

Pseudodeterministic Communication Complexity

Mika Göös
EPFL

Nathaniel Harms
University of British Columbia

Artur Riazanov
EPFL

Anastasia Sofronova
EPFL

Dmitry Sokolov
EPFL & Université de Montréal

Weiqiang Yuan
EPFL

November 20, 2025

Abstract

We exhibit an n -bit partial function with randomized communication complexity $O(\log n)$ but such that any completion of this function into a total one requires randomized communication complexity $n^{\Omega(1)}$. In particular, this shows an exponential separation between randomized and *pseudodeterministic* communication protocols. Previously, Gavinsky (2025) showed an analogous separation in the weaker model of parity decision trees. We use lifting techniques to extend his proof idea to communication complexity.

Contents

1	Introduction	1
2	Warm-Up and Proof Overview	5
3	Proof of the Main Theorem	10
4	Density Restoring Partitions and Protocol Transformation	17
5	Completing the Safe Stage Lemma	21
A	Appendix: Counting Arguments and Variations of FBPP	27
	References	31

1 Introduction

Here is a simple question about the behavior of randomized algorithms. A basic statistical task is to distinguish “few” vs. “many”, formalized by the GAP MAJORITY problem:

$$\text{GAPMAJ}(x) := \begin{cases} 1 & \text{if } |x| \geq \frac{2}{3}n, \\ 0 & \text{if } |x| \leq \frac{1}{3}n, \\ * & \text{otherwise,} \end{cases}$$

where $|x|$ denotes the number of 1s in $x \in \{0, 1\}^n$ and $*$ indicates that we put no requirement on the output of an algorithm (that is, GAPMAJ is a search problem with 0 and 1 both being acceptable outputs). This problem is difficult to solve deterministically: It requires $\Omega(n)$ queries for decision trees and parity decision trees, and $\Omega(n)$ bits of deterministic communication when we turn it into a two-player problem by composing it with an appropriate gadget. For example, composing with XOR yields the GAP HAMMING DISTANCE problem given by $\text{GAPHD}(x, y) := \text{GAPMAJ}(x \oplus y)$. On the other hand, for *randomized* decision trees (and the other models as well), the cost of GAPMAJ is only 1 query because the algorithm can sample a random coordinate and output it; when $|x| \geq \frac{2}{3}n$ or $|x| \leq \frac{1}{3}n$ this is correct with probability at least $2/3$. Our simple question is:

What is the randomized algorithm doing on the $$ inputs?*

On inputs with $|x| = n/2$, the sampling algorithm will output 0 or 1 with equal probability. But *must it* do this, or can we ask that the randomized algorithm produce consistent outputs for every input x ? That is, can we ask that the algorithm computes some *completion* of GAPMAJ into a total function: on any input x it should output either value 0 with probability at least $2/3$, or value 1 with probability at least $2/3$. Such an algorithm is called *pseudodeterministic*. Generally, a pseudodeterministic algorithm is a randomized algorithm that is required to output (with probability at least $2/3$) a single consistent output for every input; this is desirable not only for partial boolean functions, but for any problem where there is more than 1 acceptable output for each input, that is, for any search problem or relation. The goal is to combine the efficiency of randomized algorithms with the consistency of deterministic algorithms. So, are pseudodeterministic algorithms nearly as efficient as randomized ones?

We show that, for communication protocols, the answer is *no*: there is a partial boolean function with an efficient randomized protocol but such that every total completion of that function has large randomized complexity.

Theorem 1. *There is a partial communication problem $\{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1, *\}$ with randomized communication complexity $O(\log n)$ but pseudodeterministic communication complexity $\Omega(\sqrt{n})$.*

Partial n -bit boolean functions with $\text{poly log } n$ randomized communication cost are a type of “BPP search problems”: relations $R \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ computed by an efficient randomized protocol. By analogy to Turing machine complexity classes, the class of “BPP search problems” would be called FBPP. If we use FPs to denote the class of relations admitting pseudodeterministic protocols with cost $\text{poly log } n$, Theorem 1 is the first explicit example witnessing the separation

$$\text{FPs} \subsetneq \text{FBPP}, \tag{1}$$

and it does so in the most restricted setting, where the output is only a single bit. We remark, however, that (1) has a caveat: just like the Turing machine analogue, the communication class

FBPP has several natural definitions, which are not equivalent [Aar10, Gol11, ILW23, ABK24]. For one of these definitions (which requires the output to be efficiently verifiable), the inclusion (1) holds in the *opposite direction* for partial boolean functions; our Theorem 1 gives an explicit witness for (1) under the remaining definitions. Weaker versions of (1), which allow the outputs to be large and do not require an explicit example, can be proved by counting; see Appendix A for a discussion of these nuances.

Several recent works [GGMW20, GIPS21, Gav25, FGHH25, BHH⁺25] asked for lower bounds on pseudodeterministic communication complexity for BPP search problems. Many of them proved versions of (1) for weaker models of computation: Goldwasser, Grossman, Mohanty, and Woodruff [GGMW20] proved it for one-way communication protocols; Goldwasser, Impagliazzo, Pitassi, and Santhanam [GIPS21] proved it for decision trees; and Gavinsky [Gav25] proved it for parity decision trees. Our result can thus be viewed as a qualitative strengthening of these prior works (the upper bound in Theorem 1 holds even for randomized non-adaptive decision trees).

For partial functions, the only known prior separation was due to Blondal, Hatami, Hatami, Lalov, and Titiak [BHH⁺25]. They showed an $O(1)$ -vs- $\Omega(\log \log n)$ separation for GAP HAMMING DISTANCE, i.e., the XOR-lift $\text{GAPMAJ}(x \oplus y)$. This would be the ideal function to witness the separation (1) for partial functions, and we suspect it exhibits the maximum possible separation:

Conjecture 2. *The pseudodeterministic communication complexity of GAPHD is $\Omega(n)$.*

This was one of the principal inspirations for our present work, but our Theorem 1 ultimately does not prove the separation using GAPHD; we instead use a slightly more complicated lift of GAPMAJ, as explained next.

1.1 Our techniques

Our proof extends the work of Gavinsky [Gav25]. He proved an $\Omega(\sqrt{n})$ query lower bound for pseudodeterministic parity decision trees computing the GAPMAJ partial function. To turn this function into a communication problem, we compose it with the INNER PRODUCT gadget on $m = O(\log n)$ bits defined by $\text{IP}_m(a, b) := \sum_{i=1}^m a_i b_i \bmod 2$. This produces a communication problem on $O(n \log n)$ bits:

$$\text{GAPMAJ} \circ \text{IP}_m^n(x_1, \dots, x_n, y_1, \dots, y_n) := \text{GAPMAJ}(\text{IP}_m(x_1, y_1), \dots, \text{IP}_m(x_n, y_n)).$$

Here $x_i, y_i \in \{0, 1\}^m$, Alice has all the x_i inputs, and Bob has all the y_i inputs. This problem has randomized communication complexity $O(\log n)$ because the players can select $O(1)$ random indices $i \in [n]$ and solve $\text{IP}_m(x_i, y_i)$ deterministically.

Standard randomized lifting theorems [GPW20, CFK⁺21] show that for every partial function $f: \{0, 1\}^n \rightarrow \{0, 1, *\}$ the randomized communication complexity of $f \circ \text{IP}_m^n$ equals roughly the randomized query complexity of f . It is open to prove such a general lifting theorem for pseudodeterministic complexity. The immediate obstacle to applying a randomized lifting theorem to a pseudodeterministic protocol for $f \circ \text{IP}_m^n$ is that the protocol may produce different consistent outputs for inputs (x, y) and (x', y') such that $\text{IP}_m^n(x, y) = \text{IP}_m^n(x', y') \in f^{-1}(*)$. Since we cannot invoke known lifting theorems as a black box, we employ a *white-box* approach: We use tools from lifting theory to adapt Gavinsky’s argument for communication protocols. This adaptation introduces new technical challenges not present in the setting of parity decision trees. We give an overview of the proof in Section 2, including an exposition of Gavinsky’s original argument.

1.2 History and Motivations

Pseudodeterministic algorithms were introduced by Gat and Goldwasser [GG11] (under the name *Bellagio* algorithms), and independently by Huynh and Nordström [HN12] in the context of proof complexity (they called them *consistent* algorithms). Pseudodeterministic algorithms have applications to cryptography and distributed computing, and have been studied in several computational models including Turing machines [GGR13, OS17, LOS21, CLO⁺23], decision trees [GGR13, GIPS21, CDM23], and parity decision trees [Gav25]. They are natural and interesting in their own right, they provide an intermediary between randomized and deterministic complexity, and they are also related to notions of *replicability* in machine learning and the natural sciences [GIPS21, ILPS22, CMY23, BHH⁺25].

Aside from intrinsic interest in pseudodeterminism, [Gav25] points out its connection to structural questions about randomized communication (which we discuss in the open problems below), and [GIPS21] argues that understanding pseudodeterministic communication is an important step towards understanding the communication complexity of search problems (i.e., relations). Whereas the communication complexity of *functions* benefits from connections to well-understood query complexity measures (e.g., sensitivity and block-sensitivity) via lifting theorems, analogues of these measures for search problems are not well understood. Pseudodeterministic algorithms must compute *some* function, and therefore serve as an intermediate between functions and relations.

Pseudodeterminism was introduced in [HN12] for applications in proof complexity. That paper shows that *cutting plane proofs* of CNF unsatisfiability can be turned into pseudodeterministic communication protocols for the *falsified-clause* search problem where two players search for a clause that is not satisfied by a (distributed) assignment of variables. Their lower bound for this communication problem allowed them to prove time–space tradeoffs for cutting planes.

1.3 Open Problems

Improved quantitative bounds. Let us start by reiterating that we conjecture GAP HAMMING DISTANCE to require pseudodeterministic complexity $\Omega(n)$ (Conjecture 2). Proving this seems to require two steps. First, improve Gavinsky’s lower bound of $\Omega(\sqrt{n})$ for parity decision trees computing GAPMAJ to $\Omega(n)$. Second, lift this to GAPHD. Another strategy may be to directly improve the current lower bound for the pseudodeterministic communication complexity of GAPHD, which is $\Omega(\log \log n)$, and which uses very different techniques [BHH⁺25].

It is important to note that GAP HAMMING DISTANCE is, in a precise sense, the *only* partial boolean communication problem where we can hope to prove an $O(1)$ -vs- $\text{poly}(n)$ separation between randomized and pseudodeterministic communication. This is because *all* n -bit partial boolean matrices with randomized cost $O(1)$ are a submatrix of GAPHD on $O(n)$ bits (with some constant $\alpha < 1/2$ in place of $1/3$ in the gap) [LS09, FGHH25]. So a pseudodeterministic lower bound for any of these matrices implies the same lower bound for GAPHD.

Separation for a TFNP problem. The most outstanding *qualitative* problem that is left open is to prove pseudodeterministic lower bounds for a *total NP search problem* (TFNP). A relation $R \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ is in (communication analogue of) TFNP if

- R is *total*, meaning that for all inputs (x, y) there is a solution z such that $(x, y, z) \in R$; and
- R is *efficiently checkable*, meaning that there is a deterministic verifier protocol V of cost $\text{poly} \log n$ such that $V(xz, yz) = 1$ iff $(x, y, z) \in R$ for all (x, y, z) .

Conjecture 3. *There exists some TFNP communication problem with randomized communication complexity $\text{poly log } n$ but pseudodeterministic complexity $n^{\Omega(1)}$.*

The analogous separation for query complexity was the main result of the paper [GIPS21]. The communication analogue of TFNP has been studied explicitly in, e.g., [GKRS19, BFI23, GGJL25], but, implicitly, TFNP problems are ubiquitous in communication complexity: this includes all (monotone) Karchmer–Wigderson games as well as the aforementioned falsified-clause search problems, which arise when applying communication lower bounds to proof complexity. The survey [dRGR22] explains these connections and more: TFNP problems give a unified lens to study the interconnections between communication protocols, propositional proofs, and boolean circuits.

Separation for a BPP-verifiable problem. In fact, it is still open to prove a weaker result than Conjecture 3. If we define BPP search problems for communication complexity by adapting the definition of Goldreich [Gol11] for Turing machines, we get the class of relations $R \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ such that

- There is a randomized $\text{poly log}(n)$ cost protocol solving R ; and
- There is a randomized $\text{poly log}(n)$ cost *verifier* protocol V : $\Pr[V(xz, yz) = 1] \geq 2/3$ if $(x, y, z) \in R$ and $\Pr[V(xz, yz) = 1] \leq 1/3$ otherwise, for all (x, y, z) .

We can then replace the TFNP requirement in Conjecture 3 with these BPP-verifiable problems and the conjecture is still open.

Our result implies this separation for a weaker, more specialized type of verification. The partial function $\text{GAPMAJ} \circ \text{IP}$ does *not* satisfy the above definition, because the verifier cannot exist (we could use it to solve *exact* majority efficiently, which cannot be done). However, we can define the relation $R \subseteq \{0, 1\}^{nm} \times \{0, 1\}^{nm} \times [n]$ where for each x, y the valid outputs are the numbers $t \in [n]$ that are within $\pm n/10$ of the number of 1-valued IP_m gadgets. This relation can be solved in cost $O(m) = O(\log n)$ the same way as $\text{GAPMAJ} \circ \text{IP}_m$. It can also be *sort-of* verified: given $(x, y, \hat{t})$, if the correct number of 1-valued gadgets is t , there is a randomized verifier V which satisfies

- If $|t - \hat{t}| \leq n/10$ (i.e., $(x, y, \hat{t}) \in R$) then $\Pr[V(x, y, \hat{t}) = 1] \geq 2/3$; and
- If $|t - \hat{t}| \geq n/9$ then $\Pr[V(x, y, \hat{t}) = 1] \leq 1/3$.

An efficient pseudodeterministic protocol for this estimation problem could be used as a pseudodeterministic protocol for $\text{GAPMAJ} \circ \text{IP}_m$, so Theorem 1 shows that it cannot exist. However, the necessary gap between $n/10$ and $n/9$ means that this does not satisfy the definition of [Gol11].

Structure of communication protocols. Communication complexity is closely related to the size of monochromatic rectangles within the communication matrices: efficient deterministic protocols imply large monochromatic rectangles inside the matrix. The question of whether the same is true for randomized protocols that compute a total boolean function was raised in [GKPW18]. A striking version of this question is obtained for constant-cost protocols: Chattopadhyay, Lovett, and Vinyals [CLV19] and Hambardzumyan, Hatami, and Hatami [HHH23] conjecture that

Conjecture 4 ([CLV19, HHH23]). *There exists a function η such that every $N \times N$ boolean matrix with randomized communication cost c has an $\eta(c) \cdot N \times \eta(c) \cdot N$ monochromatic rectangle.*

Completions of the GAPHD matrix *cannot* have such large monochromatic rectangles [FF81], and therefore Conjecture 4 already implies lower bounds on the pseudodeterministic cost of GAPHD, as noted in [FGHH25]. Gavinsky [Gav25] points out similar implications for parity decision trees, with monochromatic affine subspaces in place of rectangles. The best progress so far towards Conjecture 4 is to find such monochromatic rectangles in matrices of bounded γ_2 -norm [BHT25].

2 Warm-Up and Proof Overview

Our proof starts with the lower bound for parity decision trees by Gavinsky [Gav25] and upgrades it to a communication lower bound using techniques from query-to-communication lifting [GLM⁺16, GPW20]. As a warm-up to our main proof, we will present an exposition of Gavinsky’s argument, but simplified to the special case of decision trees.

Of course, for decision trees (rather than parity decision trees), there is a simpler proof of a superior $\Theta(n)$ bound: in any completion f of GAPMAJ, find the input $x \in f^{-1}(0)$ with largest weight; then, observe that by fixing the 1-valued coordinates of x and letting the 0-valued coordinates vary, we obtain the OR problem on $\Omega(n)$ bits, which implies a lower bound of $\Omega(n)$ queries. However, the more complicated warm-up proof gives a technique that can be lifted to communication complexity.

2.1 Warm-Up: Lower Bound for Decision Trees

Theorem 5 (Simplified version of [Gav25]). *Any completion $f: \{0,1\}^n \rightarrow \{0,1\}$ of GAPMAJ requires randomized query complexity $\Omega(\sqrt{n})$.*

Proof sketch. The proof has two parts, encapsulated in the *Closeness Lemma* and the *Stage Lemma*. The Closeness Lemma formalizes the key observation that a shallow decision tree cannot distinguish between the uniform distribution over $\{0,1\}^n$ and the distribution where the bits are slightly biased towards 1. Specifically, suppose $\mathbf{u} \sim \{0,1\}^n$ is uniform random and $\mathbf{x}$ is obtained from $\mathbf{u}$ by setting to 1 a uniformly random set of $\sqrt{n}$ coordinates; that is, we “sprinkle” some $\sqrt{n}$ many 1s into $\mathbf{u}$. Then the probability that a depth- $o(\sqrt{n})$ decision tree T can distinguish $\mathbf{u}$ from $\mathbf{x}$ is $o(1)$; in other words, the output distributions $T(\mathbf{u})$ and $T(\mathbf{x})$ are close to each other.

Using this Closeness Lemma, we break down the argument into stages, each stage handled by the Stage Lemma. In each stage, we start with a subcube $X \subseteq \{0,1\}^n$ where at least half the strings $x \in X$ satisfy $f(x) = 0$. A subcube X is equivalent to a partial assignment of variables, where a set $F \subseteq [n]$ are “free”, while variables $[n] \setminus F$ are “fixed”, i.e., all $x \in X$ agree on the coordinates $[n] \setminus F$. Our goal in each stage is to find $\approx \sqrt{n}$ free variables to fix (i.e., to remove from F to obtain a new subcube); essentially, these newly fixed variables are the ones queried by the decision tree, along with the $\sqrt{n}$ “sprinkled” 1-bits. Our choice should satisfy the property that, after fixing these new variables, we maintain the invariant that at least half the remaining strings x have $f(x) = 0$. If we can accomplish this goal in each stage, repeating this process over $0.9\sqrt{n}$ stages yields a contradiction as illustrated in Figure 1: on one hand, we have fixed $0.9\sqrt{n} \cdot \sqrt{n} = 0.9n$ bits to 1, so f is identically 1 on all remaining strings. On the other hand, we maintained the invariant that half the remaining strings x have $f(x) = 0$, forcing a contradiction.

We find the $\approx \sqrt{n}$ variables to fix as follows. The key trick is that, by Yao’s principle, there is a deterministic $o(\sqrt{n})$ -depth decision tree T that errs with probability at most ε on the distribution

$$\frac{1}{2}(\text{unif}(X) + \sigma),$$

where σ is the “sprinkled-1s” distribution: start with a uniform $\mathbf{u} \sim X$ and construct $\mathbf{x}$ by fixing a random set of $\sqrt{n}$ free variables to 1. Crucially, T has error probability 2ε over each component distribution $\text{unif}(X)$ and σ . Since $\Pr_{\mathbf{x} \sim X}[f(\mathbf{x}) = 0] \geq 1/2$ and T has error 2ε over $\text{unif}(X)$, we can easily find a 0-leaf ℓ where $f(x) = 0$ for most x reaching that leaf; by fixing the variables queried by T on the path to ℓ , we obtain a subcube L with $\Pr_{\mathbf{x} \sim L}[f(\mathbf{x}) = 0] \geq 1/2$. We can “upgrade” the last step to let us fix an additional $\sqrt{n}$ variables to 1, by using the fact that T also has error 2ε over σ , and the Closeness Lemma, which says that the distribution over leaves of T is similar for σ as for $\text{unif}(X)$.

Full proof (skippable on first reading). We now give a more formal proof that serves as a blueprint for our communication lower bound. The following lemma formalizes the process for each stage, that is, the act of assigning $(1 + o(1))\sqrt{n}$ bits. To better match with the communication proof we view the conditioning on a partial assignment as zooming in to a subcube of $\{0, 1\}^n$.

Notation: Throughout the paper, for any string $x \in \Sigma^n$ and a set $S \subseteq [n]$, we write $x_F \in \Sigma^F$ for the substring of x on coordinates F . For $X \subseteq \Sigma^n$, we write $X_F := \{x_F \mid x \in X\}$.

Lemma 6 (Stage Lemma for Decision Trees). *Suppose $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is a completion of GAPMAJ with randomized query complexity $o(\sqrt{n})$. Let $X \subseteq \{0, 1\}^n$ be a subcube of the boolean cube, that is, there exists $F \subseteq [n]$ be such that $X_{[n] \setminus F}$ is fixed and X_F is free. Suppose that $|F| \geq 0.1n$ and $\Pr_{\mathbf{u} \sim X}[f(\mathbf{u}) = 0] \geq 1/2$. Then there exists a subcube $X' \subseteq X$ and $F' \subseteq F$ such that*

$$\Pr_{\mathbf{u}' \sim X'}[f(\mathbf{u}') = 0] \geq \frac{1}{2} \quad \text{and} \quad X_{F'} \text{ is free.}$$

Moreover, there is a set $R \subseteq F \setminus F'$ such that: $X'_R = 1_R$ and $|R| \geq 0.9|F \setminus F'|$.

This lemma suffices to prove the lower bound, as follows (see Figure 1 for an illustration):

Proof of Theorem 5 given Lemma 6. We may assume that $|f^{-1}(0)| \geq 2^{n-1}$, as otherwise we can swap the roles of 0/1 output values in the upcoming argument. Begin with $X = \{0, 1\}^n$, $F = [n]$, and $S = \emptyset$. We will apply the stage lemma iteratively as long as $|F| \geq 0.1n$ and update X, F, S in each stage, as follows: assuming the invariant $\Pr_{\mathbf{u} \sim X}[f(\mathbf{u}) = 0] \geq 1/2$, we apply the stage lemma and update

- $X \leftarrow X'$ with X' from the stage lemma, maintaining the invariant $\Pr_{\mathbf{u}' \sim X'}[f(\mathbf{u}') = 0] \geq 1/2$;
- $F \leftarrow F'$ from the stage lemma;
- $S \leftarrow S \cup R$ from the stage lemma, maintaining the invariant that all $x \in X'$ have value 1 on coordinates S .

On termination, we have preserved the invariant $\Pr_{\mathbf{u} \sim X}[f(\mathbf{u}) = 0] \geq 1/2$, but on the other hand, since $|S| \geq 0.9|[n] \setminus F| \geq 0.8n$ and all $x \in X$ have coordinates in S fixed to 1, we have $f(x) = 1$ for all $x \in X$, a contradiction. $\square$

To prove the stage lemma, we define the following distribution over inputs. Given any subcube $X \subseteq \{0, 1\}^n$ with free variables $F \subseteq [n]$, we define σ as the distribution over X obtained by choosing a uniformly random $\mathbf{u} \sim X$ and a uniformly random set $\mathbf{I} \sim \binom{F}{\sqrt{n}}$ of $\sqrt{n}$ free variables. Then we take $\mathbf{x} \in X$ to be the string obtained from $\mathbf{u}$ by setting all coordinates $i \in \mathbf{I}$ to 1.




Figure 1: Decision tree stages

This distribution has two important properties: first, it has $\sqrt{n}$ “planted” 1-valued coordinates, so that it is more biased towards 1 than a uniformly random string. Second, it is indistinguishable from the uniform distribution by $o(\sqrt{n})$ -query decision trees, as established in the *Closeness Lemma*:

Lemma 7 (Closeness Lemma for Decision Trees). *Let $X' \subseteq X$ be a subcube of X and $F' \subseteq F$ the set of its free variables. Suppose $|F \setminus F'| \leq o(\sqrt{n})$ and $|F| \geq 0.1n$. Then*

$$\Pr_{\mathbf{x} \sim \sigma}[\mathbf{x} \in X'] \geq (1 - o(1)) \Pr_{\mathbf{u} \sim X}[\mathbf{u} \in X'].$$

Proof. Let $\mathbf{I} \sim \binom{F}{\sqrt{n}}$ be as in the definition of σ . Then $\mathbf{x} \sim \sigma$ is obtained by choosing $\mathbf{u} \sim X$ and then fixing variables in $\mathbf{I}$ to 1. If $\mathbf{u} \in X'$ and the set $\mathbf{I} \subseteq F$ does not intersect the set $F \setminus F'$ of variables which are fixed by X' , then it must be the case that $\mathbf{x} \in X'$ as well. Therefore

$$\Pr_{\mathbf{x} \sim \sigma}[\mathbf{x} \in X'] \geq \Pr_{\mathbf{u} \sim X}[\mathbf{u} \in X' \wedge (\mathbf{I} \cap (F' \setminus F) = \emptyset)] = \Pr_{\mathbf{u} \sim X}[\mathbf{u} \in X'] \cdot \Pr[\mathbf{I} \cap (F' \setminus F) = \emptyset].$$

By the union bound,

$$\Pr[\mathbf{I} \cap (F' \setminus F) \neq \emptyset] \leq \sqrt{n} \cdot \frac{|F' \setminus F|}{|F|} \leq \sqrt{n} \cdot \frac{o(\sqrt{n})}{0.1n} = o(1). \quad \square$$

We now conclude the lower bound by proving the stage lemma:

Proof of Lemma 6. Suppose that f has randomized decision tree complexity $o(\sqrt{n})$; we may assume that the randomized decision tree has error probability ε . Let X be any subcube with free variables F satisfying $|F| \geq 0.1n$. Recall the distribution σ over X where $\mathbf{x} \sim \sigma$ is sampled by choosing $\mathbf{I} \sim \binom{F}{\sqrt{n}}$ and $\mathbf{x} \sim \text{unif}(\{x \in X \mid x_{\mathbf{I}} = 1_{\mathbf{I}}\})$.

By Yao's principle, there exists a deterministic decision tree T of depth $o(\sqrt{n})$, computing f with error ε over the mixture distribution

$$\frac{1}{2}(\text{unif}(X) + \sigma).$$

Observe that tree T has error at most 2ε over each individual distribution σ and $\text{unif}(X)$.

Note that each leaf ℓ of T corresponds to a subcube $L \subseteq X$ obtained by fixing the $o(\sqrt{n})$ bits queried by T on the path to ℓ ; therefore the free variables F_L of L satisfy $|F \setminus F_L| = o(\sqrt{n})$. By applying the closeness lemma, Lemma 7, to each subcube L corresponding to the 0-leaves of T , we obtain

$$\Pr_{\mathbf{x} \sim \sigma}[T(\mathbf{x}) = 0] = \sum_{\text{0-leaf } L} \Pr_{\mathbf{x} \sim \sigma}[\mathbf{x} \in L] \geq (1 - o(1)) \Pr_{\mathbf{u} \sim X}[T(\mathbf{u}) = 0] \geq (1 - o(1)) \left(\frac{1}{2} - 2\varepsilon\right) \geq 1/4,$$

where the penultimate inequality uses the assumption $\Pr_{\mathbf{u} \sim X}[f(\mathbf{u}) = 0] \geq 1/2$. Now, since T errs with probability at most 2ε on σ , we have

$$\Pr_{\mathbf{x} \sim \sigma}[f(\mathbf{x}) = 1 \mid T(\mathbf{x}) = 0] = \frac{\Pr_{\mathbf{x} \sim \sigma}[f(\mathbf{x}) = 1 \wedge T(\mathbf{x}) = 0]}{\Pr_{\mathbf{x} \sim \sigma}[T(\mathbf{x}) = 0]} \leq 8\varepsilon.$$

Hence, according to the total probability law, there exists a 0-leaf ℓ in T such that $\Pr_{\mathbf{x} \sim \sigma}[f(\mathbf{x}) = 1 \mid \mathbf{x} \in L] \leq 8\varepsilon \leq 1/2$. Again by total probability law, there exists $\mathbf{I} \in \binom{F}{\sqrt{n}}$ such that $\Pr_{\mathbf{x} \sim \sigma}[f(\mathbf{x}) = 1 \mid \mathbf{x} \in L, \mathbf{I} = \mathbf{I}] \leq 1/2$. We then observe that for $\mathbf{u} \sim X$ the distributions of $(\mathbf{u} \mid \mathbf{u}_{\mathbf{I}} = 1)$ and $(\mathbf{x} \mid \mathbf{I} = \mathbf{I})$ coincide, therefore $\Pr_{\mathbf{u} \sim X}[f(\mathbf{u}) = 1 \mid \mathbf{u} \in \ell, \mathbf{u}_{\mathbf{I}} = 1] \leq 1/2$. We then conclude by defining $X' := \{x \in L \mid x_{\mathbf{I}} = 1_{\mathbf{I}}\}$ with $R := \mathbf{I}$ being the variables forced to 1, and the new set of free variables being $F' := F_L \setminus R$ (the free variables remaining in the leaf L , minus those forced to 1), which satisfies $|R| = \sqrt{n} \geq 0.9|F \setminus F'| = 0.9(\sqrt{n} + o(\sqrt{n}))$. $\square$

2.2 Proof Overview

Let us describe how to lift the argument for decision trees in Section 2.1 to prove a lower bound for the communication complexity of $\text{GAPMAJ} \circ \text{IP}_m^n$. For the sake of brevity, we will write $\Sigma := \{0, 1\}^m$ for the alphabet of inputs to the IP_m gadget, so that both inputs to $\text{GAPMAJ} \circ \text{IP}_m^n$ are in domain Σ^n . For the remainder of the paper, $m = \log |\Sigma|$ indicates the bit-size of the IP_m gadget. Below, we use $R(\cdot)$ to denote randomized communication complexity (to within error $1/3$).

Theorem 8. *Any completion $f: \Sigma^n \times \Sigma^n \rightarrow \{0, 1\}$ of $\text{GAPMAJ} \circ \text{IP}_m^n$, with $m := 100 \log n$, has*

$$R(f) = \Omega(\sqrt{n} \log n).$$

Our proof proceeds analogously to the proof for query complexity in Section 2.1. At a high level, we have the following analogies:

- For decision trees we tracked a subcube $X \subseteq \{0, 1\}^n$, because leaves of a decision tree correspond to subcubes. Now we track a rectangle $X \times Y \subseteq \Sigma^n \times \Sigma^n$ because leaves of a communication protocol correspond to rectangles. (Gavinsky's argument for parity decision trees tracks an affine subspace $X \subseteq \mathbb{F}_2^n$.)

- We still have a set F of “free” variables, but these take on a different meaning, because the rectangle $X \times Y$ no longer corresponds exactly to a set of inputs obtained by fixing variables $[n] \setminus F$ and leaving those in F completely unrestricted; this is because communication protocols may exchange information about variables without entirely “fixing” them. Our “free variables” F instead satisfy the condition that “only a small amount of information about them is known”.
- We still have a *Stage Lemma*, which suffices to prove the theorem in almost the same way: starting with a rectangle $X \times Y$ where $f(x, y) = 0$ for most $(x, y) \in X \times Y$, we iteratively restrict to a smaller rectangle $X' \times Y'$ by “fixing” some free variables, including $\sqrt{n}$ variables fixed to 1. Unlike the decision tree proof, we can no longer fix $\sqrt{n}$ variables to 1 in *every* stage, but only *most* stages (which we call “safe stages”).
- We still have a *Closeness Lemma*, which we use to find $\sqrt{n}$ variables to fix to 1; however, the simple birthday paradox argument in [Lemma 7](#) no longer works, because a protocol does not simply query $o(\sqrt{n})$ bits, so $\sqrt{n}$ random coordinates will not simply “miss” the queried coordinates with high probability. Indeed, our new Closeness Lemma will only apply in the “safe stages”.

Let’s give a little more detail on how the Stage Lemma and Closeness Lemma differ from their decision tree analogues. We also provide a diagram of the Stage Lemma in [Figure 3](#).

Stage Lemma: Suppose for the sake of contradiction that there exists a randomized communication protocol with error ε and communication cost $o(\sqrt{nm})$. We proceed as follows:

1. Start each stage with a rectangle $X \times Y$ (originally $\Sigma^n \times \Sigma^n$) which contains mostly 0-valued inputs to f , and a set $F \subseteq [n]$ of “free variables”. We keep track of the “deficit”

$$\mathcal{D}(X \times Y, F) := 2|F|m - H_\infty(\mathbf{x}_F) - H_\infty(\mathbf{y}_F),$$

where $(\mathbf{x}, \mathbf{y}) \sim X \times Y$ and $H_\infty(\mathbf{x}) := \min_x \log(1/\Pr[\mathbf{x} = x])$ denotes min-entropy. Intuitively, the deficit quantifies how many bits of information Alice and Bob know about the free variables F . For intuition, note that if Alice and Bob were merely simulating a decision tree (solving individual gadgets one at a time) then the deficit would always be 0.

2. By Yao’s principle we obtain a deterministic protocol Π with cost $|\Pi| = o(\sqrt{nm})$ and error ε over the distribution

$$\frac{1}{2} (\text{unif}(X \times Y) + \sigma(X \times Y, F)) ,$$

where $\sigma(X \times Y, F)$ is a *sprinkled-1s* distribution ([Definition 14](#)). In the sprinkled-1s distribution, an input $(\mathbf{x}, \mathbf{y})$ is chosen uniformly from $X \times Y$, and then a random set of $\sqrt{n}$ gadgets in F are fixed to have value 1, i.e., we force the input distribution to increase the number of 1-valued gadgets by at least $\sqrt{n}$.

3. We use the fact that Π has error 2ε on each of $\text{unif}(X \times Y)$ and $\sigma(X \times Y, F)$, together with the Closeness Lemma (which says that distribution over leaves induced by $\sigma(X \times Y, F)$ and $\text{unif}(X \times Y)$ is similar), to find a smaller subrectangle $X' \times Y' \subseteq X \times Y$ and smaller set of free variables $F' \subseteq F$, where $\sqrt{n}$ new IP gadgets are fixed to 1. However, the Closeness Lemma will now apply only in “safe stages” where the deficit is low; in the unsafe stages, we skip the step of fixing $\sqrt{n}$ gadgets to 1. This is OK, because most stages will be safe.

Closeness Lemma: The Closeness Lemma should show that the distribution over leaves of the protocol is similar under the sprinkled-1s distribution $\sigma(X \times Y, F)$ and $\text{unif}(X \times Y)$, so that we can safely sprinkle $\sqrt{n}$ 1-values “without the protocol noticing”. However, to prove this claim, there are two caveats, which constitute the main technical novelty of the proof, and do not have an analogue in Gavinsky’s parity decision tree proof:

1. In each stage, we must first transform Π into a more structured protocol Π' using techniques from lifting and properties of the IP_m gadget.
2. We can prove the claim only when the deficit $\mathcal{D}(X \times Y, F)$ is *small* (i.e., there is little information known about the free variables F , so that intuitively they behave similarly to the completely free variables in the decision tree argument).

In the next section, we formalize the main claims required for this argument.

3 Proof of the Main Theorem

For the sake of contradiction, we make the following assumption throughout the proof:

Assumption 9. Assume that there is a completion of $\text{GAPMAJ} \circ \text{IP}_m^n$, denoted $f: \Sigma^n \times \Sigma^n \rightarrow \{0, 1\}$, such that the randomized communication cost of f is $o(\sqrt{n}m)$ and $\Pr_{(\mathbf{x}, \mathbf{y})}[f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2$.

As in the decision tree proof, the main theorem will follow from a Stage Lemma. To state the communication version of the Stage Lemma, we need a few definitions. First, we will use the *deficit* as a potential function throughout the stage procedure:

Definition 10 (Deficit). For a set $F \subseteq [n]$ and a random variable $\mathbf{x}$ over Σ^n , we write

$$D_\infty(\mathbf{x}_F) := |F|m - H_\infty(\mathbf{x}_F).$$

For a rectangle $X \times Y \in \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$ we define the *deficit*:

$$\mathcal{D}(X \times Y, F) := D_\infty(\mathbf{x}_F) + D_\infty(\mathbf{y}_F) = 2|F|m - H_\infty(\mathbf{x}_F) - H_\infty(\mathbf{y}_F),$$

where $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. Observe that the deficit is non-negative.

Next, we will maintain the property that, for the current rectangle $X \times Y$ in the process, the sets X and Y are pseudorandom in sense of being γ -spread:

Definition 11 (Spread Variables). A random variable $\mathbf{x} \in \Sigma^J$ is γ -spread if $\forall I \subseteq J$ it holds that $H_\infty(\mathbf{x}_I) \geq \gamma|I|m$. A set $X \subseteq \Sigma^J$ is γ -spread if $\mathbf{x} \sim \text{unif}(X)$ is γ -spread.

The reason we make this definition and maintain the spreadness property throughout the stage process is that it lets us take advantage of near-uniformity of the INNER PRODUCT gadget outputs on the free variables. Let us state this lemma now to make the motivation of this definition clear:

Lemma 12 ([GLM⁺16, Lemma 13]). Suppose $m \geq 100 \log n$. Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$ be such that the random variables $\mathbf{x}_F$ and $\mathbf{y}_F$ are 0.9-spread, for $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. Then

$$\forall z \in \{0, 1\}^F: \quad \Pr[(\text{IP}_m(\mathbf{x}_i, \mathbf{y}_i))_{i \in F} = z] \in 2^{-|F|} \cdot (1 \pm 2^{-m/20}).$$

We may now state the communication version of the Stage Lemma.

Lemma 13 (Stage Lemma). *Assume [Assumption 9](#). Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ be a rectangle and $F \subseteq [n]$, such that: $|F| \geq 0.1n$,*

$$\Pr_{\mathbf{x}, \mathbf{y} \sim X \times Y} [f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2, \quad \text{and} \quad X_F \text{ and } Y_F \text{ are 0.9-spread.}$$

Then there exist $X' \times Y' \subseteq X \times Y$ and $F' \subsetneq F$, such that: $|F \setminus F'| \geq \Omega(\sqrt{n})$,

$$\Pr_{\mathbf{x}, \mathbf{y} \sim X' \times Y'} [f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2, \quad \text{and} \quad X'_{F'} \text{ and } Y'_{F'} \text{ are 0.9-spread.}$$

Moreover, there exists a set $R \subseteq F \setminus F'$ such that either $|R| = \sqrt{n}$ (“safe stage”) or $R = \emptyset$ (“unsafe stage”), and

1. *For every $x, y \in X' \times Y'$, and every $i \in R$, $\text{IP}_m(x_i, y_i) = 1$ (i.e., the output of every gadget inside R is fixed to 1);*
2. $\mathcal{D}(X' \times Y', F') \leq \mathcal{D}(X \times Y, F) + o(\sqrt{nm}) - \Omega(|F \setminus (F' \cup R)| \cdot m)$.

Here, the constants under $\Omega(\cdot)$ are universal.

To give some intuition behind the deficit inequality, it comes from

$$\mathcal{D}(X' \times Y', F') \leq \mathcal{D}(X \times Y, F) + O(|\Pi| + |R|) - \Omega(|F \setminus (F' \cup R)| \cdot m),$$

where Π refers to the protocol obtained from Yao’s principle, and the deficit increases in each stage by $O(|\Pi| + |R|)$ because the entropy is reduced by $O(1)$ for each bit of communication, as well as for each gadget in R that we fix to 1 (we lose $O(1)$ bits of entropy instead of $O(m)$ because we leave the $2m$ input bits random); meanwhile, the process will also fix the $2m$ input bits for each gadget in $F \setminus (F' \cup R)$ in an uncontrolled way, to “restore” spreadness and decreases the deficit.

The remainder of this paper is dedicated to proving the Stage Lemma. Assuming this lemma, here is the proof of the main theorem:

Proof of [Theorem 8](#) given [Lemma 13](#). As in the proof for decision trees, we may assume without loss of generality that $\Pr_{\mathbf{x}, \mathbf{y} \sim \Sigma^n \times \Sigma^n} [f(\mathbf{x}, \mathbf{y}) = 0] \geq \frac{1}{2}$, and we will apply the Stage Lemma iteratively over many stages,

Observe that the initial value $X \times Y = \Sigma^n \times \Sigma^n$ and $F = [n]$ satisfies the condition of [Lemma 13](#), and that $X' \times Y'$ and F' obtained from [Lemma 13](#) maintain the conditions to reapply the lemma, as long as $|F'| \geq 0.1n$. We may therefore perform the following procedure.

Initialize

- $X \times Y = \Sigma^n \times \Sigma^n$ (the current rectangle);
- $F = [n]$ (the current “free variables”);
- $T = \emptyset$ (the current set of gadgets fixed to 1);
- $U = \emptyset$ (the current set of “unfree variables” other than T).

While $|F| \geq 0.1n$, apply [Lemma 13](#) to obtain $X' \times Y' \subseteq X \times Y$, $F' \subsetneq F$, and $R \subseteq F \setminus F'$, and update

$$X \times Y \leftarrow X' \times Y', \quad U \leftarrow U \cup (F \setminus (F' \cup R)), \quad F \leftarrow F', \quad T \leftarrow T \cup R.$$




Figure 2: Communication stages. The second stage is *unsafe*, we assign too many coordinates

Since F shrinks in each iteration, the process halts. Upon halting, the final rectangle $X \times Y$ satisfies

$$\Pr[f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2 \quad (2)$$

for uniformly random $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$.

Note that the total number of stages is $O(\sqrt{n})$, since $|F \setminus F'| = \Omega(\sqrt{n})$, and so the number of “free variables” decreases by $\Omega(\sqrt{n})$ on every stage.

Now consider the deficit $\mathcal{D}(X \times Y, F)$ of the final rectangle with respect to the final set of free variables. By property (2) of [Lemma 13](#), this satisfies

$$0 \leq \mathcal{D}(X \times Y, F) \leq o(\sqrt{nm}) \cdot O(\sqrt{n}) - \Omega(|U|m) \leq o(nm) - \Omega(|U|m),$$

implying $|U| \leq o(n)$ for sufficiently large n .

When the process halts, the number of free variables is at most $|F| < 0.1n$, meaning that $|T| + |U| = n - |F| \geq 0.9n$, so

$$|T| \geq 0.9n - |U| \geq 0.8n.$$

For each $i \in T$, we have ensured in property (1) of [Lemma 13](#) that $\text{IP}_m(x_i, y_i) = 1$ for all $(x, y) \in X \times Y$. Since $|T| > \frac{2}{3}n$, this means that $f(x, y) = 1$ for all $(x, y) \in X \times Y$, contradicting [Equation \(2\)](#). This concludes the proof. $\square$

3.1 Proving the Stage Lemma

The structure of the proof of the Stage Lemma is shown in [Figure 3](#), which explains how the 4 main ingredients fit together: the Protocol Transformation lemma ([Lemma 15](#)), the communication version of the Closeness Lemma ([Lemma 18](#)), the Safe Stage Lemma ([Lemma 16](#)), and the Unsafe Stage Lemma ([Lemma 17](#)). We state these lemmas formally below. Let us begin with the definition of the sprinkled-1s distribution, similar to the one we used for decision trees:

Definition 14 (Sprinkled 1s Distribution). For $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$, let $\sigma(X \times Y, F)$ be the *sprinkled-1s distribution* obtained by

1. Choosing a uniformly random subset $\mathbf{I} \sim \binom{F}{\sqrt{n}}$ of free variables, and
2. Returning $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$ conditional on the event $\text{IP}_m(\mathbf{x}_i, \mathbf{y}_i) = 1$ for all $i \in \mathbf{I}$.

We prove the following Protocol Transformation lemma in [Section 4](#).

Lemma 15 (Protocol Transformation Lemma). *Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$, let $F \subseteq [n]$, and let Π be any deterministic protocol with cost $|\Pi| = o(\sqrt{nm})$ and error ε over an arbitrary distribution μ . Suppose that X_F and Y_F are both 0.9-spread. Then there exists a protocol Π' with error at most ε over the same distribution, such that each leaf $L = X^{(L)} \times Y^{(L)}$ of Π' is associated with a subset $F_L \subseteq F$ of free variables, where:*

1. *For every leaf L , the random variable $(\mathbf{x}^{(L)}, \mathbf{y}^{(L)}) \sim \text{unif}(X^{(L)} \times Y^{(L)})$ has $\mathbf{x}_{F_L}^{(L)}$ and $\mathbf{y}_{F_L}^{(L)}$ both 0.9-spread; and*
2. *For a random leaf $\mathbf{L} := L(\mathbf{x}, \mathbf{y})$ defined as the leaf of a random $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$, with probability at least $1 - 4\varepsilon$,*

$$\mathcal{D}(\mathbf{L}, F_{\mathbf{L}}) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_{\mathbf{L}}| \cdot m),$$

where the constants under $O(\cdot)$ and $\Omega(\cdot)$ are universal.




Figure 3: Outline of the proof of the Stage Lemma.

Lemma 16 (Safe Stage Lemma). *Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$ satisfy the conditions of Lemma 13, and assume the current stage is safe, i.e., Equation (3) holds. Then the conclusion of Lemma 13 holds with $|R| = \sqrt{n}$.*

Lemma 17 (Unsafe Stage Lemma). *Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$ satisfy the conditions of Lemma 13, and assume the current stage is unsafe, i.e., Equation (3) does not hold. Then the conclusion of Lemma 13 holds with $R = \emptyset$.*

The Stage Lemma (Lemma 13) follows immediately from Lemmas 16 and 17.

3.2 Proof of the Safe Stage Lemma

In this section we state the main ingredients required for the Safe Stage Lemma, including the communication version of the Closeness Lemma, and we prove the lemma using those ingredients. We defer the proofs of each ingredient to Section 5.

We start with the rectangle $X \times Y$ and free variables $F \subseteq [n]$. Recall that Π' is the transformed protocol from the Protocol Transformation Lemma, which has error ε over the distribution

$$\frac{1}{2}(\text{unif}(X \times Y) + \sigma(X \times Y, F)).$$

Suppose that the current stage is safe, i.e., Equation (3) holds. We want to find a subset $R \subseteq F$ of the free variables where we can fix all of the IP gadget values to 1, and produce a new rectangle $X' \times Y' \subseteq X \times Y$ and new free variables $F' \subseteq F \setminus R$ satisfying the conditions to repeat the process in Lemma 13.

The crucial step is the “birthday paradox” step, where we show that we can sprinkle $\sqrt{n}$ random 1-valued gadgets without the protocol Π' noticing. This is formalized in the next lemma, which uses spreadness of $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$ on the free variables F , together with the pseudorandomness of the INNER PRODUCT gadget under the spreadness condition (Lemma 12). We prove the Closeness Lemma in Section 5.1.

Lemma 18 (Closeness Lemma). *Assume $m \geq 100 \log n$. Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and let $F \subseteq [n]$ be such that the random variables $\mathbf{x}_F$ and $\mathbf{y}_F$ are both 0.9-spread, where $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. Now let $X' \times Y' \subseteq X \times Y$ and $F' \subseteq F$ also be such that $\mathbf{x}'_{F'}, \mathbf{y}'_{F'}$ are both 0.9-spread, for $(\mathbf{x}', \mathbf{y}') \sim \text{unif}(X' \times Y')$, and suppose $|F| \geq 0.1n$ and $|F| - |F'| \leq \sqrt{n}/1000$. Then*

$$\Pr_{(\mathbf{x}, \mathbf{y}) \sim \sigma(X \times Y, F)}[(\mathbf{x}, \mathbf{y}) \in X' \times Y'] \in (1 \pm 0.1) \Pr_{(\mathbf{x}, \mathbf{y}) \in \text{unif}(X \times Y)}[(\mathbf{x}, \mathbf{y}) \in X' \times Y'].$$

Informally, this lemma claims that if we start from a big enough spread rectangle in $\Sigma^n \times \Sigma^n$, a typical leaf of a shallow protocol will not be able to distinguish between uniform distribution and “sprinkled 1s” distribution.

With the Closeness Lemma, finding the target rectangle $X' \times Y'$ together with its free variables $F' \subseteq F$ and set $R \subseteq F \setminus F'$ of the 1-valued gadgets is done in three steps, given by the next three propositions. The first proposition claims that it is possible to find a 0-leaf in the protocol, such that it is mostly correct under the sprinkled distribution, and the deficiency would not grow too much.

Proposition 19 (Sprinkle 1s). *Let $\varepsilon > 0$ be a sufficiently small constant. Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$ satisfy the conditions of Lemma 13. If the current stage is safe, then there exists a leaf L of Π' such that:*

1. L contains mostly 0 inputs under the sprinkled-1s distribution:

$$\Pr_{(\mathbf{x}, \mathbf{y}) \sim \sigma(X \times Y, F)}[f(\mathbf{x}, \mathbf{y}) = 0 \mid (\mathbf{x}, \mathbf{y}) \in L] \geq 1 - 10\varepsilon,$$

2. $|F| - |F_L| \leq \sqrt{n}/1000$.

3. $\mathcal{D}(L, F_L) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|)$.

The next proposition states that we can fix a particular outcome for sprinkling it 1s; again, without increasing the deficiency too much.

Proposition 20 (Fix 1s). *Under the conditions of Proposition 19, let L be a leaf of Π' obtained from that proposition, and let F_L be its associated set of free variables (as defined in Lemma 15). Then there exists $X' \times Y' \subseteq L$, $R \subseteq F_L$ with $|R| = \sqrt{n}$, and $F' := F_L \setminus R$ such that*

1. $\Pr[f(\mathbf{x}, \mathbf{y}) = 0] \geq 0.6$ for $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X' \times Y')$;
2. For all $j \in R$ and all $x, y \in X' \times Y'$, $\text{IP}_m(x_j, y_j) = 1$, and
3. $\mathcal{D}(X' \times Y', F') \leq \mathcal{D}(L, F_L) + 2|R|$.

While the leaf L of Π' is spread (due to the protocol transformation), in the the previous step we moved to a subrectangle $X' \times Y' \subseteq L$ by fixing the R gadgets to 1, which may not have preserved spreadness. We regain spreadness by the same technique as in the Protocol Transformation Lemma.

Proposition 21 (Spreadify). *Let $A \times B \subseteq \Sigma^n \times \Sigma^n$ and $H \subseteq [n]$ such that $\Pr_{\mathbf{a}, \mathbf{b} \sim A \times B}[f(\mathbf{a}, \mathbf{b}) = 0] \geq 0.6$. Then there exists $H' \subseteq H$ and a rectangle $A' \times B' \subseteq A \times B$ such that*

1. $\Pr_{\mathbf{a}, \mathbf{b} \sim A' \times B'}[f(\mathbf{a}, \mathbf{b}) = 0] \geq 1/2$.
2. $A'_{H'}$ and $B'_{H'}$ are 0.9-spread.
3. $\mathcal{D}(A' \times B', H') \leq \mathcal{D}(A \times B, H) - \Omega(|H \setminus H'|m)$.

We prove these propositions in Section 5.2. With these tools, we can complete the proof of the safe stage lemma.

Proof of the Safe Stage Lemma. Start with $X \times Y$ and F satisfying the conditions of Lemma 13. Let L be the leaf of Π' with free variables F_L obtained from Proposition 19. Let $X' \times Y'$, F' , and $R = F \setminus F'$ be from Proposition 20 and let $X'' \times Y'' \subseteq X' \times Y'$ and $F'' \subseteq F'$ be obtained from applying Proposition 21 to $X' \times Y'$ and F' .

From Proposition 20 we guarantee that $\text{IP}_m(x_i, y_i) = 1$ for every $(x, y) \in X'' \times Y''$ and every $i \in R$. From Propositions 20 and 21 we get

$$\begin{aligned} \mathcal{D}(X'' \times Y'', F'') &\leq \mathcal{D}(X' \times Y', F') - \Omega(|F' \setminus F''|m) && \text{(Proposition 21)} \\ &\leq \mathcal{D}(L, F_L) + 2|R| - \Omega(|F' \setminus F''|m) && \text{(Proposition 20)} \\ &\leq \mathcal{D}(X \times Y, F) + O(|\Pi|) + 2|R| - \Omega(|F' \setminus F''|m) && \text{(Proposition 19)} \\ &\leq \mathcal{D}(X \times Y, F) + o(\sqrt{nm}) - \Omega(|F' \setminus F''|m) && (|\Pi|, |R| = o(\sqrt{nm})) \\ &\leq \mathcal{D}(X \times Y, F) + o(\sqrt{nm}) - \Omega(|F \setminus (F'' \cup R)|m) \end{aligned}$$

This concludes the proof of Lemma 16. □

3.3 Proof of the Unsafe Stage Lemma

A stage is *unsafe* if Equation (3) is false. In this case, we choose any 0-valued leaf L_0 with $|F| - |F_{L_0}|$ at least the median in Equation (3), and which also contains mostly 0-valued inputs of f ; we must show this exists.

Recall that $\mathbf{L}_0 = (\mathbf{L} \mid \mathbf{L} \text{ outputs } 0)$ is a random leaf drawn by taking the unique 0-leaf containing $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$ conditional on $(\mathbf{x}, \mathbf{y})$ ending up in a 0-valued leaf. Suppose Equation (3) is false. Let $R = \emptyset$. Then, to conclude the lemma, we want to show that there exists a 0-valued leaf L_0 of Π' with $|F| - |F_{L_0}| \geq \text{median}(|F| - |F_{L_0}|) > \sqrt{n}/1000$, such that

- $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(L_0)}[f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2$, and
- $\mathcal{D}(L_0, F_{L_0}) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_{L_0}| \cdot m)$.

Let $\mathbf{L}'_0$ be an independent copy of $\mathbf{L}_0$. We upper bound the probabilities for the following events:

- (E1) $|F| - |F_{\mathbf{L}'_0}| < \text{median}(|F| - |F_{\mathbf{L}_0}|)$;
- (E2) $\text{ERROR}_{\mathbf{L}'_0} > 1/2$, where $\text{ERROR}_{\mathbf{L}'_0} := \Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(\mathbf{L}'_0)}[f(\mathbf{x}, \mathbf{y}) = 1]$.
- (E3) $\mathcal{D}(\mathbf{L}'_0, F_{\mathbf{L}'_0}) > \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_{\mathbf{L}'_0}| \cdot m)$.

For (E1), by definition, $\Pr_{\mathbf{L}'_0}[|F| - |F_{\mathbf{L}'_0}| < \text{median}(|F| - |F_{\mathbf{L}_0}|)] < 1/2$.

For (E2), since Π' has error ε over the mixture $\frac{1}{2}(\text{unif}(X \times Y) + \sigma(X \times Y, F))$, it has error at most 2ε over $\text{unif}(X \times Y)$. Therefore $\mathbb{E}_{\mathbf{L}'_0}[\text{ERROR}_{\mathbf{L}'_0}] \leq 4\varepsilon$, so it follows by Markov inequality that $\Pr_{\mathbf{L}'_0}[\text{ERROR}_{\mathbf{L}'_0} \geq 20\varepsilon] \leq \frac{1}{5}$. For a small enough ε , we can conclude that:

$$\Pr_{\mathbf{L}'_0} \left[\text{ERROR}_{\mathbf{L}'_0} > \frac{1}{2} \right] = \Pr_{\mathbf{L}'_0} \left[\Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(\mathbf{L}'_0)}[f(\mathbf{x}, \mathbf{y}) = 0] < 1/2 \right] \leq \frac{1}{5}.$$

For (E3), we are guaranteed from Lemma 15 that

$$\mathcal{D}(\mathbf{L}, F_{\mathbf{L}}) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_{\mathbf{L}}| m) \quad (4)$$

with probability at least $1 - 4\varepsilon$ over a random leaf $\mathbf{L}$ chosen as the leaf containing $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. We want this to hold for the random 0-leaf $\mathbf{L}'_0$. Since $\Pr[f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2$ and the protocol has error at most 2ε over $(\mathbf{x}, \mathbf{y})$, $\mathbf{L}$ is a 0-leaf with probability at least $1/2 - 2\varepsilon$. Then Equation (4) must hold with probability at least $1 - 10\varepsilon$ for $\mathbf{L}'_0$ instead of $\mathbf{L}$, when ε is sufficiently small.

Applying the union bound, we can conclude that with positive probability, none of the three events occur. Then there exists a 0-leaf L_0 that satisfies the properties stated in the claim.

4 Density Restoring Partitions and Protocol Transformation

4.1 Density Restoring Partitions

The following is [GPW20, Lemma 3.5] (with straightforward adaptation to allow for $\mathbf{x}$ non-uniform, see also [CFK⁺21]):

Lemma 22 (Density Restoring Partition). *Let $\mathbf{x} \in \Sigma^n$ be a random variable with support X , let $F \subseteq [n]$, and let $\gamma \in (0, 1)$. Then there exists a partition $X = \bigsqcup_{j=1}^r X_j$ with associated sets $I_j \subseteq F$ and values $\alpha_j \in \Sigma^{I_j}$ such that:*

- $X_j := \{x \in X \mid x_{I_j} = \alpha_j\} \setminus \bigcup_{i < j} X_i$;
- $(\mathbf{x}_{F \setminus I_j} \mid \mathbf{x} \in X_j)$ is γ -spread;
- $D_\infty(\mathbf{x}_{F \setminus I_j} \mid \mathbf{x} \in X_j) \leq D_\infty(\mathbf{x}_F) - (1 - \gamma)|I_j|m + \delta_j$ where $\delta_j = \log \left(\frac{1}{\Pr[\mathbf{x} \in \bigcup_{k \geq j} X_k]} \right)$.

We define a procedure to apply the density restoring partition to rectangles $X \times Y$:

Lemma 23 (Density Restoring Partition for Rectangles). *Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$, $F \subseteq [n]$, $\varepsilon, \gamma \in (0, 1)$. Then there exist partitions $X = \bigsqcup_{i=1}^a X_i$ and $Y = \bigsqcup_{j=1}^b Y_j$ where each rectangle $X_i \times Y_j \subseteq X \times Y$ is associated with a set $F_{i,j} \subseteq F$ of “free variables” which satisfy the following properties:*

1. *For all i, j , the random variables $(\mathbf{x}_{F_{i,j}} \mid \mathbf{x} \in X_i)$ and $(\mathbf{y}_{F_{i,j}} \mid \mathbf{y} \in Y_j)$ are both γ -spread, where $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$;*
2. *With probability at least $1 - 2\varepsilon$ over $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$, the unique rectangle $X_i \times Y_j$ containing $(\mathbf{x}, \mathbf{y})$ satisfies*

$$\mathcal{D}(X_i \times Y_j, F_{i,j}) \leq \mathcal{D}(X \times Y, F) - (1 - \gamma)|F \setminus F_{i,j}|m + 2 \log(1/\varepsilon).$$

Proof. Begin by applying the density restoring partition of [Lemma 22](#) to each of X and Y to obtain

$$X = \bigsqcup_i X_i, \quad \text{and} \quad Y = \bigsqcup_j Y_j,$$

where for each i, j there exist sets $I_i, J_j \subseteq F$ and assignments $\alpha_i \in \Sigma^{I_i}$, $\beta_j \in \Sigma^{J_j}$ such that

$$X_i := \{x \in X \mid x_{I_i} = \alpha_i\} \setminus \bigcup_{k < i} X_k, \quad \text{and} \quad Y_j := \{y \in Y \mid y_{J_j} = \beta_j\} \setminus \bigcup_{k < j} Y_k,$$

where $(\mathbf{x}_{F \setminus I_i} \mid \mathbf{x} \in X_i)$ and $(\mathbf{y}_{F \setminus J_j} \mid \mathbf{y} \in Y_j)$ are both γ -spread for $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. We define the set $F_{i,j}$ associated with $X_i \times Y_j$ as

$$F_{i,j} := F \setminus (I_i \cup J_j),$$

so that the “free variables” are those which are not assigned by either one of the density restoring partitions. Since $F_{i,j} \subseteq F \setminus I_i$ and $F_{i,j} \subseteq F \setminus J_j$, the random variables

$$(\mathbf{x}_{F_{i,j}} \mid \mathbf{x} \in X_i) \quad \text{and} \quad (\mathbf{y}_{F_{i,j}} \mid \mathbf{y} \in Y_j)$$

remain γ -spread, establishing the first property of the lemma. Let us verify the second property. For every fixed rectangle $X_i \times Y_j$, we have from [Lemma 22](#) that

$$D_\infty(\mathbf{x}_{F \setminus I_i} \mid \mathbf{x} \in X_i) \leq D_\infty(\mathbf{x}_F) - (1 - \gamma)|I_i|m + \log \left(\frac{1}{\Pr[\mathbf{x}' \in \bigcup_{k \geq i} X_k]} \right),$$

where $\mathbf{x}' \sim \text{unif}(X)$. Since $F_{i,j} = F \setminus (I_i \cup J_j)$, we can use [Proposition 25](#) (stated below) to bound $D_\infty(\mathbf{x}_{F_{i,j}}) \leq D_\infty(\mathbf{x}_{F \setminus I_i})$. Then the remaining task to bound the final term.

Claim 24. *With probability at least $1 - \varepsilon$ over $\mathbf{x} \sim \text{unif}(X)$, the unique value $\mathbf{i} \in [a]$ such that $\mathbf{x} \in X_{\mathbf{i}}$ satisfies $\log \left(\frac{1}{\Pr[\mathbf{x}' \in \bigcup_{k \geq \mathbf{i}} X_k]} \right) < \log(1/\varepsilon)$.*

Proof of claim. We may think of the distribution over the index $\mathbf{i} \in [a]$ defined by $p(\mathbf{i}) := \Pr[\mathbf{i} = \mathbf{i}] = \Pr[\mathbf{x} \in X_{\mathbf{i}}]$. Then for all $t \in [0, 1]$,

$$\Pr_{\mathbf{x}} \left[\log \left(\frac{1}{\Pr[\mathbf{x}' \in \bigcup_{k \geq \mathbf{i}} X_k]} \right) > t \right] = \Pr_{\mathbf{i}} \left[\log \left(\frac{1}{\sum_{k \geq \mathbf{i}} p(k)} \right) > t \right] = \Pr_{\mathbf{i}} \left[\sum_{k \geq \mathbf{i}} p(k) < 2^{-t} \right].$$

Let i^* be the smallest value such that $\sum_{k \geq i^*} p(k) < 2^{-t}$. If $\sum_{k \geq \mathbf{i}} p(k) < 2^{-t}$ then $\mathbf{i} \geq i^*$ so the probability of this event is at most $\sum_{k \geq i^*} p(k) < 2^{-t}$. Setting $t = \log(1/\varepsilon)$ produces the desired bound. $\square$

Applying the same reasoning to $\mathbf{y}$ and using the union bound over $\mathbf{x}$ and $\mathbf{y}$, we have with probability at least $1 - 2\varepsilon$ the inequality

$$D_{\infty}(\mathbf{x}_{F \setminus J} \mid \mathbf{x} \in X_{\mathbf{i}}) + D_{\infty}(\mathbf{y}_{F \setminus J} \mid \mathbf{y} \in Y_{\mathbf{j}}) \leq D_{\infty}(\mathbf{x}_F) + D_{\infty}(\mathbf{y}_F) - (1 - \gamma)(|I_{\mathbf{i}}| + |J_{\mathbf{j}}|)m + 2\log(1/\varepsilon),$$

which concludes the proof since $|I_{\mathbf{i}}| + |J_{\mathbf{j}}| \geq |I_{\mathbf{i}} \cup J_{\mathbf{j}}|$. $\square$

We used the following proposition in the proof above:

Proposition 25. *Let $\mathbf{x}$ be a random variable over Σ^n and let $J \subseteq F \subseteq [n]$. Then*

$$D_{\infty}(\mathbf{x}_{F \setminus J}) \leq D_{\infty}(\mathbf{x}_F).$$

Proof. First, observe

$$\max_{z \in \Sigma^{F \setminus J}} \Pr[\mathbf{x}_{F \setminus J} = z] = \max_{z \in \Sigma^{F \setminus J}} \sum_{w \in \Sigma^F, w_J = z} \Pr[\mathbf{x}_F = w] \leq |\Sigma|^{|J|} \max_{w \in \Sigma^F} \Pr[\mathbf{x}_F = w],$$

so that

$$H_{\infty}(\mathbf{x}_{F \setminus J}) \geq H_{\infty}(\mathbf{x}_F) - |J|m.$$

Then

$$D_{\infty}(\mathbf{x}_{F \setminus J}) = (|F| - |J|)m - H_{\infty}(\mathbf{x}_{F \setminus J}) \leq |F|m - H_{\infty}(\mathbf{x}_F) = D_{\infty}(\mathbf{x}_F). \quad \square$$

4.2 Protocol Transformation

To prove the **Protocol Transformation Lemma**, we must relate the growth of the deficit function to the communication complexity of the protocol Π .

Proposition 26. *Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$, let $F \subseteq [n]$, and let Π be any communication protocol with cost d . Then for any $\varepsilon > 0$, with probability at least $1 - 2\varepsilon$ over a random leaf $\mathbf{L}$ of Π chosen as the leaf reached by $(\mathbf{x}', \mathbf{y}') \sim X \times Y$,*

$$\mathcal{D}(\mathbf{L}, F) \leq \mathcal{D}(X \times Y, F) + 2d + 2\log(1/\varepsilon).$$

To prove the statement, we will need a chain rule for min-entropy:

Lemma 27 (see e.g. [Vad12, Lemma 6.30]). Let $(\mathbf{a}, \mathbf{b})$ be distributed over $A \times B$ with $|A| = 2^\ell$ and $H_\infty(\mathbf{a}, \mathbf{b}) \geq k$, and let $\mathcal{A}$ denote the marginal distribution of $\mathbf{a}$. Then for every $\varepsilon > 0$

$$\Pr_{\mathbf{a}' \sim \mathcal{A}}[H_\infty(\mathbf{b} \mid \mathbf{a} = \mathbf{a}') \geq k - \ell - \log 1/\varepsilon] \geq 1 - \varepsilon.$$

Proof. Fix some $a \in A, b \in B$ then $\Pr[\mathbf{a} = a \wedge \mathbf{b} = b] / \Pr[\mathbf{a} = a] = \Pr[\mathbf{b} = b \mid \mathbf{a} = a]$, so $H_\infty(\mathbf{b} \mid \mathbf{a} = a) \geq H_\infty(\mathbf{a}, \mathbf{b}) - \log 1 / \Pr[\mathbf{a} = a]$. Then defining $p(a) := \Pr[\mathbf{a} = a]$ we have

$$\Pr[\log 1/p(\mathbf{a}) \geq \ell + \log 1/\varepsilon] = \Pr[p(\mathbf{a}) \leq \varepsilon/|A|] = \sum_{a: p(a) \leq \varepsilon/|A|} p(a) \leq \varepsilon. \quad \square$$

Proof of Proposition 26. The claim is equivalent to the statement that, with probability at least $1 - 2\varepsilon$ over a random leaf $\mathbf{L}$,

$$H_\infty(\mathbf{x}_F \mid (\mathbf{x}, \mathbf{y}) \in \mathbf{L}) + H_\infty(\mathbf{y}_F \mid (\mathbf{x}, \mathbf{y}) \in \mathbf{L}) \geq H_\infty(\mathbf{x}_F) + H_\infty(\mathbf{y}_F) - 2d - 2\log(1/\varepsilon).$$

Let $t(x, y)$ denote the transcript of protocol Π on input x, y , so that $t(x, y) \in \{0, 1\}^d$, and observe that transcript $t(x, y)$ is in one-to-one correspondence with the leaf L of (x, y) . Consider the random variable $(t(\mathbf{x}, \mathbf{y}), \mathbf{x}_F)$. By the chain rule for min-entropy (Lemma 27), with probability at least $1 - \varepsilon$ over a random transcript $\mathbf{t}' = t(\mathbf{x}', \mathbf{y}')$

$$H_\infty(\mathbf{x}_F \mid t(\mathbf{x}, \mathbf{y}) = \mathbf{t}') \geq H_\infty(\mathbf{x}_F) - d - \log(1/\varepsilon).$$

Using the same argument for $\mathbf{y}$ and applying the union bound, we get the desired conclusion. $\square$

We may now prove the protocol transformation lemma.

Proof of Protocol Transformation Lemma. Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$, let $F \subseteq [n]$, and let Π be a deterministic protocol with cost $o(\sqrt{nm})$ and error ε over μ . We define the protocol Π' in Algorithm 1, which we think of as outputting a transcript (i.e., a leaf L) along with a value.

Algorithm 1 The transformed protocol Π'

Input: $(x, y) \in X \times Y$

Output: Leaf $L'(x, y)$ and value $b = \Pi(x, y)$

- 1: Run protocol Π on (x, y) to obtain a leaf $L = L(x, y)$ and a value b .
- 2: Apply the density restoring partition for rectangles (Lemma 23) to $L = U^{(L)} \times V^{(L)}$ and F , to obtain

$$U^{(L)} = \bigsqcup_i U_i^{(L)}, \quad \text{and} \quad V^{(L)} = \bigsqcup_j V_j^{(L)}.$$

- 3: Alice sends value i for the unique $U_i^{(L)}$ containing x .
 - 4: Bob sends value j for the unique $V_j^{(L)}$ containing y .
 - 5: The players output leaf $U_i^{(L)} \times V_j^{(L)}$ and value b .
-

By definition, on any input the protocol outputs the same value as the original protocol Π , so its output is correct with the probability $1 - \varepsilon$ on μ .

For each leaf $L' = U_i^{(L)} \times V_j^{(L)}$ of Π' , we associate the free variables $F_{L'} = F_{i,j}^{(L)} \subseteq F$ given by **Lemma 23**. The first desired property of **Lemma 15** is that for every leaf $U_i^{(L)} \times V_j^{(L)}$ of Π' , the variables

$$(\mathbf{x}_{F_{i,j}^{(L)}} \mid (\mathbf{x}, \mathbf{y}) \sim U_i^{(L)} \times V_j^{(L)}) \quad \text{and} \quad (\mathbf{y}_{F_{i,j}^{(L)}} \mid (\mathbf{x}, \mathbf{y}) \sim U_i^{(L)} \times V_j^{(L)})$$

are 0.9-spread. This is immediately guaranteed by **Lemma 23**. The second desired property of **Lemma 15** is that a random leaf $\mathbf{L}' = U_i^{(\mathbf{L})} \times V_j^{(\mathbf{L})}$ satisfies

$$\mathcal{D}(\mathbf{L}', F_{\mathbf{L}'}) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_{\mathbf{L}'}|)m$$

with probability at least $1 - 4\varepsilon$, where constants under $O(\cdot)$ and $\Omega(\cdot)$ (C_1 and C_2 , respectively) are universal. The random leaf $\mathbf{L}'$ is obtained by first selecting a random leaf $\mathbf{L}$ of Π by taking the leaf containing $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$, and then conditional on $(\mathbf{x}, \mathbf{y}) \in \mathbf{L}$ taking the part $U_i^{(\mathbf{L})} \times V_j^{(\mathbf{L})}$ containing $(\mathbf{x}, \mathbf{y})$.

By **Proposition 26**, with probability at least $1 - 2\varepsilon$ over $\mathbf{L}$, we have

$$\mathcal{D}(\mathbf{L}, F) \leq \mathcal{D}(X \times Y, F) + 2|\Pi| + 2\log(1/\varepsilon).$$

Now for any fixed leaf L of Π , **Lemma 23** guarantees that with probability at least $1 - 2\varepsilon$ over $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$ conditional on $(\mathbf{x}, \mathbf{y}) \in L$, we have

$$\mathcal{D}(U_i^{(L)} \times V_j^{(L)}, F_{i,j}^{(L)}) \leq \mathcal{D}(L, F) - 0.1|F \setminus F_{i,j}^{(L)}|m + 2\log(1/\varepsilon).$$

Therefore, by the union bound, we may combine these inequalities to obtain, with probability at least $1 - 4\varepsilon$ over the random leaf $\mathbf{L}'$ of Π' ,

$$\mathcal{D}(\mathbf{L}', F_{\mathbf{L}'}) \leq \mathcal{D}(X \times Y, F) + 2|\Pi| - 0.1|F \setminus F_{\mathbf{L}'}|m + 4\log(1/\varepsilon).$$

For any fixed $\varepsilon > 0$ and sufficiently large n , $4\log(1/\varepsilon) < |\Pi|$, so we get the conclusion with $C_1 = 3$ and $C_2 = 0.1$. $\square$

5 Completing the Safe Stage Lemma

5.1 Closeness Lemma

In this section, we show that the sprinkled-1s distribution in **Definition 14** is close to uniform for the “good leaves”, i.e., the leaves L where we fix at most $\sqrt{n}/1000$ new variables ($|F| - |F_L| \leq \sqrt{n}/1000$). The intuition of this proof is that if the communication protocol “fixes” at most $\sqrt{n}/1000$ variables within one stage, then this is similar to “querying” at most $\sqrt{n}/1000$ gadgets, and therefore by a birthday paradox argument we may sprinkle a random set of $\sqrt{n}$ gadgets fixed to 1 without affecting the distribution over leaves of the protocol. We require the following birthday paradox bounds:

Claim 28. *Let $\mathbf{I} \sim \binom{[\ell]}{s}$ be a uniformly random set of s elements drawn from $[\ell]$ and let $T \subseteq [\ell]$ be such that $s \cdot |T| \leq \ell/100$. Then*

$$\Pr[\mathbf{I} \cap T = \emptyset] \geq 1 - 1/100 \quad \text{and} \quad \mathbb{E}[2^{|\mathbf{I} \cap T|}] \leq 1 + 1/20.$$

Proof. Write $\mathbf{I} = \{\mathbf{i}_1, \dots, \mathbf{i}_k\}$. For any fixed k we estimate

$$\Pr[|\mathbf{I} \cap T| \geq k] = \Pr[\exists K \subseteq [s] : |K| = k \wedge \forall j \in K, \mathbf{i}_j \in T] \leq \binom{s}{k} \left(\frac{|T|}{\ell}\right)^k \leq \frac{(s \cdot |T|)^k}{\ell^k} \leq 100^{-k}.$$

Then $\Pr[\mathbf{I} \cap T = \emptyset] \geq 1 - 1/100$ and

$$\mathbb{E}[2^{|\mathbf{I} \cap T|}] \leq 1 + \sum_{k=1}^{\ell} 2^k \Pr[|\mathbf{I} \cap T| \geq k] \leq 1 + \sum_{k=1}^{\ell} \frac{1}{50^k} \leq 1 + \frac{1}{20}. \quad \square$$

Lemma 18 (Closeness Lemma). *Assume $m \geq 100 \log n$. Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and let $F \subseteq [n]$ be such that the random variables $\mathbf{x}_F$ and $\mathbf{y}_F$ are both 0.9-spread, where $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. Now let $X' \times Y' \subseteq X \times Y$ and $F' \subseteq F$ also be such that $\mathbf{x}'_{F'}, \mathbf{y}'_{F'}$ are both 0.9-spread, for $(\mathbf{x}', \mathbf{y}') \sim \text{unif}(X' \times Y')$, and suppose $|F| \geq 0.1n$ and $|F| - |F'| \leq \sqrt{n}/1000$. Then*

$$\Pr_{(\mathbf{x}, \mathbf{y}) \sim \sigma(X \times Y, F)}[(\mathbf{x}, \mathbf{y}) \in X' \times Y'] \in (1 \pm 0.1) \Pr_{(\mathbf{x}, \mathbf{y}) \in \text{unif}(X \times Y)}[(\mathbf{x}, \mathbf{y}) \in X' \times Y'].$$

Proof. Within this proof we will use measure notation. For any subset $I \subseteq F$ of size $|I| = \sqrt{n}$, let ψ_I denote the measure on rectangles defined by our distributions $\sigma(X, Y, F)$ conditioned on choosing the set I of gadgets to fix to 1:

$$\forall A \times B \subseteq X \times Y: \quad \psi_I(A \times B) := \Pr_{(\mathbf{x}, \mathbf{y}) \sim \sigma(X \times Y, F)}[(\mathbf{x}, \mathbf{y}) \in A \times B \mid \mathbf{I} = I],$$

where $\mathbf{I} \sim \binom{F}{\sqrt{n}}$ denotes the random set of coordinates used to generate $(\mathbf{x}, \mathbf{y})$ in the distribution $\sigma(X \times Y, F)$.

According the definition of $\sigma(X, Y, F)$, the lemma's conclusion is equivalent to:

$$\mathbb{E}_{\mathbf{I} \sim \binom{F}{\sqrt{n}}}[\psi_I(X' \times Y')] \in \frac{|X' \times Y'|}{|X \times Y|}(1 \pm 0.1). \quad (5)$$

To prove the equation, we first observe that ψ_I is uniform over its support

$$\text{supp}(\psi_I) = \{(x, y) \in X \times Y \mid \forall i \in I, \text{IP}_m(x_i, y_i) = 1\},$$

so

$$\psi_I(X' \times Y') = \frac{|\text{supp}(\psi_I) \cap X' \times Y'|}{|\text{supp}(\psi_I)|}.$$

Next, we calculate a formula for $|\text{supp}(\psi_I) \cap A \times B|$ for any rectangle $A \times B \subseteq X \times Y$. For convenience, we define $\delta = 2^{-m/20} = n^{-5}$ for $m = 100 \log n$. Observe that $(1 - 2^{-m})\sqrt{n} \geq 1 - \sqrt{n} \cdot 2^{-m} \geq 1 - \delta$. Consider $(\mathbf{a}, \mathbf{b}) \sim \text{unif}(A \times B)$. Let's assume that $\mathbf{a}_S, \mathbf{b}_S$ are both 0.9-spread, for some set $S \subseteq F$. Then, due to the pseudorandomness lemma, [Lemma 12](#), the sequence of inner products $(\text{IP}_m(\mathbf{a}_i, \mathbf{b}_i))_{i \in S}$ is nearly uniform; specifically,

$$\forall z \in \{0, 1\}^S: \quad \Pr[(\text{IP}_m(\mathbf{a}_i, \mathbf{b}_i))_{i \in S} = z] \in \frac{1}{2^{|S|}}(1 \pm 2^{-m/20}) = \frac{1}{2^{|S|}}(1 \pm \delta).$$

Therefore,

$$\begin{aligned}
|\text{supp}(\psi_I) \cap A \times B| &= |A \times B| \cdot \Pr[(\mathbf{a}, \mathbf{b}) \in \text{supp}(\psi_I)] \\
&= |A \times B| \cdot \Pr[(\text{IP}_m(\mathbf{a}_i, \mathbf{b}_i))_{i \in I} = 1^I] \\
&\leq |A \times B| \cdot \Pr[(\text{IP}_m(\mathbf{a}_i, \mathbf{b}_i))_{i \in I \cap S} = 1^{I \cap S}] \\
&\leq |A \times B| \cdot 2^{-|I \cap S|} \cdot (1 + \delta).
\end{aligned} \tag{6}$$

In the case $I \subseteq S$, repeating the above calculation also yields

$$|\text{supp}(\psi_I) \cap A \times B| \in |A \times B| \cdot 2^{-|I|} \cdot (1 \pm \delta). \tag{7}$$

With these calculations, we may now complete the proof. For the rectangle $X \times Y$, $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$ have $\mathbf{x}_F, \mathbf{y}_F$ both 0.9-spread and that $I \subseteq F$ always, so by Equation (7), we have

$$|\text{supp}(\psi_I)| = |\text{supp}(\psi_I) \cap X \times Y| \in |X \times Y| \cdot 2^{-|I|} \cdot (1 \pm \delta),$$

using the fact that $\delta < 0.1$. Therefore, we can write the left-hand side of Equation (5) as

$$\begin{aligned}
\mathbb{E}_{\mathbf{I} \sim (\frac{F}{\sqrt{n}})}[\psi_I(X' \times Y')] &= \mathbb{E}_{\mathbf{I} \sim (\frac{F}{\sqrt{n}})} \left[\frac{|\text{supp}(\psi_I) \cap X' \times Y'|}{|\text{supp}(\psi_I)|} \right] \\
&\in \mathbb{E}_{\mathbf{I} \sim (\frac{F}{\sqrt{n}})}[|\text{supp}(\psi_I) \cap X' \times Y'|] \cdot \frac{2^{\sqrt{n}}}{|X \times Y|} \cdot (1 \pm \delta).
\end{aligned} \tag{8}$$

For the rectangle $X' \times Y'$, the random variables $\mathbf{x}'_{F'}, \mathbf{y}'_{F'}$ are both 0.9-spread for $(\mathbf{x}', \mathbf{y}') \sim \text{unif}(X' \times Y')$. In this case, we do not have $I \subseteq F'$ always, but we do have it with large enough probability, so we can write the following lower bound

$$\mathbb{E}_{\mathbf{I} \sim (\frac{F}{\sqrt{n}})}[|\text{supp}(\psi_I) \cap X' \times Y'|] \geq \Pr[\mathbf{I} \subseteq F'] \cdot \frac{|X' \times Y'|}{2^{\sqrt{n}}} \cdot (1 - \delta) \geq (1 - 1/100) \cdot \frac{|X' \times Y'|}{2^{\sqrt{n}}} \cdot (1 - \delta),$$

where we use the fact that $|\mathbf{I}| \cdot |F \setminus F'| \leq n/1000 \leq |F|/100$, together with Claim 28, to obtain

$$\Pr[\mathbf{I} \subseteq F'] = \Pr[\mathbf{I} \cap (F \setminus F') = \emptyset] \geq 1 - \frac{1}{100}.$$

Using Equation (6) and Claim 28, we also get an upper bound

$$\mathbb{E}_{\mathbf{I} \sim (\frac{F}{\sqrt{n}})}[|\text{supp}(\psi_I) \cap X' \times Y'|] \leq \frac{|X' \times Y'|}{2^{\sqrt{n}}} \cdot (1 + \delta) \cdot \mathbb{E}[2^{|\mathbf{I} \setminus F'|}] \leq \frac{|X' \times Y'|}{2^{\sqrt{n}}} \cdot (1 + \delta) \cdot (1 + \frac{1}{20}).$$

By combining these bounds with Equation (8), we establish Equation (5). $\square$

5.2 Sprinkling and Fixing the 1s

Throughout this section, we are in the context of the Safe Stage Lemma, meaning that we have rectangle $X \times Y$ and free variables $F \subseteq [n]$ satisfying the conditions of the Stage Lemma; Π is the protocol assumed in Assumption 9, Π' is the transformed protocol from the Protocol Transformation Lemma, and Equation (3) holds.

5.2.1 Finding a Leaf

Proposition 19 (Sprinkle 1s). *Let $\varepsilon > 0$ be a sufficiently small constant. Let $X \times Y \subseteq \Sigma^n \times \Sigma^n$ and $F \subseteq [n]$ satisfy the conditions of [Lemma 13](#). If the current stage is safe, then there exists a leaf L of Π' such that:*

1. L contains mostly 0 inputs under the sprinkled-1s distribution:

$$\Pr_{(\mathbf{x}, \mathbf{y}) \sim \sigma(X \times Y, F)}[f(\mathbf{x}, \mathbf{y}) = 0 \mid (\mathbf{x}, \mathbf{y}) \in L] \geq 1 - 10\varepsilon,$$

2. $|F| - |F_L| \leq \sqrt{n}/1000$.
3. $\mathcal{D}(L, F_L) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|)$.

Proof. We say a leaf L of the protocol Π' is *good* if it satisfies the following three conditions:

1. F is a 0-leaf of Π' (i.e., the protocol outputs 0 if it reaches L);
2. $|F| - |F_L| \leq \sqrt{n}/1000$;
3. $\mathcal{D}(L, F_L) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_L|m)$, from [Lemma 15](#).

Let $L_1, \dots, L_N$ be all the good leaves of Π' , with associated variables F_{L_i} . First, we lower bound the probability to be in a good leaf, and then prove that a typical good leaf does not contain too many errors, which yields the proposition.

Let $\mathbf{L}$ be a random leaf of Π' , chosen as the unique leaf containing $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)$. The current stage is safe, meaning that [Equation \(3\)](#) holds, so $\Pr[|F| - |F_{\mathbf{L}}| \leq \sqrt{n}/1000 \mid \mathbf{L} \text{ prints } 0] \geq 1/2$. By [Lemma 15](#),

$$\Pr[\mathcal{D}(\mathbf{L}, F_{\mathbf{L}}) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F \setminus F_{\mathbf{L}}|m)] \geq 1 - 4\varepsilon.$$

Since Π' errs with probability at most ε over the mixture distribution $\frac{1}{2}(\text{unif}(X \times Y) + \sigma(X \times Y, F))$, it errs with probability at most 2ε over $\text{unif}(X \times Y)$. By assumption, $\Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)}[f(\mathbf{x}, \mathbf{y}) = 0] \geq 1/2$, so by the union bound,

$$\Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)}[L(\mathbf{x}, \mathbf{y}) \text{ is good}] \geq \Pr[|F| - |F_{L(\mathbf{x}, \mathbf{y})}| \leq \sqrt{n}/1000 \text{ and } \Pi'(\mathbf{x}, \mathbf{y}) = 0] - 4\varepsilon \geq 1/4 - 5\varepsilon.$$

We rewrite the inequality as

$$\sum_{i \in [N]} \Pr_{(\mathbf{x}, \mathbf{y}) \sim X \times Y}[(\mathbf{x}, \mathbf{y}) \in L_i] \geq 1/4 - 5\varepsilon. \quad (9)$$

For brevity, we write $\sigma := \sigma(X \times Y, F)$. By the [Closeness Lemma](#), for every $i \in [N]$,

$$\Pr_{(\mathbf{u}, \mathbf{v}) \sim \sigma}[(\mathbf{u}, \mathbf{v}) \in L_i] \in (1 \pm 1/10) \Pr_{(\mathbf{x}, \mathbf{y}) \sim X \times Y}[(\mathbf{x}, \mathbf{y}) \in L_i]. \quad (10)$$

Let $p_L := \Pr_{(\mathbf{u}, \mathbf{v}) \sim \sigma}[f(\mathbf{u}, \mathbf{v}) = 1 \mid (\mathbf{u}, \mathbf{v}) \in L]$. Since Π' has error ε over the mixture, it has error 2ε

over σ , so

$$\begin{aligned}
2\varepsilon &\geq \Pr_{(\mathbf{u}, \mathbf{v}) \sim \sigma} [f(\mathbf{u}, \mathbf{v}) = 1 \wedge \Pi'(\mathbf{u}, \mathbf{v}) = 0] \\
&\geq \sum_{L \text{ 0-leaf of } \Pi'} p_L \Pr_{(\mathbf{u}, \mathbf{v}) \sim \sigma} [(\mathbf{u}, \mathbf{v}) \in L] \\
&\geq \sum_{i \in [N]} p_{L_i} \Pr_{(\mathbf{u}, \mathbf{v}) \sim \sigma} [(\mathbf{u}, \mathbf{v}) \in L_i] \\
&\geq 0.9 \sum_{i \in [N]} p_{L_i} \Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)} [(\mathbf{x}, \mathbf{y}) \in L_i] \quad (\text{by Equation (10).})
\end{aligned}$$

Thus, for a random L_i chosen as the unique leaf containing $(\mathbf{x}', \mathbf{y}')$ drawn from the uniform distribution on $\bigcup_i L_i$,

$$\begin{aligned}
\mathbb{E}[p_{L_i}] &= \sum_{i \in [N]} p_{L_i} \Pr_{(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X \times Y)} [(\mathbf{x}, \mathbf{y}) \in L_i \mid (\mathbf{x}, \mathbf{y}) \in \bigcup_i L_i] \\
&\leq 2\varepsilon \cdot \frac{10}{9} \cdot \Pr \left[(\mathbf{x}, \mathbf{y}) \in \bigcup_i L_i \right]^{-1} \\
&\leq 2\varepsilon \cdot \frac{10}{9} \cdot \frac{1}{1/4 - 5\varepsilon} \quad (\text{by Equation (9)}) \\
&\leq 10\varepsilon,
\end{aligned}$$

for sufficiently small $\varepsilon > 0$. Therefore, there exists a good leaf L_i with $p_{L_i} \leq 10\varepsilon$. For this leaf,

$$\mathcal{D}(L_i, F_{L_i}) \leq \mathcal{D}(X \times Y, F) + O(|\Pi|) - \Omega(|F| - |F_{L_i}|)m \leq \mathcal{D}(X \times Y, F) + O(|\Pi|),$$

as desired. $\square$

5.2.2 Fixing the Gadgets

Proposition 20 (Fix 1s). *Under the conditions of Proposition 19, let L be a leaf of Π' obtained from that proposition, and let F_L be its associated set of free variables (as defined in Lemma 15). Then there exists $X' \times Y' \subseteq L$, $R \subseteq F_L$ with $|R| = \sqrt{n}$, and $F' := F_L \setminus R$ such that*

1. $\Pr[f(\mathbf{x}, \mathbf{y}) = 0] \geq 0.6$ for $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(X' \times Y')$;
2. For all $j \in R$ and all $x, y \in X' \times Y'$, $\text{IP}_m(x_j, y_j) = 1$, and
3. $\mathcal{D}(X' \times Y', F') \leq \mathcal{D}(L, F_L) + 2|R|$.

Proof. Let $\sigma := \sigma(X \times Y, F)$. By Proposition 19 there exists a leaf $L = X^{(L)} \times Y^{(L)}$ of Π' satisfying $|F| - |F_L| \leq \sqrt{n}/1000$ and

$$\Pr_{(\mathbf{u}, \mathbf{v}) \sim \sigma} [f(\mathbf{u}, \mathbf{v}) = 1 \mid (\mathbf{u}, \mathbf{v}) \in L] \leq 1/5.$$

Recall from Definition 14 that $(\mathbf{u}, \mathbf{v}) \sim \sigma$ is obtained by choosing $\mathbf{I} \sim \binom{F}{\sqrt{n}}$ and taking $(\mathbf{u}, \mathbf{v})$ uniform over the set $\{(u, v) \in X \times Y \mid \text{IP}_m(u_i, v_i) = 1, \forall i \in \mathbf{I}\}$. So there exists a choice $\mathbf{I} = R \in \binom{F}{\sqrt{n}}$ such that

$$\Pr[f(\mathbf{u}, \mathbf{v}) = 1 \mid (\mathbf{u}, \mathbf{v}) \in L \wedge \mathbf{I} = R] \leq 1/5.$$

Let $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(L)$ and let $(\mathbf{u}', \mathbf{v}')$ be distributed identically to $(\mathbf{u}, \mathbf{v} \mid (\mathbf{u}, \mathbf{v}) \in L \wedge \mathbf{I} = R)$, which is equivalent to $(\mathbf{u}', \mathbf{v}')$ distributed as $(\mathbf{x}, \mathbf{y} \mid \text{IP}_m(x_i, y_i) = 1, \forall i \in R)$. In calculations below, we write $F := F_L$ for brevity. We first want to show

$$\text{H}_\infty(\mathbf{u}'_F, \mathbf{v}'_F) \geq \text{H}_\infty(\mathbf{x}_F, \mathbf{y}_F) - |R| - 1. \quad (11)$$

From [Lemma 12](#),

$$\Pr[\forall i \in R : \text{IP}_m(\mathbf{x}_i, \mathbf{y}_i) = 1] \in 2^{-|R|}(1 \pm 2^{-m/20}),$$

so

$$\max_{(u,v) \in \Sigma^F \times \Sigma^F} \Pr[(\mathbf{u}'_F, \mathbf{v}'_F) = (u, v)] \leq \max_{(x,y) \in \Sigma^F \times \Sigma^F} \Pr[(\mathbf{x}_F, \mathbf{y}_F) = (x, y)] \cdot 2^{|R|}(1 - 2^{-m/20})^{-1},$$

which establishes [Equation \(11\)](#). The set $\{x, y \mid \text{IP}_m(x_i, y_i) = 1, \forall i \in R\}$ is not a rectangle; to obtain a rectangle we will find variable settings $x_R^0, y_R^0 \in \Sigma^R$ so that the rectangle $X' \times Y' := \{(x, y) \in L \mid (x_R, y_R) = (x_R^0, y_R^0)\}$ satisfies the desired condition. Let $(\mathbf{x}_R^0, \mathbf{y}_R^0)$ be distributed identically to $(\mathbf{u}'_R, \mathbf{v}'_R)$. Then by the chain rule for min-entropy ([Lemma 27](#)), for all $\delta \in (0, 1)$,

$$\Pr_{(\mathbf{x}_R^0, \mathbf{y}_R^0)} \left[\text{H}_\infty(\mathbf{u}'_{F \setminus R}, \mathbf{v}'_{F \setminus R} \mid (\mathbf{u}'_R, \mathbf{v}'_R) = (\mathbf{x}_R^0, \mathbf{y}_R^0)) \geq \text{H}_\infty(\mathbf{u}'_F, \mathbf{v}'_F) - 2|R|m - \log(1/\delta) \right] \geq 1 - \delta.$$

The last step is to ensure that the rectangle will be mostly 0-valued.

$$\mathbb{E}_{\mathbf{x}_R^0, \mathbf{y}_R^0} \left[\Pr_{\mathbf{u}', \mathbf{v}'} [f(\mathbf{u}', \mathbf{v}') = 1 \mid (\mathbf{u}'_R, \mathbf{v}'_R) = (\mathbf{x}_R^0, \mathbf{y}_R^0)] \right] = \Pr[f(\mathbf{u}', \mathbf{v}') = 1] \leq 1/5,$$

so by Markov's inequality we have

$$\Pr_{\mathbf{x}_R^0, \mathbf{y}_R^0} [\Pr[f(\mathbf{u}', \mathbf{v}') = 1 \mid (\mathbf{u}'_R, \mathbf{v}'_R) = (\mathbf{x}_R^0, \mathbf{y}_R^0)] > 0.4] \leq \frac{1}{5} \cdot \frac{10}{4} = 1/2.$$

By the union bound (using $\delta = 1/4$), there exist x_R^0, y_R^0 such that $\Pr[f(\mathbf{u}', \mathbf{v}') = 0 \mid (\mathbf{u}'_R, \mathbf{v}'_R) = (x_R^0, y_R^0)] \geq 6/10$ and

$$\text{H}_\infty(\mathbf{x}'_{F \setminus R}, \mathbf{v}'_{F \setminus R} \mid (\mathbf{u}'_R, \mathbf{v}'_R) = (x_R^0, y_R^0)) \geq \text{H}_\infty(\mathbf{u}'_F, \mathbf{v}'_F) - 2|R|m - 2 \geq \text{H}_\infty(\mathbf{x}_F, \mathbf{y}_F) - 2|R|m - |R| - 3.$$

Therefore the rectangle $X' \times Y'$ for $X' := \{x \in X^{(L)} \mid x_R = x_R^0\}$ and $Y' := \{y \in Y^{(L)} \mid y_R = y_R^0\}$ satisfies

$$\begin{aligned} \mathcal{D}(X' \times Y', F \setminus R) &= 2(|F| - |R|)m - \text{H}_\infty(\mathbf{u}'_{F \setminus R} \mid \mathbf{u}'_R = x_R^0) - \text{H}_\infty(\mathbf{v}'_{F \setminus R} \mid \mathbf{v}'_R = y_R^0) \\ &\leq 2(|F| - |R|)m - \text{H}_\infty(\mathbf{x}_F, \mathbf{y}_F) + 2|R|m + |R| + 3 \\ &\leq 2|F|m - \text{H}_\infty(\mathbf{x}_F, \mathbf{y}_F) + 2|R| \\ &= \mathcal{D}(L, F) + 2|R|. \square \end{aligned}$$

5.2.3 Restoring Spreadness

Proposition 21 (Spreadify). *Let $A \times B \subseteq \Sigma^n \times \Sigma^n$ and $H \subseteq [n]$ such that $\Pr_{\mathbf{a}, \mathbf{b} \sim A \times B} [f(\mathbf{a}, \mathbf{b}) = 0] \geq 0.6$. Then there exists $H' \subseteq H$ and a rectangle $A' \times B' \subseteq A \times B$ such that*

1. $\Pr_{\mathbf{a}, \mathbf{b} \sim A' \times B'}[f(\mathbf{a}, \mathbf{b}) = 0] \geq 1/2$.
2. $A'_{H'}$ and $B'_{H'}$ are 0.9-spread.
3. $\mathcal{D}(A' \times B', H') \leq \mathcal{D}(A \times B, H) - \Omega(|H \setminus H'|m)$.

Proof. For parameter $\varepsilon > 0$ to be fixed later, we apply the density restoring partition for the rectangle $A \times B$ (Lemma 23) to obtain partitions $A = \bigsqcup_i A_i$ and $B = \bigsqcup_j B_j$ with each rectangle $A_i \times B_j$ associated with a set $H_{i,j} \subseteq H$ of free variables. For each i, j , the random variables $\mathbf{x}_{H_{i,j}}$ and $\mathbf{y}_{H_{i,j}}$ are 0.9-spread, where $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(A_i \times B_j)$. With probability at least $1 - 2\varepsilon$ over $(\mathbf{x}, \mathbf{y}) \sim \text{unif}(A \times B)$, the unique rectangle $A_i \times B_j$ containing $(\mathbf{x}, \mathbf{y})$ satisfies

$$\mathcal{D}(A_i \times B_j, H_{i,j}) \leq \mathcal{D}(A \times B, H) - 0.1|H \setminus H_{i,j}|m + 2\log(1/\varepsilon). \quad (12)$$

For each i, j , define

$$p_{i,j} := \Pr_{(\mathbf{a}', \mathbf{b}') \sim \text{unif}(A_i \times B_j)}[f(\mathbf{a}', \mathbf{b}') = 1].$$

Let (i, j) the random variable chosen as the unique values such that $A_i \times B_j$ contains $(\mathbf{a}, \mathbf{b}) \sim \text{unif}(A \times B)$. By assumption,

$$\mathbb{E}[p_{i,j}] = \sum_{i,j} \Pr[(\mathbf{x}, \mathbf{y}) \in A_i \times B_j] \cdot p_{i,j} = \Pr_{(\mathbf{a}, \mathbf{b}) \sim \text{unif}(A \times B)}[f(\mathbf{a}, \mathbf{b}) = 1] \leq 4/10,$$

so by Markov's inequality, $\Pr[p_{i,j} > 1/2] \leq 8/10$. Setting $\varepsilon = 1/20$, the probability that $p_{i,j} > 1/2$ or that Equation (12) fails for i, j is at most $8/10 + 1/10 = 9/10$, so there exists rectangle $A' \times B' = A_i \times B_j$ with free variables $H' := H_{i,j}$ satisfying the desired conditions. $\square$

A Appendix: Counting Arguments and Variations of FBPP

A.1 Variations of FBPP

As mentioned in the introduction, there is a variety of natural ways to define a class FBPP of “BPP-search problems”, i.e., search problems with efficient randomized protocols. Here are four definitions of increasing restrictiveness:

- (i) Direct translation of the definition of BPP for decision problems. A sequence of relations $R_n \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ is in FBPP if and only if $\forall n$ there exists a randomized protocol Π with cost $\text{poly} \log n$ such that $\forall x, y \in \{0, 1\}^n$, $\Pr[(x, y, \Pi(x, y)) \in R_n] \geq 2/3$.

The issue with this definition is that, unlike for BPP, the constant $2/3$ is no longer arbitrary: choosing a different constant will change the class, because for search problems, the error cannot generally be boosted (see examples below).

- (ii) Fix the issue by demanding that there is an efficient protocol achieving any error $\varepsilon > 0$. The sequence $R_n \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ is in FBPP if and only if $\forall n, \forall \varepsilon > 0$ there exists a randomized protocol with cost $\text{poly}(\log n, 1/\varepsilon)$ such that $\forall x, y$, $\Pr[(x, y, \Pi(x, y)) \in R_n] \geq 1 - \varepsilon$. This is the approach taken in [Aar10, ABK24] for Turing machines.
- (iii) Demand that the cost of the protocol is $\text{poly}(\log n, \log(1/\varepsilon))$ instead of $\text{poly}(\log n, 1/\varepsilon)$, so that dependence on ε is the same as it is for decision problems.

(iv) Require that solutions are not only efficient to find, but also efficient to verify. This is the approach taken by [Gol11], and this is also analogous to the definition of efficient deterministic search problems¹ in the context of TFNP. According to this definition, R_n is in FBPP if and only if both the following conditions hold:

- The condition from definition (i) above, where the relation can be computed with cost $\text{poly log } n$ and success probability is $2/3$; and,
- For every n there exists a randomized *verifier* protocol V which outputs a value in $\{0, 1\}$ such that, for all $(x, y, z) \in \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$, when Alice is given (x, z) and Bob is given (y, z) , they can run protocol V which satisfies $\Pr[V(x, y, z) = 1_{(x, y, z) \in R_n}] \geq 2/3$.

Note that each of these definitions is more strict than the previous one: If a relation satisfies definition (ii) then it also satisfies definition (i), etc.

Pseudodeterminism. It is most natural to compare pseudodeterminism with definition (iii). This is because, if a problem admits an efficient $\text{poly log } n$ pseudodeterministic protocol, it also satisfies definition (iii). With a pseudodeterministic protocol, we can take a majority vote on the output of $O(\log(1/\varepsilon))$ independent runs; with probability $1 - \varepsilon$, the canonical output for given input (x, y) will be the majority of outputs.

Partial boolean functions. Each of the above definitions is *more strict* than the previous one, but this is only true for problems with large outputs. When we consider partial boolean functions, i.e., problems where the valid outputs are single bits (the main topic of this paper), these 4 definitions collapse into only 2, because the first 3 definitions are equivalent. Given a protocol with the guarantees in definition (i), we can perform majority-vote error boosting. If there is only 1 valid output for given (x, y) then the majority vote will take this value; if there are 2 valid outputs for given (x, y) then the output is valid with probability 1.

If we then take definition (iii) as the definition of FBPP for partial functions, our Theorem 1 proves $\text{FBPP} \subsetneq \text{FPs}$ for partial functions.

Furthermore, when the number of valid outputs is small, definition (iv) is a special case of efficient pseudodeterminism. This is because, for relations $R_n \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ with only a small $\leq \text{poly log}(n)$ number of possible outputs, the verifier can be used to create a pseudodeterministic protocol, by iterating over every possible output in a fixed order:

Claim 29. *Let $R_n \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ be a relation that satisfies definition (iv) above, and also $|\{z \in \{0, 1\}^* \mid \exists x, y : (x, y, z) \in R_n\}| \leq \text{poly log } n$. Then there is a pseudodeterministic protocol for R_n with cost $\text{poly log } n$.*

A.2 Separations via counting

Let us explain how to prove Equation (1) for definitions (i)–(iii) of FBPP given above, using a counting argument that works equally well for boolean circuits and other models of computation. An initial sketch of the argument for definition (i) is as follows:

¹We remark that, confusingly, the name FP is used for two different classes of problems: relations where a solution can be found efficiently, and relations where a solution can be found *and* verified efficiently.

Define a relation $S \subseteq \{0, 1\}^n \times \{0, 1\}^n \times [3]$ by choosing $S(x, y) := \{i : (x, y, i) \in S\}$ as a random subset of $[3]$ of size 2 (i.e., we forbid one random output element). On any input, outputting a random number in $[3]$ will solve S with probability $2/3$. A simple counting argument shows that the pseudodeterministic complexity is $\Omega(n)$ with high probability over the choice of S .

For definition (i), we can make this argument explicit (see [Appendix A.3](#)), i.e., we exhibit a specific relation satisfying definition (i) but which does not have an efficient pseudodeterministic protocol. But let us now give the full counting argument for definition (iii).

Proposition 30. *There exists a sequence $R_n \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^*$ of relations such that for every $\varepsilon > 0$, $R_\varepsilon(R_n) = O(\log(1/\varepsilon))$, but which has pseudodeterministic cost $\Omega(n)$.*

Proof. Fix any n . Let $\mathcal{R}$ be the set of all relations R such that $|R(x, y) \cap \{0, 1\}^k| = 2^k - 1$ for all $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$ and $k \in [n]$. That is, in any relation $R \in \mathcal{R}$ we have removed exactly one of the 2^k valid outputs for every output length $k \in [n]$. Each relation $R \in \mathcal{R}$ has randomized cost $O(\log(1/\varepsilon))$ since for every ε we can take $k = \lceil \log(1/\varepsilon) \rceil$ and output a random value $z \sim \{0, 1\}^k$; then $(x, y, z) \in R$ with probability at least $1 - 2^{-k} \geq 1 - \varepsilon$.

Fix any function $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^*$ such that the length of any output is at most $|f(x, y)| \leq c$ for some c . Let $\mathcal{R}_f \subseteq \mathcal{R}$ be the set of relations consistent with f (i.e., $f(x, y) \in R$ for all x, y). Then

$$\frac{|\mathcal{R}_f|}{|\mathcal{R}|} = \prod_{x, y} \left(1 - 2^{-|f(x, y)|}\right) \leq (1 - 2^{-c})^{2^{2n}} \leq e^{-2^{2n-c}}.$$

Every pseudodeterministic protocol with cost c computes one of these functions f . Using Newman's theorem to bound the number of random bits in any randomized protocol by $\log n + O(1)$, the number of pseudodeterministic protocols with cost c is at most

$$\left((2^{2n})^{2^c} \cdot (2^{c+1})^{2^c} \right)^{O(n)}.$$

Then the fraction of relations consistent with any pseudodeterministic protocol is at most

$$2^{O(n(2^{n+c} + c2^c)) - \Omega(2^{2n-c})} < 1$$

when $c = o(n)$. □

A.3 Explicit Separation for Definition (i)

Theorem 31. *There exists an explicit two-party search problem $f \subseteq \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^2$ with randomized communication complexity $O(1)$ but pseudodeterministic complexity $\Omega(n)$.*

We first claim that it suffices to find an f and two different constant error parameters $0 < \varepsilon_2 < \varepsilon_1 < \frac{1}{3}$, such that $R_{\varepsilon_1}(f) = O(1)$ while $R_{\varepsilon_2}(f) = \Omega(n)$ where we denote by R_ε the ε -error communication complexity and we omit ε when $\varepsilon = 1/3$. Let us also denote by Ps_ε the ε -error pseudodeterministic complexity. Indeed, if we can find such an f , we have $R(f) \leq R_{\varepsilon_1}(f) = O(1)$, and on the other hand,

$$Ps(f) \geq \Omega(Ps_{\varepsilon_2}(f)/\log(1/\varepsilon_2)) \geq \Omega(R_{\varepsilon_2}(f)/\log(1/\varepsilon_2)) = \Omega(n),$$

where we used that pseudodeterministic complexity admits error reduction.

We now construct f as follows: Let $n = 2m$ without loss of generality. Given $x, y \in \{0, 1\}^n$ as the input for Alice and Bob respectively, we interpret $x = (x^1, x^2)$ as the concatenation of two m -bit strings $x^1, x^2 \in \{0, 1\}^m$, and similarly for $y = (y^1, y^2)$. Then we define $f(x, y) := \{0, 1\}^2 \setminus \{(\text{IP}(x^1, y^1), \text{IP}(x^2, y^2))\}$ as the set of all pairs of bits excluding $(\text{IP}(x^1, y^1), \text{IP}(x^2, y^2))$. We prove that $R_{1/4}(f) = O(1)$, while $R_{0.1}(f) = \Omega(n)$.

We first show $R_{1/4}(f) = O(1)$: Observe that $|f(x, y)| = 3$ for all $x, y \in \{0, 1\}^n$, the simple protocol that outputs a uniform random pair $(\mathbf{u}, \mathbf{v}) \sim \{0, 1\}^2$ succeeds with probability $3/4$.

It remains to show $R_{0.1}(f) = \Omega(n)$. For the sake of contradiction, suppose that $R_{0.1}(f) = o(n)$. By Yao's principle, there exists a deterministic protocol Π of cost $d = o(n)$ that computes f with error at most 0.1 with respect to $\text{unif}(\{0, 1\}^n \times \{0, 1\}^n)$. We will then construct a deterministic protocol Π' of cost $d' = O(d) = o(n)$ that computes n -bit IP with error at most 0.4 with respect to $\text{unif}(\{0, 1\}^m \times \{0, 1\}^m)$, a contradiction to the following folklore lower bound for distributional communication complexity of IP.

Lemma 32 ([CG88]). *The distributional communication complexity of the n -bit IP with respect to $\mu := \text{unif}(\{0, 1\}^{n/2} \times \{0, 1\}^{n/2})$ is $D_{0.4}(\text{IP}_{n/2}, \mu) = \Omega(n)$.*

Before presenting the protocol Π' , we define the following notation:

$$\begin{aligned} p_{00} &:= \Pr[\mathbf{u} \neq \text{IP}(\mathbf{x}^1, \mathbf{y}^1) \wedge \mathbf{v} \neq \text{IP}(\mathbf{x}^2, \mathbf{y}^2)], \\ p_{01} &:= \Pr[\mathbf{u} \neq \text{IP}(\mathbf{x}^1, \mathbf{y}^1) \wedge \mathbf{v} = \text{IP}(\mathbf{x}^2, \mathbf{y}^2)], \\ p_{10} &:= \Pr[\mathbf{u} = \text{IP}(\mathbf{x}^1, \mathbf{y}^1) \wedge \mathbf{v} \neq \text{IP}(\mathbf{x}^2, \mathbf{y}^2)], \end{aligned}$$

where $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2), \mathbf{y} = (\mathbf{y}^1, \mathbf{y}^2)$ are uniform n -bit strings and $(\mathbf{u}, \mathbf{v}) := \Pi(\mathbf{x}, \mathbf{y})$. It follows that $p_{00} + p_{01} + p_{10} = \Pr[(\mathbf{u}, \mathbf{v}) \in f(\mathbf{x}, \mathbf{y})] \geq 0.9$.

Next, we specify the protocol Π' as follows: Alice holds $\mathbf{x}' \sim \{0, 1\}^m$ and Bob holds $\mathbf{y}' \sim \{0, 1\}^m$. To compute $\text{IP}(\mathbf{x}', \mathbf{y}')$, they first sample two uniform random m -bit strings $\mathbf{x}'', \mathbf{y}'' \sim \{0, 1\}^m$. Consider the following cases:

- $p_{00} \geq 0.3$: In this case, we have either $p_{00} + p_{01} \geq 0.6$ or $p_{00} + p_{10} \geq 0.6$. Suppose the former holds without loss of generality, as the other case can be dealt with a similar argument. Both parties simulate Π on $\mathbf{x} := (\mathbf{x}', \mathbf{x}'')$ and $\mathbf{y} := (\mathbf{y}', \mathbf{y}'')$, obtain $(\mathbf{u}, \mathbf{v}) := \Pi(\mathbf{x}, \mathbf{y})$. Finally, they output $\neg \mathbf{u}$. We conclude that Π' succeeds with probability

$$\Pr[\Pi'(\mathbf{x}', \mathbf{y}') = \text{IP}(\mathbf{x}', \mathbf{y}')] = \Pr[\mathbf{u} \neq \text{IP}(\mathbf{x}', \mathbf{y}')] = p_{00} + p_{01} \geq 0.6.$$

- $p_{00} < 0.3$: In this case, we have $p_{01} + p_{10} \geq 0.6$. Similar to the previous case, both parties simulate Π on $\mathbf{x} := (\mathbf{x}', \mathbf{x}'')$ and $\mathbf{y} := (\mathbf{y}', \mathbf{y}'')$, obtain $(\mathbf{u}, \mathbf{v}) := \Pi(\mathbf{x}, \mathbf{y})$. Since $(\mathbf{x}'', \mathbf{y}'')$ is the common information between both parties, they can compute $\mathbf{v}' := \text{IP}(\mathbf{x}'', \mathbf{y}'')$ without extra communication. Finally, they output $\neg(\mathbf{u} \oplus \mathbf{v} \oplus \mathbf{v}')$. We conclude that Π' succeeds with probability

$$\Pr[\Pi(\mathbf{x}', \mathbf{y}') = \text{IP}(\mathbf{x}', \mathbf{y}')] = \Pr[\mathbf{u} \oplus \mathbf{v} \oplus \text{IP}(\mathbf{x}', \mathbf{y}') \oplus \text{IP}(\mathbf{x}'', \mathbf{y}'') = 1] = p_{01} + p_{10} \geq 0.6.$$

References

- [Aar10] Scott Aaronson. BQP and the polynomial hierarchy. In *Proceedings of the 42nd Symposium on Theory of Computing (STOC)*, pages 141–150. ACM, 2010. doi:[10.1145/1806689.1806711](https://doi.org/10.1145/1806689.1806711).
- [ABK24] Scott Aaronson, Harry Buhrman, and William Kretschmer. A qubit, a coin, and an advice string walk into a relational problem. In *Proceedings of the 15th Innovations in Theoretical Computer Science (ITCS)*, volume 287 of *LIPIcs*, pages 1:1–1:24. Schloss Dagstuhl, 2024. doi:[10.4230/LIPIcs.ITCS.2024.1](https://doi.org/10.4230/LIPIcs.ITCS.2024.1).
- [BFI23] Sam Buss, Noah Fleming, and Russell Impagliazzo. TFNP characterizations of proof systems and monotone circuits. In *Proceedings of the 14th Innovations in Theoretical Computer Science (ITCS)*, volume 251 of *LIPIcs*, pages 30:1–30:40. Schloss Dagstuhl, 2023. doi:[10.4230/LIPIcs.ITCS.2023.30](https://doi.org/10.4230/LIPIcs.ITCS.2023.30).
- [BHH⁺25] Ari Blondal, Hamed Hatami, Pooya Hatami, Chavdar Lalov, and Sivan Treiakh. Borsuk-Ulam and replicable learning of large-margin halfspaces. Technical report, arXiv, 2025. doi:[10.48550/arXiv.2503.15294](https://doi.org/10.48550/arXiv.2503.15294).
- [BHT25] Igor Balla, Lianna Hambardzumyan, and Istvá Tomon. Factorization norms and an inverse theorem for maxcut. Technical report, arXiv, 2025. doi:[10.48550/arXiv.2506.23989](https://doi.org/10.48550/arXiv.2506.23989).
- [CDM23] Arkadev Chattopadhyay, Yogesh Dahiya, and Meena Mahajan. Query complexity of search problems. In *Proceedings of the 48th Mathematical Foundations of Computer Science (MFCS)*, pages 34–1. Schloss Dagstuhl, 2023. doi:[10.4230/LIPIcs.MFCS.2023.34](https://doi.org/10.4230/LIPIcs.MFCS.2023.34).
- [CFK⁺21] Arkadev Chattopadhyay, Yuval Filmus, Sajin Koroth, Or Meir, and Toniann Pitassi. Query-to-communication lifting using low-discrepancy gadgets. *SIAM Journal on Computing*, 50(1):171–210, 2021. doi:[10.1137/19M1310153](https://doi.org/10.1137/19M1310153).
- [CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988. doi:[10.1137/0217015](https://doi.org/10.1137/0217015).
- [CLO⁺23] Lijie Chen, Zhenjian Lu, Igor Oliveira, Hanlin Ren, and Rahul Santhanam. Polynomial-time pseudodeterministic construction of primes. In *Proceedings of the 64th Symposium on Foundations of Computer Science (FOCS)*, pages 1261–1270. IEEE, 2023. doi:[10.1109/FOCS57990.2023.00074](https://doi.org/10.1109/FOCS57990.2023.00074).
- [CLV19] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. In *Proceedings of the 34th Computational Complexity Conference (CCC)*, pages 14:1–14:11. Schloss Dagstuhl, 2019. doi:[10.4230/LIPIcs.CCC.2019.14](https://doi.org/10.4230/LIPIcs.CCC.2019.14).
- [CMY23] Zachary Chase, Shay Moran, and Amir Yehudayoff. Stability and replicability in learning. In *Proceedings of the 64th Symposium on Foundations of Computer Science (FOCS)*, pages 2430–2439. IEEE, 2023. doi:[10.1109/FOCS57990.2023.00148](https://doi.org/10.1109/FOCS57990.2023.00148).

- [dRGR22] Susanna de Rezende, Mika Göös, and Robert Robere. Proofs, circuits, and communication. *SIGACT News*, 53(1), 2022. doi:[10.1145/3532737.3532745](https://doi.org/10.1145/3532737.3532745).
- [FF81] Peter Frankl and Zoltán Füredi. A short proof for a theorem of Harper about Hamming-spheres. *Discrete Mathematics*, 34(3):311–313, 1981. doi:[10.1016/0012-365X\(81\)90009-1](https://doi.org/10.1016/0012-365X(81)90009-1).
- [FGHH25] Yuting Fang, Mika Göös, Nathaniel Harms, and Pooya Hatami. Constant-cost communication is not reducible to k-hamming distance. In *Proceedings of the 57th Symposium on Theory of Computing (STOC)*, pages 565–571. ACM, 2025. doi:[10.1145/3717823.3718129](https://doi.org/10.1145/3717823.3718129).
- [Gav25] Dmytro Gavinsky. Unambiguous parity-query complexity. *Random Structures & Algorithms*, 66(3):e70010, 2025. doi:[10.1002/rsa.70010](https://doi.org/10.1002/rsa.70010).
- [GG11] Eran Gat and Shafi Goldwasser. Probabilistic search algorithms with unique answers and their cryptographic applications. In *Electronic Colloquium on Computational Complexity (ECCC)*, number TR11-136, 2011. URL: <https://eccc.weizmann.ac.il/report/2011/136/>.
- [GGJL25] Mika Göös, Tom Gur, Siddhartha Jain, and Jiawei Li. Quantum communication advantage in TFNP. In *Proceedings of the 57th Symposium on Theory of Computing (STOC)*, pages 1465–1475. ACM, 2025. doi:[10.1145/3717823.3718155](https://doi.org/10.1145/3717823.3718155).
- [GGMW20] Shafi Goldwasser, Ofer Grossman, Sidhanth Mohanty, and David Woodruff. Pseudo-deterministic streaming. In *Proceedings of the 11th Innovations in Theoretical Computer Science (ITCS)*, volume 151 of *LIPIcs*, pages 79:1–79:25. Schloss Dagstuhl, 2020. doi:[10.4230/LIPIcs.ITCS.2020.79](https://doi.org/10.4230/LIPIcs.ITCS.2020.79).
- [GGR13] Oded Goldreich, Shafi Goldwasser, and Dana Ron. On the possibilities and limitations of pseudodeterministic algorithms. In *Proceedings of the 4th Innovations in Theoretical Computer Science (ITCS)*, pages 127–138, 2013. doi:[10.1145/2422436.2422453](https://doi.org/10.1145/2422436.2422453).
- [GIPS21] Shafi Goldwasser, Russell Impagliazzo, Toniann Pitassi, and Rahul Santhanam. On the pseudo-deterministic query complexity of NP search problems. In *Proceedings of the 36th Computational Complexity Conference (CCC)*, pages 36:1–36:22. Schloss Dagstuhl, 2021. doi:[10.4230/LIPIcs.CCC.2021.36](https://doi.org/10.4230/LIPIcs.CCC.2021.36).
- [GKPW18] Mika Göös, Pritish Kamath, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for P^{NP} . *Computational Complexity*, 28(1):113–144, 2018. doi:[10.1007/s00037-018-0175-5](https://doi.org/10.1007/s00037-018-0175-5).
- [GKRS19] Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *Proceedings of the 10th Innovations in Theoretical Computer Science (ITCS)*, volume 124 of *LIPIcs*, pages 38:1–38:19. Schloss Dagstuhl, 2019. doi:[10.4230/LIPIcs.ITCS.2019.38](https://doi.org/10.4230/LIPIcs.ITCS.2019.38).
- [GLM⁺16] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. Rectangles are nonnegative juntas. *SIAM Journal on Computing*, 45(5):1835–1869, 2016. doi:[10.1137/15M103145X](https://doi.org/10.1137/15M103145X).

- [Gol11] Oded Goldreich. In a world of $P=BPP$. In *Studies in Complexity and Cryptography*, pages 191–232. Springer, 2011. doi:[10.1007/978-3-642-22670-0_20](https://doi.org/10.1007/978-3-642-22670-0_20).
- [GPW20] Mika Göös, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for BPP. *SIAM Journal on Computing*, 49(4):FOCS17–441–FOCS17–461, 2020. doi:[10.1137/17M115339X](https://doi.org/10.1137/17M115339X).
- [HHH23] Lianna Hambardzumyan, Hamed Hatami, and Pooya Hatami. Dimension-free bounds and structural results in communication complexity. *Israel Journal of Mathematics*, 253(2):555–616, 2023. doi:[10.1007/s11856-022-2365-8](https://doi.org/10.1007/s11856-022-2365-8).
- [HN12] Trinh Huynh and Jakob Nordström. On the virtue of succinct proofs: Amplifying communication complexity hardness to time-space trade-offs in proof complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, pages 233–248. ACM, 2012. doi:[10.1145/2213977.2214000](https://doi.org/10.1145/2213977.2214000).
- [ILPS22] Russell Impagliazzo, Rex Lei, Toniann Pitassi, and Jessica Sorrell. Reproducibility in learning. In *Proceedings of the 54th Symposium on Theory of Computing (STOC)*, pages 818–831. ACM, 2022. doi:[10.1145/3519935.3519973](https://doi.org/10.1145/3519935.3519973).
- [ILW23] Rahul Ilango, Jiatu Li, and Ryan Williams. Indistinguishability obfuscation, range avoidance, and bounded arithmetic. In *Proceedings of the 55th Symposium on Theory of Computing (STOC)*, pages 1076–1089. ACM, 2023. doi:[10.1145/3564246.3585187](https://doi.org/10.1145/3564246.3585187).
- [LOS21] Zhenjian Lu, Igor Oliveira, and Rahul Santhanam. Pseudodeterministic algorithms and the structure of probabilistic time. In *Proceedings of the 53rd Symposium on Theory of Computing (STOC)*, pages 303–316. ACM, 2021. doi:[10.1145/3406325.3451085](https://doi.org/10.1145/3406325.3451085).
- [LS09] Nati Linial and Adi Shraibman. Learning complexity vs communication complexity. *Combinatorics, Probability and Computing*, 18(1-2):227–245, 2009. doi:[10.1017/S0963548308009656](https://doi.org/10.1017/S0963548308009656).
- [OS17] Igor Oliveira and Rahul Santhanam. Pseudodeterministic constructions in subexponential time. In *Proceedings of the 49th Symposium on Theory of Computing (STOC)*, pages 665–677. ACM, 2017. doi:[10.1145/3055399.3055500](https://doi.org/10.1145/3055399.3055500).
- [Vad12] Salil Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1–3):1–336, 2012. doi:[10.1561/0400000010](https://doi.org/10.1561/0400000010).