[\pi' = \pi] \geq \left(1 - \frac{1}{2m}\right)^m \geq 1 - \frac{m}{2m} = \frac{1}{2},$$

as required. Moreover, this event is independent of π , since ρ is independent of π . We show that π' is a $(4 \cdot m^2 \cdot q_0, 4 \cdot m^2 \cdot q_1)$ -fixing projection. To this end, we should show that for every projection π' , every bit $\sigma \in \{0, 1\}$, and every variable x_i ,

$$\Pr[\pi'(x_i) \notin \{0, 1\} \text{ and } \pi'_{\text{var}(\pi'(x_i)) \leftarrow \sigma} = \pi'] \leq 4 \cdot m^2 \cdot q_\sigma \cdot \Pr[\pi' = \pi']. \quad (6)$$

This is implied by the following inequality, which we will prove for all $j \in [m]$:

$$\Pr[\pi'(x_i) \in \{y_j, \overline{y_j}\} \text{ and } \pi'_{y_j \leftarrow \sigma} = \pi'] \leq 4 \cdot m \cdot q_\sigma \cdot \Pr[\pi' = \pi']. \quad (7)$$

Let $j \in [m]$ and let ρ_{-j} be the restriction of ρ to $\{y_1, \dots, y_m\} \setminus \{y_j\}$. First, observe that if $\rho(y_j) = \sigma$ then $\pi' = \rho \circ \pi = \rho_{-j} \circ \pi_{y_j \leftarrow \sigma}$, and therefore

$$\Pr[\pi' = \pi'] \geq \Pr[\rho(y_j) = \sigma \text{ and } \rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'] = \frac{1}{4m} \cdot \Pr[\rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'],$$

where the equality holds since $\rho(y_j)$ is independent of ρ_{-j} and π . In the rest of this section we will prove the following inequality, which together with the last inequality will imply Equation (7):

$$\Pr\left[\pi'(x_i) \in \{y_j, \bar{y}_j\} \text{ and } \pi'_{y_j \leftarrow \sigma} = \pi'\right] \leq q_\sigma \cdot \Pr[\rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi']. \quad (8)$$

To this end, observe that the event $\pi'(x_i) \in \{y_j, \bar{y}_j\}$ happens if and only if $\pi(x_i) \in \{y_j, \bar{y}_j\}$ and $\rho(y_j) = y_j$, and therefore

$$\begin{aligned} & \Pr\left[\pi'(x_i) \in \{y_j, \bar{y}_j\} \text{ and } \pi'_{y_j \leftarrow \sigma} = \pi'\right] \\ &= \Pr\left[\pi(x_i) \in \{y_j, \bar{y}_j\} \text{ and } \rho(y_j) = y_j \text{ and } \pi'_{y_j \leftarrow \sigma} = \pi'\right]. \end{aligned}$$

Next, observe that if $\rho(y_j) = y_j$ then $\pi'_{y_j \leftarrow \sigma} = \rho_{-j} \circ \pi_{y_j \leftarrow \sigma}$. It follows that the latter expression is equal to

$$\begin{aligned} & \Pr\left[\pi(x_i) \in \{y_j, \bar{y}_j\} \text{ and } \rho(y_j) = y_j \text{ and } \rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'\right] \\ & \leq \Pr\left[\pi(x_i) \in \{y_j, \bar{y}_j\} \text{ and } \rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'\right] \\ & = \Pr\left[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'\right] \cdot \Pr\left[\rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'\right] \\ & \leq q_\sigma \cdot \Pr\left[\rho_{-j} \circ \pi_{y_j \leftarrow \sigma} = \pi'\right], \end{aligned}$$

where the last inequality follows by applying Claim 4.8 with \mathcal{E} being the set of projections π that satisfy $\rho_{-j} \circ \pi = \pi'$. This concludes the proof of Equation (8).

Remark 4.9. We can improve the bound m^2 in the statement of Lemma 4.4 to mk , where k is the maximal number of variables $1, \dots, m$ which could appear in any position of π . The reason is that in the latter case, the transition from Equation (7) to Equation (6) incurs a factor of k rather than m . This is useful, for example, since the random projections π of Section 7.1 have the feature that for each $i \in [n]$ there is a unique $j \in [m]$ such that $\pi(x_i) \in \{0, 1, j, \bar{j}\}$, and so for these projections $k = 1$.

5 Joining projections

In this section, we define a *join* operation on fixing and hiding projections, and show that it preserves the corresponding properties. This operation provides a convenient tool for constructing random projections, and will be used in our applications.

Definition 5.1. Let α be a random projection from x_1, \dots, x_{n_a} to y_1, \dots, y_{m_a} , and let β be a random projection from w_1, \dots, w_{n_b} to z_1, \dots, z_{m_b} . The *join* $\alpha \uplus \beta$ is the random projection from $x_1, \dots, x_{n_a}, w_1, \dots, w_{n_b}$ to $y_1, \dots, y_{m_a}, z_1, \dots, z_{m_b}$ obtained by sampling α, β independently and joining them together in the obvious way.

Lemma 5.2. *If α and β are (q_0, q_1) -fixing projections, then so is $\alpha \uplus \beta$.*

Proof. Let α and β be (q_0, q_1) -fixing projections, and let $\gamma = \alpha \uplus \beta$. We prove that γ is a (q_0, q_1) -fixing projection. Let α be a projection from x_1, \dots, x_{n_a} to y_1, \dots, y_{m_a} , let β be a projection from

w_1, \dots, w_{n_b} to z_1, \dots, z_{m_b} , and let $\gamma = \alpha \uplus \beta$. We should show that for every $\sigma \in \{0, 1\}$ and every input variable u (either x_i or w_i) it holds that

$$\Pr\left[\gamma(u) \notin \{0, 1\} \text{ and } \gamma_{\text{var}(\gamma(u)) \leftarrow \sigma} = \gamma\right] \leq q_\sigma \cdot \Pr[\gamma = \gamma].$$

Let $\sigma \in \{0, 1\}$ be a bit, and let u be an input variable. Assume that $u = x_i$ for some $i \in [n_a]$ (if $u = w_i$ for some $i \in [n_b]$, the proof is similar). The above equation is therefore equivalent to

$$\Pr\left[\alpha(x_i) \notin \{0, 1\} \text{ and } \alpha_{\text{var}(\alpha(x_i)) \leftarrow \sigma} = \alpha \text{ and } \beta_{\text{var}(\alpha(x_i)) \leftarrow \sigma} = \beta\right] \leq q_\sigma \cdot \Pr[\alpha = \alpha \text{ and } \beta = \beta].$$

Since the ranges of α and β are disjoint, the variable $\text{var}(\alpha(x_i))$ does not appear in the range of β , and therefore $\beta_{\text{var}(\alpha(x_i)) \leftarrow \sigma} = \beta$. The independence of α and β shows that the above inequality is equivalent to

$$\Pr\left[\alpha(x_i) \notin \{0, 1\} \text{ and } \alpha_{\text{var}(\alpha(x_i)) \leftarrow \sigma} = \alpha\right] \cdot \Pr[\beta = \beta] \leq q_\sigma \cdot \Pr[\alpha = \alpha] \cdot \Pr[\beta = \beta],$$

which follows from α being (q_0, q_1) -fixing. □

Lemma 5.3. *If α and β are (q_0, q_1) -hiding projections, then so is $\alpha \uplus \beta$.*

Proof. Let α and β be (q_0, q_1) -hiding projections, and let $\gamma = \alpha \uplus \beta$. We prove that γ is a (q_0, q_1) -hiding projection. Let α be a projection from x_1, \dots, x_{n_a} to y_1, \dots, y_{m_a} , let β be a projection from w_1, \dots, w_{n_b} to z_1, \dots, z_{m_b} , and let $\gamma = \alpha \uplus \beta$. We should show that for every $\sigma \in \{0, 1\}$, every input variable u (either x_i or w_i), and every output variable v (either y_j or z_j) it holds that

$$\Pr[\gamma(u) \in \{v, \bar{v}\} \mid \gamma_{v \leftarrow \sigma} = \gamma] \leq q_\sigma.$$

Let $\sigma \in \{0, 1\}$ be a bit, let u be an input variable, and let v be an output variable. Assume that $u = x_i$ for some $i \in [n_a]$ (if $u = w_i$ for some $i \in [n_b]$, the proof is similar). In this case, we may assume that $v = y_j$ for some $j \in [m_a]$, since the probability on the left-hand side is 0. Thus, the above equation is equivalent to

$$\Pr\left[\alpha(x_i) \in \{y_j, \bar{y}_j\} \mid \alpha_{y_j \leftarrow \sigma} = \alpha \text{ and } \beta_{y_j \leftarrow \sigma} = \beta\right] \leq q_\sigma.$$

Since α and β are independent, whenever the conditioned event has positive probability, it holds that

$$\Pr\left[\alpha(x_i) \in \{y_j, \bar{y}_j\} \mid \alpha_{y_j \leftarrow \sigma} = \alpha \text{ and } \beta_{y_j \leftarrow \sigma} = \beta\right] = \Pr[\alpha(x_i) \in \{y_j, \bar{y}_j\} \mid \alpha_{x_i \leftarrow \sigma} = \alpha] \leq q_\sigma,$$

where the inequality holds since α is (q_0, q_1) -hiding. □

6 Hiding projections from complexity measures

In this section we define generalizations of the complexity measures due to Khrapchenko [Khr72] and Ambainis [Amb02], and show how they can be used to construct hiding projections.

Along the way, we prove several auxiliary results on the generalized complexity measures. Since these results are not necessary for proving our main results, we have relegated them to appendices.

6.1 The min-entropy Khrapchenko bound

Khrapchenko [Khr72] defined the following complexity measure for Boolean functions, and proved that it is a lower bound on formula complexity.

Definition 6.1. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. For every two sets $A \subseteq f^{-1}(1)$ and $B \subseteq f^{-1}(0)$, let $E(A, B)$ be the set of pairs $(a, b) \in A \times B$ such that a and b differ on exactly one coordinate. The *Khrapchenko bound* of f is

$$\text{Khr}(f) = \max_{A \subseteq f^{-1}(1), B \subseteq f^{-1}(0)} \frac{|E(A, B)|^2}{|A| \cdot |B|}.$$

Theorem 6.2 ([Khr72]). *For every non-constant function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $L(f) \geq \text{Khr}(f)$.*

The Khrapchenko measure can be viewed as follows: Consider the subgraph of the Hamming cube that consists only of the cut between A and B . Then, the measure $\frac{|E(A, B)|^2}{|A| \cdot |B|}$ is the product of the average degree of a vertex in A (which is $\frac{|E(A, B)|}{|A|}$) and the average degree of a vertex in B (which is $\frac{|E(A, B)|}{|B|}$). Note that the average degree of A can also be described as the average, over all strings $a \in A$, of the number of coordinates $i \in [n]$ such that if we flip the i -th bit of a we get a string in B . We generalize the Khrapchenko measure as follows:

- Whereas the Khrapchenko bound maximizes over all *cuts* (A, B) of the Hamming cube with $A \subseteq f^{-1}(1)$ and $B \subseteq f^{-1}(0)$, we are maximizing over all *distributions* over edges (\mathbf{a}, \mathbf{b}) of the Hamming cube, where $\mathbf{a} \in f^{-1}(1)$ and $\mathbf{b} \in f^{-1}(0)$.
- Whereas the Khrapchenko bound considers the average number of coordinates i as described above, we consider the min-entropy of the coordinate i that is associated with the random edge (\mathbf{a}, \mathbf{b}) .
- Whereas the Khrapchenko bound considers functions whose inputs are binary strings, we consider inputs that are strings over arbitrary finite alphabets.

Our generalization uses the following notion: The *conditional min-entropy* of a random variable \mathbf{x} conditioned on a random variable \mathbf{y} is $H_\infty(\mathbf{x} \mid \mathbf{y}) = \min_{x, y} \log \frac{1}{\Pr[\mathbf{x}=x \mid \mathbf{y}=y]}$. We can now define our generalization formally:

Definition 6.3. Let Σ be a finite alphabet, and let $f: \Sigma^n \rightarrow \{0, 1\}$ be a Boolean function. We say that a distribution (\mathbf{a}, \mathbf{b}) on $f^{-1}(1) \times f^{-1}(0)$ is a *Khrapchenko distribution* for f if \mathbf{a} and \mathbf{b} always differ on a unique coordinate $i \in [n]$. We define the *min-entropy Khrapchenko bound* of f , denoted $\text{Khr}_{H_\infty}(f)$, to be the maximum of the quantity

$$2^{H_\infty(i \mid \mathbf{a}) + H_\infty(i \mid \mathbf{b})}$$

over all Khrapchenko distributions (\mathbf{a}, \mathbf{b}) for f .

Observe that when $\Sigma = \{0, 1\}$, the min-entropy $H_\infty(i \mid \mathbf{a})$ is exactly the min-entropy of a random neighbor of \mathbf{a} . In particular, when (\mathbf{a}, \mathbf{b}) is the uniform distribution over a set of edges, the min-entropy $H_\infty(i \mid \mathbf{a})$ is the logarithm of the minimal degree of the vertex \mathbf{a} . Moreover, if

the latter set of edges also induces a regular graph, then the measure $2^{H_\infty(\mathbf{i}|\mathbf{a})+H_\infty(\mathbf{i}|\mathbf{b})}$ coincides exactly with the original measure $\frac{|E(A,B)|^2}{|A|\cdot|B|}$ of Khrapchenko. More generally, when $\Sigma = \{0, 1\}$, the bound $\text{Khr}_{H_\infty}(f)$ is within a constant factor of the original Khrapchenko bound, as we show in Appendix B.

Proposition 6.4. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $\frac{\text{Khr}(f)}{4} \leq \text{Khr}_{H_\infty}(f) \leq \text{Khr}(f)$.*

Unfortunately, when Σ is a larger alphabet then the connection between $\text{Khr}_{H_\infty}(f)$ and the measure $\frac{|E(A,B)|^2}{|A|\cdot|B|}$ is not so clean. Specifically, the min-entropy $H_\infty(\mathbf{i} | \mathbf{a})$ has no clear connection to the degree of \mathbf{a} , since the vertex \mathbf{a} may have multiple neighbors that correspond to the same coordinate \mathbf{i} . Nevertheless, even when Σ is a large alphabet, the min-entropy Khrapchenko bound $\text{Khr}_{H_\infty}(f)$ still gives a lower bound on the formula complexity $L(f)$. In order to formalize the latter statement, we first need to deal with the following small issue: formula complexity is not defined for functions over a *non-binary* alphabet Σ . Thus, for the purpose of formula complexity, we view a function $f: \Sigma^n \rightarrow \{0, 1\}$ over a non-binary finite alphabet Σ as if its input is a binary string in $\{0, 1\}^{n \cdot \lceil \log |\Sigma| \rceil}$ that encodes a string in Σ^n via some fixed encoding (the choice of the encoding does not matter for what follows). We now have the following result, proved in Appendix C.⁴

Proposition 6.5. *Let Σ be a finite alphabet, and let $f: \Sigma^n \rightarrow \{0, 1\}$ be a Boolean function. Then $L(f) \geq \text{Khr}_{H_\infty}(f)$.*

The min-entropy Khrapchenko bound can be used for constructing hiding projections. Specifically, we have the following result.

Lemma 6.6. *Let Σ be a finite alphabet, let $f: \Sigma^n \rightarrow \{0, 1\}$ be a Boolean function. There is a random projection π to a single variable y such that π is q -hiding for $q = \sqrt{1/\text{Khr}_{H_\infty}(f)}$, and $f|_\pi$ is a non-constant function with probability 1.*

Lemma 6.6 follows immediately from Lemmas 6.12 and 6.14, which are proved in Section 6.2 below. We will use this construction in our applications in Section 7.

Previous works on the Khrapchenko bound. Several versions of the Khrapchenko bound appeared in the literature: Zwick [Zwi91] generalized the Khrapchenko bound such that different input coordinates can be given different weights, and Koutsoupias [Kou93] gave a spectral generalization of the bound. The paper of Håstad [Hås98] observed that his analogue of Lemma 3.8 can be viewed as a generalization of the Khrapchenko bound. Ganor, Komargodski, and Raz [GKR12] considered a variant of the Khrapchenko bound in which the edges of the Boolean hypercube are replaced with random walks on the noisy hypercube. Of particular relevance is a paper of Laplante, Lee, and Szegedy [LLS06] that defined a complexity measure that is very similar to our min-entropy Khrapchenko bound, except that the entropy is replaced with Kolmogorov complexity.

An entropy Khrapchenko bound. It is possible to generalize the complexity measure Khr_{H_∞} by replacing the min-entropy in Definition 6.3 with Shannon entropy. Such a measure would still lower bound formula complexity — specifically, the proof of Proposition 6.5 would go through without a change. However, we do not know how to use such a measure for constructing hiding projections as in Lemma 6.6. We note that it is easy to prove that such a measure is an upper bound on Khr_{H_∞} .

⁴We note that Lemma 6.6 along with Lemma 4.4 and Lemma 3.8 imply Proposition 6.5 up to a constant factor (i.e., $L(f) \geq \frac{1}{64}\text{Khr}_{H_\infty}(f)$).

6.2 The soft-adversary method

Ambainis [Amb02] defined the following complexity measure for Boolean functions, called the *unweighted quantum adversary bound* and proved that it is a lower bound on quantum query complexity (for a definition of quantum query complexity, see [Amb02] or [BdW02]).

Definition 6.7. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. Let $R \subseteq f^{-1}(1) \times f^{-1}(0)$, and let A, B be the projections into the first and second coordinates, respectively. Let $R(a, B) = \{b \in B : (a, b) \in R\}$, let $R_i(a, B) = \{b \in B : (a, b) \in R, a_i \neq b_i\}$, and define $R(A, b), R_i(A, b)$ analogously. The *unweighted quantum adversary bound* of f is

$$\text{Adv}_u(f) = \max_{R \subseteq f^{-1}(1) \times f^{-1}(0)} \frac{\min_{a \in A} |R(a, B)| \cdot \min_{b \in B} |R(A, b)|}{\max_{\substack{a \in A \\ i \in [n]}} |R_i(a, B)| \cdot \max_{\substack{b \in B \\ i \in [n]}} |R_i(A, b)|}.$$

Theorem 6.8 ([Amb02]). *The quantum query complexity of a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is $\Omega(\sqrt{\text{Adv}_u(f)})$.*

In contrast to Khrapchenko’s bound which only considers inputs differing in a single position, the unweighted quantum adversary bound allows the relation R to contain inputs at arbitrary distance. On the other hand, the unweighted quantum adversary method is not a strict generalization of Khrapchenko’s bound since Khrapchenko’s bound measures average degree whereas the unweighted quantum adversary bound measures minimum degree. Nonetheless, since changing average degree to minimum degree incurs only a constant factor to the Khrapchenko measure (as stated in Proposition 6.4), the two bounds are related as follows (for a formal proof, see Appendix B).

Proposition 6.9. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $\text{Adv}_u(f) \geq \frac{\text{Khr}(f)}{4}$.*

We generalize the unweighted quantum adversary method to a min-entropy variant along the lines of the min-entropy Khrapchenko bound. In this case, we can stick to the binary input alphabet.

Definition 6.10. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. We define the *soft-adversary bound* of f , denoted Adv_s , to be the maximum of the quantity

$$\min_{\substack{a \in \text{sup } \mathbf{a} \\ i \in [n]}} \frac{1}{\Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{a} = a]} \cdot \min_{\substack{b \in \text{sup } \mathbf{b} \\ i \in [n]}} \frac{1}{\Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{b} = b]}$$

over all distributions (\mathbf{a}, \mathbf{b}) supported on $f^{-1}(1) \times f^{-1}(0)$.

We chose to call this notion “soft-adversary bound” since we view it as a variant of the unweighted adversary bound in which, instead of choosing for each pair (a, b) whether it belongs to R or not (a “hard” decision), we assign each pair (a, b) a probability of being in the relation R (a “soft” decision). Indeed, the soft-adversary bound generalizes the unweighted quantum adversary method.

Lemma 6.11. *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. Then $\text{Adv}_s(f) \geq \text{Adv}_u(f)$.*

Proof. Let $R \subseteq f^{-1}(1) \times f^{-1}(0)$ be a relation that attains $\text{Adv}_u(f)$. Let (\mathbf{a}, \mathbf{b}) be a uniform element of R . Using $\Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{a} = a] = |R_i(a, B)|/|R(a, B)|$ and $\Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{b} = b] = |R_i(A, b)|/|R(A, b)|$, it is easy to check that the quantity in the definition of $\text{Adv}_s(f)$ is lower bounded by $\text{Adv}_u(f)$. \square

The soft-adversary bound also generalizes the min-entropy Khrapchenko bound.

Lemma 6.12. *Let Σ be a finite alphabet, and let $f: \Sigma^n \rightarrow \{0, 1\}$ be a Boolean function. Then $\text{Adv}_s(f) \geq \text{Khr}_{H_\infty}(f)$, where for the sake of defining $\text{Adv}_s(f)$, we treat f as a function on $\{0, 1\}^{n \cdot \lceil \log |\Sigma| \rceil}$ as described in Section 6.1.*

Proof. Let (\mathbf{a}, \mathbf{b}) be a Khrapchenko distribution that attains $\text{Khr}_{H_\infty}(f)$. Let $i \in [n]$ be the unique index at which \mathbf{a}, \mathbf{b} differ, and let $\mathbf{a}_{i,j}, \mathbf{b}_{i,j}$ be the j -th bits of the binary encoding of $\mathbf{a}_i, \mathbf{b}_i \in \Sigma$. For any $a \in \text{sup } \mathbf{a}$ and $(i, j) \in [n] \times [\lceil \log |\Sigma| \rceil]$, we have

$$\Pr[\mathbf{a}_{i,j} \neq \mathbf{b}_{i,j} \mid \mathbf{a} = a] \leq \Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{a} = a] = \Pr[\mathbf{i} = i \mid \mathbf{a} = a] \leq 2^{-H_\infty(i|\mathbf{a})}.$$

It follows that the first factor in the definition of Adv_s is at least $2^{H_\infty(i|\mathbf{a})}$. An identical calculation shows that the second factor is at least $2^{H_\infty(i|\mathbf{b})}$, and so the entire expression is at least $\text{Khr}_{H_\infty}(f)$. \square

We note that the soft-adversary bound is a lower bound on formula complexity. The reason is that this bound is a special case of the complexity measure WA^2 due to [ŠS06]. Since [LLS06] show that WA^2 is a lower bound on formula complexity, the same holds for the soft-adversary bound. We provide more details in Appendix D. ⁵

Proposition 6.13. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $L(f) \geq \text{Adv}_s(f)$.*

We turn to show how to use the soft-adversary bound for constructing hiding projections, generalizing Lemma 6.6.

Lemma 6.14. *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. There is a random projection π from x_1, \dots, x_n to a single variable y such that π is q -hiding for $q = \sqrt{1/\text{Adv}_s(f)}$, and $f|_\pi$ is a non-constant function with probability 1.*

Proof. Let (\mathbf{a}, \mathbf{b}) be a distribution supported on $f^{-1}(1) \times f^{-1}(0)$ which attains $\text{Adv}_s(f)$. We construct a q -hiding projection π from x_1, \dots, x_n to a single variable y for $q = \sqrt{1/\text{Adv}_s(f)}$, such that $\pi_{y \leftarrow 1} = \mathbf{a}$ and $\pi_{y \leftarrow 0} = \mathbf{b}$. Note that this implies in particular that $f|_\pi$ is a non-constant function, since $\mathbf{a} \in f^{-1}(1)$ and $\mathbf{b} \in f^{-1}(0)$. For every input variable x_i of f , we define π as follows:

- If $\mathbf{a}_i = \mathbf{b}_i$ then $\pi(x_i) = \mathbf{a}_i$.
- If $\mathbf{a}_i = 1$ and $\mathbf{b}_i = 0$ then $\pi(x_i) = y$.
- If $\mathbf{a}_i = 0$ and $\mathbf{b}_i = 1$ then $\pi(x_i) = \bar{y}$.

It is not hard to see that indeed $\pi_{y \leftarrow 1} = \mathbf{a}$ and $\pi_{y \leftarrow 0} = \mathbf{b}$.

We now show that π is $\sqrt{1/\text{Adv}_s(f)}$ -hiding. To this end, we show that π is (q_0, q_1) -hiding for

$$q_0 = \max_{\substack{b \in \text{sup } \mathbf{b} \\ i \in [n]}} \Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{b} = b], \quad q_1 = \max_{\substack{a \in \text{sup } \mathbf{a} \\ i \in [n]}} \Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{a} = a].$$

⁵We note that Lemma 6.14 along with Lemma 4.4 and Lemma 3.8 implies Proposition 6.13 up to a constant factor (i.e., $L(f) \geq \frac{1}{64} \text{Adv}_s(f)$).

Note that indeed $\sqrt{q_0 q_1} = \sqrt{1/\text{Adv}_s(f)}$. We should show that for every projection π and every input variable x_i ,

$$\Pr[\pi(x_i) \in \{y, \bar{y}\} \mid \pi_{y \leftarrow 1} = \pi] \leq \max_{\substack{a \in \text{sup } \mathbf{a} \\ i \in [n]}} \Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{a} = a],$$

and a similar inequality for $\pi_{y \leftarrow 0}$. Since $\pi_{y \leftarrow 1} = \mathbf{a}$, we only need to consider projections π that are equal to some string a in the support of \mathbf{a} . Furthermore, $\pi(x_i) \in \{y, \bar{y}\}$ if and only if $\mathbf{a}_i \neq \mathbf{b}_i$. The desired inequality immediately follows. The inequality for $\pi_{y \leftarrow 1}$ is proved in an identical fashion. \square

Lemma 6.6 immediately follows by Lemma 6.12.

Other generalizations of the quantum adversary method. Several generalizations of Ambainis' original quantum adversary bound have appeared in the literature. Špalek and Szegedy [ŠS06] showed that all of these methods are equivalent, and so they are known collectively as the strong quantum adversary bound. The formulation of the strong quantum adversary bound in Ambainis [Amb06] makes it clear that it subsumes our min-entropy version, as we explain in Appendix D. Laplante, Lee and Szegedy [LLS06] showed that the strong quantum adversary bound lower bounds formula complexity, and came up with an even stronger measure $\max\text{PI}^2$ that still lower bounds formula complexity, but no longer lower bounds quantum query complexity.

Høyer, Lee and Špalek [HyLŠ07] generalized the strong quantum adversary bound, coming up with a new measure known as the general adversary bound, which lower bounds both quantum query complexity and (after squaring) formula complexity. Reichardt [Rei09, Rei11, LMR⁺11, Rei14] showed that the general adversary bound in fact *coincides* with quantum query complexity (up to constant factors). These results are described in a recent survey by Li and Shirley [LS20].

Remark 6.15. We note that $\text{Adv}_s(f)$ is always upper bounded by n^2 . The reason is that, as noted above, $\sqrt{\text{Adv}_s(f)}$ is a lower bound on the complexity measure of WA of [ŠS06]. Since the latter measure is a lower bound on quantum query complexity, which in turn is upper bounded by n , it follows that $\text{Adv}_s(f) \leq n^2$. In particular, the soft-adversary bound $\text{Adv}_s(f)$ on its own cannot directly prove formula lower bounds that are better than quadratic.

7 Applications

In this section, we apply our shrinkage theorems to obtain new results regarding the KRW conjecture and the formula complexity of \mathbf{AC}^0 . First, in Section 7.1, we prove the KRW conjecture for inner functions for which the min-entropy Khrapchenko bound is tight. We use this version of the KRW conjecture in Section 7.2 to prove cubic formula lower bounds for a function in \mathbf{AC}^0 . Finally, we rederive some closely related known results in Section 7.3.

7.1 Application to the KRW conjecture

Given two Boolean functions $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and $g: \{0, 1\}^n \rightarrow \{0, 1\}$, their (block-)composition is the function $f \diamond g: (\{0, 1\}^n)^m \rightarrow \{0, 1\}$ defined by

$$(f \diamond g)(x_1, \dots, x_m) = f(g(x_1), \dots, g(x_m)),$$

where $x_1, \dots, x_m \in \{0, 1\}^n$. It is easy to see that $L(f \diamond g) \leq L(f) \cdot L(g)$. The KRW conjecture [KRW95] asserts that this is roughly optimal, namely, that $L(f \diamond g) \gtrsim L(f) \cdot L(g)$. In this section, we prove the following related result.

Theorem 7.1. *Let $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and $g: \{0, 1\}^n \rightarrow \{0, 1\}$ be non-constant Boolean functions. Then,*

$$L(f \diamond g)^{1+O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \geq \frac{1}{O(m^4)} \cdot (L(f) - O(1)) \cdot \text{Adv}_s(g).$$

Proof. Let π be the q -hiding projection constructed in Lemma 6.14 for $q = \sqrt{1/\text{Adv}_s(g)}$, and recall that $g|_\pi$ is non-constant with probability 1. Let π^m be the m -fold join of π with itself, which is also q -hiding according to Lemma 5.3. Applying Corollary 4.6 and using the estimate $x + \sqrt{x} = O(x+1)$, we see that $s = L(f \diamond g)$ satisfies

$$\mathbb{E}[L((f \diamond g)|_{\pi^m})] = m^4 \cdot \frac{1}{\text{Adv}_s(g)} \cdot s^{1+O\left(\frac{1}{\sqrt{\log s}}\right)} + O(1).$$

On the other hand, it is not hard to see that

$$(f \diamond g)|_{\pi^m} = f \diamond g|_{\pi},$$

and since $g|_\pi$ is non-constant, f reduces to $(f \diamond g)|_{\pi^m}$. In particular,

$$L(f) \leq \mathbb{E}[L((f \diamond g)|_{\pi^m})] = m^4 \cdot \frac{1}{\text{Adv}_s(g)} \cdot s^{1+O\left(\frac{1}{\sqrt{\log s}}\right)} + O(1).$$

We obtain the theorem by rearranging. □

A direct consequence of the theorem is the following corollary, which is a special case of the KRW conjecture for inner functions g for which the soft-adversary bound is almost tight.

Corollary 7.2. *Let $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and $g: \{0, 1\}^n \rightarrow \{0, 1\}$ be non-constant Boolean functions such that $\text{Adv}_s(g) \geq L(g)^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}$. Then,*

$$L(f \diamond g) \geq \frac{1}{O(m^4)} \cdot L(f)^{1-O\left(\frac{1}{\sqrt{\log L(f)}}\right)} \cdot L(g)^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}.$$

Proof. By substituting the assumption on g in the bound of Theorem 7.1, we obtain that

$$L(f \diamond g)^{1+O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \geq \frac{1}{O(m^4)} \cdot (L(f) - O(1)) \cdot (L(g))^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}.$$

Moreover, since $L(f \diamond g) \leq L(f) \cdot L(g)$, it holds that

$$\begin{aligned} L(f \diamond g)^{1+O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} &\leq L(f \diamond g) \cdot (L(f) \cdot L(g))^{O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \\ &= L(f \diamond g) \cdot L(f)^{O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \cdot L(g)^{O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \\ &\leq L(f \diamond g) \cdot L(f)^{O\left(\frac{1}{\sqrt{\log L(f)}}\right)} \cdot L(g)^{O\left(\frac{1}{\sqrt{\log L(g)}}\right)}. \end{aligned}$$

The corollary follows by combining the two bounds. □

We have the next immediate corollary.

Corollary 7.3. *Let $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and $g: \Sigma^r \rightarrow \{0, 1\}$ be non-constant Boolean functions. Then,*

$$L(f \diamond g)^{1+O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \geq \frac{1}{O(m^4)} \cdot (L(f) - O(1)) \cdot \text{Khr}_{H_\infty}(g).$$

Moreover, if $\text{Khr}_{H_\infty}(g) \geq L(g)^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}$, then

$$L(f \diamond g) \geq \frac{1}{O(m^4)} \cdot L(f)^{1-O\left(\frac{1}{\sqrt{\log L(f)}}\right)} \cdot L(g)^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}.$$

Remark 7.4. We believe that a result along the lines of Corollary 7.3 could also be proved using the techniques of [DM18]. However, we have not verified it.

In the special case where $\Sigma = \{0, 1\}$, we can get analogues of Theorem 7.1 and Corollary 7.2 for the original Khrapchenko bound, by combining the above results with Proposition 6.9.

Corollary 7.5. *Let $f: \{0, 1\}^m \rightarrow \{0, 1\}$ and $g: \{0, 1\}^n \rightarrow \{0, 1\}$ be non-constant Boolean functions. Then,*

$$L(f \diamond g)^{1+O\left(\frac{1}{\sqrt{\log L(f \diamond g)}}\right)} \geq \frac{1}{O(m^4)} \cdot (L(f) - O(1)) \cdot \text{Khr}(g).$$

Moreover, if $\text{Khr}(g) \geq L(g)^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}$, then

$$L(f \diamond g) \geq \frac{1}{O(m^4)} \cdot L(f)^{1-O\left(\frac{1}{\sqrt{\log L(f)}}\right)} \cdot L(g)^{1-O\left(\frac{1}{\sqrt{\log L(g)}}\right)}.$$

7.2 Formula lower bounds for \mathbf{AC}^0

In this section, we derive our second main application: cubic formula lower bounds for \mathbf{AC}^0 . Formally, we have the following result.

Theorem 7.6 (Restatement of Theorem 1.1). *There exists an infinite family of Boolean functions $F_n: \{0, 1\}^n \rightarrow \{0, 1\}$, for $n \in \mathbb{N}$, that is computable by uniform formulas with unbounded fan-in of depth 4 and size $O\left(\frac{n^3}{\log^3 n}\right)$ such that $L(F_n) = \Omega(n^{3-o(1)})$.*

The function F_n is constructed similarly to the Andreev function [And87], with the Parity function replaced with the surjectivity function [BM12], defined next. Fix a parameter $s \in \mathbb{N}$, define $\Sigma = [2s + 1]$, and let $r = 3s + 1$. The surjectivity function $\text{Surj}: \Sigma^r \rightarrow \{0, 1\}$ interprets its input as a function from $[r]$ to Σ , and outputs whether the function is surjective. In other words, $\text{Surj}(\sigma_1, \dots, \sigma_r) = 1$ if and only if every symbol in Σ appears in $(\sigma_1, \dots, \sigma_r)$. Observe that the input length of the function Surj measured in bits is

$$n = r \cdot \lceil \log |\Sigma| \rceil = (3s + 1) \cdot \lceil \log(2s + 1) \rceil,$$

where the $\lceil \log |\Sigma| \rceil$ factor appears since it takes $\lceil \log |\Sigma| \rceil$ bits to encode a symbol in Σ . We now define the function Surj for any sufficiently large input length n by choosing s to be the largest number such that $n \geq (3s + 1) \cdot \lceil \log(2s + 1) \rceil$, and defining the function $\text{Surj}_n: \{0, 1\}^n \rightarrow \{0, 1\}$ as

above by interpreting its input as a string in Σ^r . It can be verified that for this choice of s it holds that

$$(3s + 1) \cdot \lceil \log(2s + 1) \rceil = \Theta(n)$$

(provided that n is sufficiently large), and in particular $s = \Theta(\frac{n}{\log n})$. In order to prove Theorem 7.6, we will use Theorem 7.1 with the inner function g being Surj_n . To this end, we first analyze $\text{Khr}_{H_\infty}(\text{Surj}_n)$.

Lemma 7.7. *For every sufficiently large input length n , it holds that $\text{Khr}_{H_\infty}(\text{Surj}_n) = \Omega(\frac{n^2}{\log^2 n})$.*

Proof. Let n be sufficiently large such that Surj_n is well-defined, and let s be the corresponding parameter. We define a Khrapchenko distribution (\mathbf{a}, \mathbf{b}) for Surj_n as follows. The input $\mathbf{b} \in \text{Surj}_n^{-1}(0)$ is a uniformly distributed string in Σ^{3s+1} in which $s + 1$ of the symbols in Σ appear exactly twice, $s - 1$ of the symbols in Σ appear exactly once, and the remaining symbol in Σ does not appear at all. The string \mathbf{a} is sampled by choosing uniformly at random a coordinate \mathbf{i} such that $\mathbf{b}_{\mathbf{i}}$ is one of the symbols that appear twice in \mathbf{b} , and replacing $\mathbf{b}_{\mathbf{i}}$ with the unique symbol in Σ that does not appear in \mathbf{b} .

We turn to bound $H_\infty(\mathbf{i} \mid \mathbf{b})$ and $H_\infty(\mathbf{i} \mid \mathbf{a})$. First, observe that conditioned on any choice of \mathbf{b} , the coordinate \mathbf{i} is uniformly distributed among $2s + 2$ coordinates, and therefore

$$H_\infty(\mathbf{i} \mid \mathbf{b}) = \log(2s + 2).$$

In order to bound $H_\infty(\mathbf{i} \mid \mathbf{a})$, observe that \mathbf{a} is distributed like a uniformly distributed string in Σ^{3s+1} in which $s + 1$ of the symbols in Σ appear exactly once, and the remaining s symbols in Σ appear exactly twice. Conditioned on any choice of \mathbf{a} , the coordinate \mathbf{i} is uniformly distributed over the $s + 1$ coordinates of symbols that appear exactly once. It follows that

$$H_\infty(\mathbf{i} \mid \mathbf{a}) \geq \log(s + 1).$$

We conclude that

$$\text{Khr}_{H_\infty}(\text{Surj}_n) \geq 2^{H_\infty(\mathbf{i} \mid \mathbf{a}) + H_\infty(\mathbf{i} \mid \mathbf{b})} \geq (s + 1) \cdot (2s + 2) = \Omega\left(\frac{n^2}{\log^2 n}\right). \quad \square$$

Lemma 7.7 implies, via Proposition 6.5, that $L(\text{Surj}_n) = \tilde{\Omega}(n^2)$. The same result also follows from the lower bound $Q(\text{Surj}_n) = \Omega(n/\log n)$ proved in [BM12], where $Q(f)$ is the quantum query complexity of f , via the relation $Q(f) = O(\sqrt{L(f)})$ [ACR⁺10, Rei11].

We turn to prove Theorem 7.6 using our special case of the KRW conjecture (Corollary 7.2).

Proof of Theorem 7.6. We would like to construct, for every sufficiently large $n \in \mathbb{N}$, a function $F: \{0, 1\}^n \rightarrow \{0, 1\}$ that is computable by a formula with unbounded fan-in of depth 4 and size $O(\frac{n^3}{\log^3 n})$ such that $L(F) = \Omega(n^{3-o(1)})$. We start by constructing the function F for input lengths n of a special form, and then extend the construction to all sufficiently large input lengths.

First, assume that the input length n is of the form $2^k \cdot (k + 1)$, where $k \in \mathbb{N}$ is sufficiently large such that Surj_{2^k} is well-defined. The function $F: \{0, 1\}^n \rightarrow \{0, 1\}$ takes two inputs: the truth table of a function $f: \{0, 1\}^k \rightarrow \{0, 1\}$, and k strings $x_1, \dots, x_k \in \{0, 1\}^{2^k}$. On such inputs, the function F outputs

$$F(f, x_1, \dots, x_k) = (f \diamond \text{Surj}_{2^k})(x_1, \dots, x_k) = f(\text{Surj}_{2^k}(x_1), \dots, \text{Surj}_{2^k}(x_k)).$$

It is easy to see that the function F has input length n . We show that F can be computed by a formula with unbounded fan-in of depth 4 and size $O\left(\frac{n^3}{\log^3 n}\right)$. We start by constructing a formula for Surj_{2^k} . Recall that Surj_{2^k} takes as input (the binary encoding of) a string $(\sigma_1, \dots, \sigma_r) \in \Sigma^r$, where $|\Sigma|, r = O\left(\frac{2^k}{k}\right)$. For simplicity, let us assume that every symbol in Σ is encoded by exactly one binary string, and that if the binary input to Surj_{2^k} contains a binary string in $\{0, 1\}^{\lceil \log |\Sigma| \rceil}$ that does not encode any symbol in Σ , then we do not care what the formula outputs. Now, observe that

$$\text{Surj}_n(\sigma_1, \dots, \sigma_r) = \bigwedge_{\gamma \in \Sigma} \bigvee_{j=1}^r (\sigma_j = \gamma). \quad (9)$$

It is not hard to see that the expression on the right-hand side can be implemented by a formula with unbounded fan-in of depth 3 and size $|\Sigma| \cdot r \cdot \lceil \log |\Sigma| \rceil = O\left(\frac{2^{2k}}{k}\right)$. Next, observe that

$$F(f, x_1, \dots, x_k) = \bigvee_{y \in \{0, 1\}^k} \left[(f(y) = 1) \wedge \left(\bigwedge_{i=1}^k \text{Surj}_{2^k}(x_i) = y_i \right) \right].$$

Using the foregoing formula for surjectivity, it is not hard to see that the expression on the right-hand side can be implemented by a formula with unbounded fan-in of depth 4 and size

$$O\left(2^k \cdot k \cdot \frac{2^{2k}}{k}\right) = O\left(2^{3k}\right) = O\left(\frac{n^3}{\log^3 n}\right),$$

as required.

Finally, we prove that $L(F) = \Omega(n^{3-o(1)})$. To this end, let us hardwire the input f to be some function from $\{0, 1\}^k$ to $\{0, 1\}$ such that $L(f) = \Omega\left(\frac{2^k}{\log k}\right)$ (such a function exists by a well-known counting argument, see [Juk12, Theorem 1.23]). After hardwiring the input f , the function F becomes exactly the function $f \diamond \text{Surj}_{2^k}$. Now, by Theorem 7.1 it follows that

$$\begin{aligned} L(F)^{1+O\left(\frac{1}{\sqrt{\log L(F)}}\right)} &\geq L(f \diamond \text{Surj}_{2^k})^{1+O\left(\frac{1}{\sqrt{\log L(f \diamond \text{Surj}_{2^k})}}\right)} \\ &\geq \frac{1}{O(k^4)} \cdot (L(f) - O(1)) \cdot \text{Adv}_s(\text{Surj}_{2^k}) \\ &\geq \frac{1}{O(k^4)} \cdot (L(f) - O(1)) \cdot \text{Khr}_{H_\infty}(\text{Surj}_{2^k}) \\ &= \Omega\left(\frac{2^{3k}}{k^6 \log k}\right) \\ &= \Omega\left(\frac{n^3}{\log^{9+o(1)} n}\right). \end{aligned}$$

Since $L(F) = O(n^3)$, we obtain that

$$L(F) \geq \Omega\left(\frac{n^3}{\log^{9+o(1)} n}\right) \Big/ L(F)^{O\left(\frac{1}{\sqrt{\log L(F)}}\right)} \geq \Omega\left(\frac{n^3}{\log^{9+o(1)} n}\right) \Big/ n^{O\left(\frac{1}{\sqrt{\log n}}\right)} = \Omega(n^{3-o(1)}),$$

as required.

It remains to deal with input lengths n that are not of the form $2^k \cdot (k + 1)$. For such input lengths n , we choose k to be the largest natural number such that $2^k \cdot (k + 1) \leq n$, and proceed as before. It can be verified that for this choice of k it holds that $2^k \cdot (k + 1) = \Theta(n)$, and therefore all the foregoing asymptotic bounds continue to hold. \square

Remark 7.8. We note that our methods cannot prove formula lower bounds that are better than cubic. As explained in Remark 6.15, it holds that $\text{Adv}_s(f)$ is upper bounded by n^2 . In particular, this implies that $\text{Chr}_{H_\infty}(f)$ is upper bounded by n^2 . Thus, one cannot expect to obtain a lower bound that is better than cubic by combining these measures with Andreev’s argument.

7.3 Reproving known formula lower bounds

The proof of Theorem 7.6 combines a lower bound on $\text{Adv}_s(\text{Surj}_n)$ with an upper bound on $L(\text{Surj}_n)$. More generally, we can prove the following result along similar lines.

Theorem 7.9. *Let $g: \{0, 1\}^n \rightarrow \{0, 1\}$ be an arbitrary function, and let k be an integer satisfying $k = \log n + O(1)$. Let $F: \{0, 1\}^{2^k + nk} \rightarrow \{0, 1\}$ be a function on $N = \Theta(n \log n)$ variables, whose input consists of a function $f: \{0, 1\}^k \rightarrow \{0, 1\}$ and k strings $x_1, \dots, x_k \in \{0, 1\}^n$, given by*

$$F(f, x_1, \dots, x_k) = (f \diamond g)(x_1, \dots, x_k).$$

We call F the g -based Andreev function.

1. If $\text{Adv}_s(g) = \Omega(n^{2-o(1)})$ then $L(F) = \Omega(N^{3-o(1)})$.
2. If $L(g) = O(n^{2+o(1)})$ then $L(F) = O(N^{3+o(1)})$.

The proof of Theorem 7.9 is very similar to the proof of Theorem 7.6, and so we leave it to the reader (the main difference is that we use Corollary 7.5 instead of Theorem 7.1). Using Theorem 7.9, we can derive two known special cases: the original Andreev function (in which g is Parity), and the variant considered in [GTN19]. In particular, using the known facts that $\text{Chr}(\text{Parity}) \geq n^2$ and $\text{Chr}(\text{Majority}) \geq \Omega(n^2)$ [Khr72], we obtain the following results.

Corollary 7.10. *The formula complexity of the Parity_n -based Andreev function is $\Theta(N^{3 \pm o(1)})$.*

Corollary 7.11. *For $m \geq 3$ odd, let $\text{Majority}_n: \{0, 1\}^m \rightarrow \{0, 1\}$ be the Majority function. The formula complexity of the Majority_n -based Andreev function is $\Omega(N^{3-o(1)})$.*

The best upper bound on the formula complexity of Majority_n is only $O(n^{3.91})$ [Ser16], and so we do not expect Corollary 7.11 to be tight.

Acknowledgements

A.T. would like to thank Igor Carboni Oliveira for bringing the question of proving formula size lower bounds for \mathbf{AC}^0 to his attention. We are also grateful to Robin Kothari for posing this open question on “Theoretical Computer Science Stack Exchange” [Kot11], and to Kaveh Ghasemloo and Stasys Jukna for their feedback on this question. We would like to thank Anna Gál for very helpful discussions, and Gregory Rosenthal for helpful comments on the manuscript.

References

- [ACR⁺10] A. Ambainis, A. M. Childs, B. W. Reichardt, R. Špalek, and S. Zhang. Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM J. Comput.*, 39(6):2513–2530, 2010. doi:10.1137/080712167.
- [Ajt83] Miklós Ajtai. Σ_1^1 -formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983.
- [Amb02] Andris Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. System Sci.*, 64(4):750–767, 2002. Special issue on STOC 2000 (Portland, OR). doi:10.1006/jcss.2002.1826.
- [Amb06] Andris Ambainis. Polynomial degree vs. quantum query complexity. *J. Comput. System Sci.*, 72(2):220–238, 2006. doi:10.1016/j.jcss.2005.06.006.
- [And87] Alexander E. Andreev. On a method for obtaining more than quadratic effective lower bounds for the complexity of π -schemes. *Moscow University Mathematics Bulletin*, 42(1):24–29, 1987.
- [BB94] Maria Luisa Bonet and Samuel R. Buss. Size-depth tradeoffs for Boolean formulae. *Inf. Process. Lett.*, 49(3):151–155, 1994.
- [BBCR13] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. *SIAM J. Comput.*, 42(3):1327–1363, 2013.
- [BdW02] Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: a survey. *Theor. Comput. Sci.*, 288(1):21–43, 2002.
- [BM12] Paul Beame and Widad Machmouchi. The quantum query complexity of AC0. *Quantum Info. Comput.*, 12(7–8):670–676, July 2012.
- [Bre74] Richard P. Brent. The parallel evaluation of general arithmetic expressions. *J. ACM*, 21(2):201–206, 1974.
- [CKK12] Andrew M. Childs, Shelby Kimmel, and Robin Kothari. The quantum query complexity of read-many formulas. In Leah Epstein and Paolo Ferragina, editors, *Algorithms – ESA 2012*, pages 337–348, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [CT91] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, 1991.
- [DM18] Irit Dinur and Or Meir. Toward the KRW composition conjecture: Cubic formula lower bounds via communication complexity. *Computational Complexity*, 27(3):375–462, 2018.
- [dRMN⁺20] Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, and Robert Robere. KRW composition theorems via lifting. *Electron. Colloquium Comput. Complex.*, 27:99, 2020. URL: <https://eccc.weizmann.ac.il/report/2020/099>.

- [EIRS01] Jeff Edmonds, Russell Impagliazzo, Steven Rudich, and Jiri Sgall. Communication complexity towards lower bounds on circuit depth. *Computational Complexity*, 10(3):210–246, 2001.
- [FSS84] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
- [GH92] Mikael Goldmann and Johan Håstad. A simple lower bound for monotone clique using a communication game. *Inf. Process. Lett.*, 41(4):221–226, 1992.
- [GKR12] Anat Ganor, Ilan Komargodski, and Ran Raz. The spectrum of small demorgan formulas. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:174, 2012.
- [GMWW17] Dmitry Gavinsky, Or Meir, Omri Weinstein, and Avi Wigderson. Toward better formula lower bounds: The composition of a function and a universal relation. *SIAM J. Comput.*, 46(1):114–131, 2017.
- [GP18] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM J. Comput.*, 47(5):1778–1806, 2018.
- [GTN19] Anna Gál, Avishay Tal, and Adrian Trejo Nuñez. Cubic formula size lower bounds based on compositions with majority. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pages 35:1–35:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ITCS.2019.35.
- [Hås86] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *STOC*, pages 6–20, 1986.
- [Hås87] Johan Torkel Håstad. *Computational limitations for small-depth circuits*. MIT press, 1987.
- [Hås93] Johan Håstad. The shrinkage exponent is 2. In *Proceedings of the 1993 IEEE 34th Annual Foundations of Computer Science, SFCS '93*, pages 114–123, USA, 1993. IEEE Computer Society. doi:10.1109/SFCS.1993.366876.
- [Hås98] Johan Håstad. The shrinkage exponent of de Morgan formulas is 2. *SIAM J. Comput.*, 27(1):48–64, 1998.
- [HJP95] Johan Håstad, Stasys Jukna, and Pavel Pudlák. Top-down lower bounds for depth-three circuits. *Computational Complexity*, 5(2):99–112, 1995.
- [HRST17] Johan Håstad, Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. An average-case depth hierarchy theorem for boolean circuits. *J. ACM*, 64(5):35:1–35:27, 2017. doi:10.1145/3095799.
- [HW93] Johan Håstad and Avi Wigderson. Composition of the universal relation. In *Advances in computational complexity theory, AMS-DIMACS*, 1993.

- [HyLŠ07] Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 526–535. ACM, New York, 2007. doi:10.1145/1250790.1250867.
- [IN93] Russell Impagliazzo and Noam Nisan. The effect of random restrictions on formula size. *Random Struct. Algorithms*, 4(2):121–134, 1993.
- [Juk12] Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012.
- [Khr72] V. M. Khrapchenko. A method of obtaining lower bounds for the complexity of π -schemes. *Mathematical Notes Academy of Sciences USSR*, 10:474–479, 1972.
- [KM18] Sajin Koroth and Or Meir. Improved composition theorems for functions and relations. In *RANDOM*, 2018.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.
- [Kot11] Robin Kothari. Formula size lower bounds for AC0 functions, 2011. Question on Theoretical Computer Science Stack Exchange. URL: <https://cstheory.stackexchange.com/questions/7156/formula-size-lower-bounds-for-ac0-functions>.
- [Kou93] Elias Koutsoupias. Improvements on khrapchenko’s theorem. *Theor. Comput. Sci.*, 116(2):399–403, 1993.
- [KRW95] Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.
- [KW90] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.*, 3(2):255–265, 1990.
- [LLS06] Sophie Laplante, Troy Lee, and Mario Szegedy. The quantum adversary method and classical formula size lower bounds. *Comput. Complexity*, 15(2):163–196, 2006. doi:10.1007/s00037-006-0212-7.
- [LMR⁺11] Troy Lee, Rajat Mittal, Ben W. Reichardt, Robert Špalek, and Mario Szegedy. Quantum query complexity of state conversion. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science—FOCS 2011*, pages 344–353. IEEE Computer Soc., Los Alamitos, CA, 2011. doi:10.1109/FOCS.2011.75.
- [LS20] Lily Li and Morgan Shirley. The general adversary bound: A survey, January 2020. Final project. URL: http://www.henryyuen.net/fall2019/projects/general_adversary.pdf.
- [Mei20] Or Meir. Toward better depth lower bounds: Two results on the multiplexor relation. *Comput. Complex.*, 29(1):4, 2020.

- [MS20] Ivan Mihajlin and Alexander Smal. Toward better depth lower bounds: the XOR-KRW conjecture. *Electron. Colloquium Comput. Complex.*, 27:116, 2020. URL: <https://eccc.weizmann.ac.il/report/2020/116>.
- [Nec66] E. I. Neciporuk. On a Boolean function. *Soviet Mathematics Doklady*, 7(4):999–1000, 1966.
- [PR17] Toniann Pitassi and Robert Robere. Strongly exponential lower bounds for monotone computation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1246–1255, 2017.
- [PZ93] Mike Paterson and Uri Zwick. Shrinkage of de Morgan formulae under restriction. *Random Struct. Algorithms*, 4(2):135–150, 1993.
- [Raz90] Alexander A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.
- [Rei09] Ben W. Reichardt. Span programs and quantum query complexity: the general adversary bound is nearly tight for every Boolean function. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2009*, pages 544–551. IEEE Computer Soc., Los Alamitos, CA, 2009. doi:10.1109/FOCS.2009.55.
- [Rei11] Ben Reichardt. Reflections for quantum query algorithms. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 560–569. SIAM, 2011. doi:10.1137/1.9781611973082.44.
- [Rei14] Ben W. Reichardt. Span programs are equivalent to quantum query algorithms. *SIAM J. Comput.*, 43(3):1206–1219, 2014. doi:10.1137/100792640.
- [RM99] Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, 1999.
- [RST15] Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. An average-case depth hierarchy theorem for Boolean circuits. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1030–1048, 2015.
- [RW92] Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992.
- [Ser16] I. S. Sergeev. Complexity and depth of formulas for symmetric Boolean functions. *Moscow University Mathematics Bulletin*, 71(3):127–130, 2016.
- [Spi71] Philip M. Spira. On time-hardware complexity tradeoffs for Boolean functions. In *Proceedings of the Fourth Hawaii International Symposium on System Sciences*, pages 525–527, 1971.
- [ŠS06] Robert Špalek and Mario Szegedy. All quantum adversary methods are equivalent. *Theory Comput.*, 2:1–18, 2006. doi:10.4086/toc.2006.v002a001.

- [Sub61] Bella Abramovna Subbotovskaya. Realizations of linear functions by formulas using +,.,-. *Soviet Mathematics Doklady*, 2:110–112, 1961.
- [Tal14] Avishay Tal. Shrinkage of de Morgan formulae by spectral techniques. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 551–560, 2014.
- [Yao85] Andrew Chi-Chih Yao. Separating the polynomial-time hierarchy by oracles (preliminary version). In *FOCS*, pages 1–10. IEEE Computer Society, 1985.
- [Zwi91] Uri Zwick. An extension of khrapchenko’s theorem. *Inf. Process. Lett.*, 37(4):215–217, 1991.

A Proof of Claim 4.8

In this appendix we prove Claim 4.8, restated next.

Claim 4.8. *Let π be a (q_0, q_1) -hiding projection, and let \mathcal{E} be a random set of projections that is independent of π . Then, for every $\sigma \in \{0, 1\}$, it holds that*

$$\Pr[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \pi_{y_j \leftarrow \sigma} \in \mathcal{E}] \leq q_\sigma.$$

Proof. Let π and \mathcal{E} be as in the claim, and let $\sigma \in \{0, 1\}$. It holds that

$$\begin{aligned} & \Pr[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \pi_{y_j \leftarrow \sigma} \in \mathcal{E}] \\ &= \sum_{\pi} \Pr[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \pi_{y_j \leftarrow \sigma} = \pi \text{ and } \pi \in \mathcal{E}] \cdot \Pr[\pi_{y_j \leftarrow \sigma} = \pi \mid \pi_{y_j \leftarrow \sigma} \in \mathcal{E}] \\ &= \sum_{\pi} \Pr[\pi(x_i) \in \{y_j, \bar{y}_j\} \mid \pi_{y_j \leftarrow \sigma} = \pi] \cdot \Pr[\pi_{y_j \leftarrow \sigma} = \pi \mid \pi_{y_j \leftarrow \sigma} \in \mathcal{E}] \\ & \hspace{20em} (\pi \text{ and } \mathcal{E} \text{ are independent}) \\ &\leq \sum_{\pi} q_\sigma \cdot \Pr[\pi_{y_j \leftarrow \sigma} = \pi \mid \pi_{y_j \leftarrow \sigma} \in \mathcal{E}] \hspace{2em} (\pi \text{ is } (q_0, q_1)\text{-hiding}) \\ &= q_\sigma. \hspace{10em} \square \end{aligned}$$

B Proof of Propositions 6.4 and 6.9

In this appendix we prove Propositions 6.4 and 6.9, restated next.

Proposition 6.4. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $\frac{\text{Khr}(f)}{4} \leq \text{Khr}_{H_\infty}(f) \leq \text{Khr}(f)$.*

Proposition 6.9. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $\frac{\text{Khr}(f)}{4} \leq \text{Adv}_u(f)$.*

We relate $\text{Khr}(f)$ to $\text{Khr}_{H_\infty}(f)$ using an auxiliary measure $\text{Khr}_{\min}(f)$:

$$\text{Khr}_{\min}(f) = \max_{A \subseteq f^{-1}(1), B \subseteq f^{-1}(0)} \left(\min_{a \in A} |E(a, B)| \cdot \min_{b \in B} |E(A, b)| \right),$$

where $E(a, B) = E(\{a\}, B)$ and similarly $E(A, b) = E(A, \{b\})$.

We first show that $\text{Khr}_{\min}(f)$ and $\text{Khr}(f)$ are equal up to constants.

Claim B.1. For any $f: \Sigma^n \rightarrow \{0, 1\}$ it holds that $\frac{\text{Khr}(f)}{4} \leq \text{Khr}_{\min}(f) \leq \text{Khr}(f)$.

Proof. The inequality $\text{Khr}_{\min}(f) \leq \text{Khr}(f)$ is simple since $\min_{a \in A} |E(a, B)| \leq |E(A, B)|/|A|$ and similarly $\min_{b \in B} |E(A, b)| \leq |E(A, B)|/|B|$.

The other direction is more subtle. For ease of notation, for any sets $A \subseteq f^{-1}(1)$ and $B \subseteq f^{-1}(0)$ we denote $\text{Khr}(A, B) = \frac{|E(A, B)|^2}{|A||B|}$. Thus,

$$\text{Khr}(f) = \max_{A \subseteq f^{-1}(1), B \subseteq f^{-1}(0)} \text{Khr}(A, B). \quad (10)$$

Assume that A and B are sets that maximize $\text{Khr}(A, B)$ in Equation (10). We show that

$$\forall a \in A: |E(a, B)| \geq \frac{|E(A, B)|}{2|A|}, \quad (11)$$

$$\forall b \in B: |E(A, b)| \geq \frac{|E(A, B)|}{2|B|}. \quad (12)$$

In words, the min-degree is at least half the average degree. Before showing why Equations (11) and (12) hold, we show that they imply the statement of the claim. Indeed,

$$\text{Khr}_{\min}(f) \geq \left(\min_{a \in A} |E(a, B)| \cdot \min_{b \in B} |E(A, b)| \right) \geq \frac{|E(A, B)|}{2|A|} \cdot \frac{|E(A, B)|}{2|B|} = \frac{\text{Khr}(A, B)}{4} = \frac{\text{Khr}(f)}{4}.$$

It remains to show that Equations (11) and (12) hold. We focus on Equation (11) due to symmetry. Assume by contradiction that there exists an $a \in A$ for which

$$|E(a, B)| < \frac{|E(A, B)|}{2|A|}.$$

It must be the case that $|A| > 1$, as otherwise A contains only one element and $|E(a, B)| = |E(A, B)|/|A|$ for this element. Consider now the set $A' = A \setminus \{a\}$ — which is non-empty by the above discussion. We claim that $\text{Khr}(A', B) > \text{Khr}(A, B)$, contradicting the choice of A, B . Indeed,

$$\begin{aligned} \text{Khr}(A', B) &= \frac{|E(A', B)|^2}{|A'||B|} \\ &= \frac{(|E(A, B)| - |E(a, B)|)^2}{(|A| - 1)|B|} \\ &> \frac{(|E(A, B)| - \frac{|E(A, B)|}{2|A|})^2}{(|A| - 1)|B|} && \text{(By assumption on } a) \\ &= \frac{|E(A, B)|^2 \cdot (1 - \frac{1}{2|A|})^2}{|A||B| \cdot (1 - \frac{1}{|A|})} \\ &= \text{Khr}(A, B) \cdot \frac{(1 - \frac{1}{2|A|})^2}{(1 - \frac{1}{|A|})} \\ &> \text{Khr}(A, B). \quad \square \end{aligned}$$

It turns out that over the binary alphabet, $\text{Khr}_{\min}(f)$ and $\text{Khr}_{H_\infty}(f)$ coincide.

Claim B.2. For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $\text{Khr}_{H_\infty}(f) = \text{Khr}_{\min}(f)$.

Proof. We first show that $\text{Khr}_{H_\infty}(f) \geq \text{Khr}_{\min}(f)$. Let A, B be sets that maximize the expression $(\min_{a \in A} |E(a, B)| \cdot \min_{b \in B} |E(A, b)|)$. We take (\mathbf{a}, \mathbf{b}) to be a uniformly distributed over $E(A, B)$. It is not hard to see that

$$2^{H_\infty(i|\mathbf{a})} = 2^{H_\infty(\mathbf{b}|\mathbf{a})} = \min_{a \in A} |E(a, B)|,$$

and similarly $2^{H_\infty(i|\mathbf{b})} = \min_{a \in A} |E(a, B)|$. We thus get $\text{Khr}_{H_\infty}(f) \geq \text{Khr}_{\min}(f)$.

Next, we show the other direction, $\text{Khr}_{\min}(f) \geq \text{Khr}_{H_\infty}(f)$. Let (\mathbf{a}, \mathbf{b}) be a random variable distributed according to a Khrapchenko distribution for f that attains the maximum of $2^{H_\infty(\mathbf{a}|\mathbf{b})+H_\infty(\mathbf{b}|\mathbf{a})}$ over all such distributions. Let $A := \text{supp}(\mathbf{a})$ and $B := \text{supp}(\mathbf{b})$ be the supports of \mathbf{a} and \mathbf{b} , respectively. By definition of H_∞ we have $2^{H_\infty(\mathbf{a}|\mathbf{b})} = \frac{1}{\max_{a,b} \Pr[\mathbf{a}=a|\mathbf{b}=b]}$. Rearranging, we get that for any b in the support of \mathbf{b} it holds that

$$\Pr[\mathbf{a} = a | \mathbf{b} = b] \leq 1/2^{H_\infty(\mathbf{a}|\mathbf{b})}.$$

In particular, since these probabilities sum to 1, it must be the case that there are at least $2^{H_\infty(\mathbf{a}|\mathbf{b})}$ neighbors of b in A — i.e., $|E(A, b)| \geq 2^{H_\infty(\mathbf{a}|\mathbf{b})}$ for all $b \in B$. Similarly, $|E(a, B)| \geq 2^{H_\infty(\mathbf{b}|\mathbf{a})}$ for all $a \in A$. The two sets A and B show that $\text{Khr}_{\min}(f) \geq 2^{H_\infty(\mathbf{a}|\mathbf{b})} \cdot 2^{H_\infty(\mathbf{b}|\mathbf{a})} = \text{Khr}_{H_\infty}(f)$. \square

Proposition 6.4 follows by combining these two claims. Proposition 6.9 follows from Claim B.2 using the following simple observation.

Claim B.3. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $\text{Adv}_u(f) \geq \text{Khr}_{\min}(f)$.*

Proof. Let A, B be sets that attain $\text{Khr}_{\min}(f)$. Define R to be the set of pairs $(a, b) \in A \times B$ that differ at a single coordinate. Thus $|R(a, B)| = |E(a, B)|$ and $|R(A, b)| = |E(A, b)|$. Moreover, $|R_i(a, B)| \leq 1$ since there is a unique b that differs from a only on the i -th coordinate. Similarly, $|R_i(A, b)| \leq 1$. The claim now immediately follows by comparing the definitions of $\text{Khr}_{\min}(f)$ and $\text{Adv}_u(f)$. \square

C Proof of Proposition 6.5

In this appendix, we prove Proposition 6.5, restated next.

Proposition 6.5. *Let Σ be a finite alphabet, and let $f: \Sigma^n \rightarrow \{0, 1\}$ be a Boolean function. Then $L(f) \geq \text{Khr}_{H_\infty}(f)$.*

Preliminaries

While the proof of Proposition 6.5 is fairly simple, it uses some basic concepts from information theory which we review next.

Definition C.1. Let $\mathbf{x}, \mathbf{y}, \mathbf{z}$ be discrete random variables.

- The *entropy* of \mathbf{x} is $H(\mathbf{x}) = \mathbb{E}_{x \leftarrow \mathbf{x}} \left[\log \frac{1}{\Pr[x=x]} \right]$.
- The *conditional entropy* of \mathbf{x} given \mathbf{y} is $H(\mathbf{x} | \mathbf{y}) = \mathbb{E}_{y \leftarrow \mathbf{y}} [H(\mathbf{x} | \mathbf{y} = y)]$.
- The *mutual information* between \mathbf{x} and \mathbf{y} is $I(\mathbf{x}; \mathbf{y}) = H(\mathbf{x}) - H(\mathbf{x} | \mathbf{y})$.

- The *conditional mutual information* between \mathbf{x} and \mathbf{y} given \mathbf{z} is

$$I(\mathbf{x}; \mathbf{y} \mid \mathbf{z}) = H(\mathbf{x} \mid \mathbf{z}) - H(\mathbf{x} \mid \mathbf{y}, \mathbf{z}).$$

We use the following basic facts from information theory (see [CT91] for proofs).

Fact C.2. *Let $\mathbf{x}, \mathbf{y}, \mathbf{z}$ be discrete random variables with finite supports, and let X denote the support of \mathbf{x} .*

- *It holds that $0 \leq H(\mathbf{x}) \leq \log |X|$.*
- *It holds that $0 \leq H_\infty(\mathbf{x} \mid \mathbf{y}) \leq H(\mathbf{x} \mid \mathbf{y}) \leq H(\mathbf{x})$, where $H(\mathbf{x} \mid \mathbf{y}) = 0$ if and only if \mathbf{y} determines \mathbf{x} (i.e., \mathbf{x} is a function of \mathbf{y}).*
- *If \mathbf{y} determines \mathbf{x} conditioned on \mathbf{z} (i.e., \mathbf{x} is a function of \mathbf{y} and \mathbf{z}) then $H(\mathbf{y} \mid \mathbf{z}) \geq H(\mathbf{x} \mid \mathbf{z})$.*
- *It holds that $0 \leq I(\mathbf{x}; \mathbf{y}) \leq H(\mathbf{x})$. Similarly, it holds that $0 \leq I(\mathbf{x}; \mathbf{y} \mid \mathbf{z}) \leq H(\mathbf{x} \mid \mathbf{z})$, where $I(\mathbf{x}; \mathbf{y} \mid \mathbf{z}) = 0$ if and only if \mathbf{x} and \mathbf{y} are independent conditioned on any value of \mathbf{z} .*
- *It holds that $I(\mathbf{x}; \mathbf{y}) = I(\mathbf{y}; \mathbf{x})$. Similarly, $I(\mathbf{x}; \mathbf{y} \mid \mathbf{z}) = I(\mathbf{y}; \mathbf{x} \mid \mathbf{z})$.*
- The chain rule: *It holds that*

$$I(\mathbf{x}; \mathbf{y}, \mathbf{z} \mid \mathbf{w}) = I(\mathbf{x}; \mathbf{z} \mid \mathbf{w}) + I(\mathbf{x}; \mathbf{y} \mid \mathbf{z}, \mathbf{w}).$$

Finally, we use the following result from the theory of interactive information complexity.

Claim C.3 ([BBCR13, Fact 4.15]). *Let Π be a deterministic protocol. Suppose we invoke Π on random inputs \mathbf{x} and \mathbf{y} for Alice and Bob, respectively, and let ℓ denote the random leaf that Π reaches on those inputs. Then,*

$$I(\mathbf{x}, \mathbf{y}; \ell) \geq I(\mathbf{x}; \ell \mid \mathbf{y}) + I(\mathbf{y}; \ell \mid \mathbf{x}).$$

Proof sketch. Let \mathbf{t} denote the transcript of the protocol that is associated with ℓ . We prove that $I(\mathbf{x}, \mathbf{y}; \mathbf{t}) \geq I(\mathbf{x}; \mathbf{t} \mid \mathbf{y}) + I(\mathbf{y}; \mathbf{t} \mid \mathbf{x})$, as this is equivalent to the claim. Suppose Alice speaks first, and denote the (random) bit she sends by \mathbf{t}_1 . We show that $I(\mathbf{x}, \mathbf{y}; \mathbf{t}_1) \geq I(\mathbf{x}; \mathbf{t}_1 \mid \mathbf{y}) + I(\mathbf{y}; \mathbf{t}_1 \mid \mathbf{x})$. Using the chain rule, the external information of Π_1 can be written as

$$I(\mathbf{x}, \mathbf{y}; \mathbf{t}_1) = I(\mathbf{y}; \mathbf{t}_1) + I(\mathbf{x}; \mathbf{t}_1 \mid \mathbf{y}) \geq I(\mathbf{x}; \mathbf{t}_1 \mid \mathbf{y}) = I(\mathbf{x}; \mathbf{t}_1 \mid \mathbf{y}) + I(\mathbf{y}; \mathbf{t}_1 \mid \mathbf{x}),$$

where the last equality follows since $I(\mathbf{y}; \mathbf{t}_1 \mid \mathbf{x}) = 0$, as Alice's message \mathbf{t}_1 is independent of \mathbf{y} given her input \mathbf{x} . Proceeding by induction on the coordinates of \mathbf{t} using the chain rule finishes the proof. \square

The proof of Proposition 6.5

Our proof of Proposition 6.5 generalizes similar arguments in [KW90, GMWW17]. The proof employs the KW relation KW_f , which we extend to non-binary alphabets as follows: Alice and Bob get strings in $\{0, 1\}^{n \cdot \lceil \log |\Sigma| \rceil}$, and would like to find a coordinate in $n \cdot \lceil \log |\Sigma| \rceil$ on which the strings differ.

Let Π be a protocol that solves KW_f . We prove that $L(\Pi) \geq \text{Khr}_{H_\infty}(f)$. Let (\mathbf{a}, \mathbf{b}) be a Khrapchenko distribution for f that attains $\text{Khr}_{H_\infty}(f)$, and let \mathbf{j} be the unique coordinate in $[n]$ on which \mathbf{a} and \mathbf{b} differ. Let ℓ be the leaf that Π reaches on input (\mathbf{a}, \mathbf{b}) . By Fact C.2, it holds that

$$I(\ell; \mathbf{a}, \mathbf{b}) \leq H(\ell) \leq \log |L(\Pi)|.$$

On the other hand, Claim C.3 implies that

$$I(\ell; \mathbf{a}, \mathbf{b}) \geq I(\ell; \mathbf{a} \mid \mathbf{b}) + I(\ell; \mathbf{b} \mid \mathbf{a}).$$

Next, observe that \mathbf{a} and \mathbf{b} together determine ℓ , and also that ℓ determines \mathbf{i} . This implies that

$$\begin{aligned} I(\ell; \mathbf{a} \mid \mathbf{b}) &= H(\ell \mid \mathbf{b}) - H(\ell \mid \mathbf{a}, \mathbf{b}) && \text{(by definition)} \\ &= H(\ell \mid \mathbf{b}) && (\mathbf{a} \text{ and } \mathbf{b} \text{ determine } \ell) \\ &\geq H(\mathbf{i} \mid \mathbf{b}) && (\ell \text{ determines } \mathbf{i}). \end{aligned}$$

Similarly, it can be shown that $I(\ell; \mathbf{b} \mid \mathbf{a}) \geq H(\mathbf{i} \mid \mathbf{a})$. By combining the foregoing equations, it follows that

$$\log |L(\Pi)| \geq I(\ell; \mathbf{a} \mid \mathbf{b}) + I(\ell; \mathbf{b} \mid \mathbf{a}) \geq H(\mathbf{i} \mid \mathbf{b}) + H(\mathbf{i} \mid \mathbf{a}),$$

and thus

$$|L(\Pi)| \geq 2^{H(\mathbf{i} \mid \mathbf{a}) + H(\mathbf{i} \mid \mathbf{b})} \geq 2^{H_\infty(\mathbf{i} \mid \mathbf{a}) + H_\infty(\mathbf{i} \mid \mathbf{b})} = \text{Khr}_{H_\infty}(f),$$

as required.

D Proof of Proposition 6.13

In this appendix, we prove Proposition 6.13, restated next.

Proposition 6.13. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $L(f) \geq \text{Adv}_s(f)$.*

We prove Proposition 6.13 by relating $\text{Adv}_s(f)$ to Ambainis' weighted quantum adversary bound [Amb06], which we now define.

Definition D.1. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. Let $R \subseteq f^{-1}(1) \times f^{-1}(0)$, and let A, B be the projections into the first and second coordinates, respectively. A *weighting scheme* consists of an assignment of weights $w(a, b) > 0$ for all $(a, b) \in R$ and $w'(a, b, i), w'(b, a, i) > 0$ for all $(a, b) \in R$ and $i \in [n]$ such that $a_i \neq b_i$. A weighting scheme is *valid* if for all $(a, b) \in R$ and $i \in [n]$ such that $a_i \neq b_i$, it holds that

$$w'(a, b, i)w'(b, a, i) \geq w(a, b)^2.$$

The *weighted quantum adversary bound* of f , denoted $WA^2(f)$, is the maximal value of

$$\min_{a \in A, i \in [n]} \frac{\sum_{b: (a,b) \in R} w(a, b)}{\sum_{b: (a,b) \in R, a_i \neq b_i} w'(a, b, i)} \cdot \min_{b \in B, i \in [n]} \frac{\sum_{a: (a,b) \in R} w(a, b)}{\sum_{a: (a,b) \in R, a_i \neq b_i} w'(b, a, i)}$$

over all $R \subseteq f^{-1}(1) \times f^{-1}(0)$ and all valid weighting schemes w, w' .

Špalek and Szegedy showed that $WA(f) = \sqrt{WA^2(f)}$ is equivalent to several other strengthenings of the unweighted quantum adversary bound, which are collectively denoted by $\text{sumPI}(f)$ by Laplante et al. [LLS06]. Laplante et al. also define a more general bound $\text{maxPI}(f)$ which satisfies $L(f) \geq \text{maxPI}^2(f) \geq \text{sumPI}^2(f)$. Hence to prove Proposition 6.13, it suffices to show that $WA^2(f) \geq \text{Adv}_s(f)$.

Claim D.2. *For any $f: \{0, 1\}^n \rightarrow \{0, 1\}$ it holds that $WA^2(f) \geq \text{Adv}_s(f)$.*

Proof. Let (\mathbf{a}, \mathbf{b}) be a distribution on $f^{-1}(1) \times f^{-1}(0)$ that attains $\text{Adv}_s(f)$. Let R be its support. For $(a, b) \in R$, define $w(a, b) = \Pr[(\mathbf{a}, \mathbf{b}) = (a, b)]$. For $(a, b) \in R$ and $i \in [n]$ such that $a_i \neq b_i$, define $w'(a, b, i) = w'(b, a, i) = w(a, b)$. The weighting scheme is trivially valid.

Given $a \in A$ and $i \in [n]$, clearly $\Pr[\mathbf{a} = a] = \sum_{b: (a,b) \in R} w(a, b)$ and $\Pr[\mathbf{a} = a, \mathbf{a}_i \neq \mathbf{b}_i] = \sum_{b: (a,b) \in R, a_i \neq b_i} w'(a, b, i)$. Therefore

$$\frac{1}{\Pr[\mathbf{a}_i \neq \mathbf{b}_i \mid \mathbf{a} = a]} = \frac{\sum_{b: (a,b) \in R} w(a, b)}{\sum_{b: (a,b) \in R, a_i \neq b_i} w'(a, b, i)}.$$

Thus the first factor in the definition of $WA^2(f)$ equals the first factor in the definition of $\text{Adv}_s(f)$. The same holds for the second factors, and we conclude that $WA^2(f) \geq \text{Adv}_s(f)$. \square