Communication Complexity: Lecture 5

Nathan Harms

February 6, 2026

1 Public vs. Private Randomness

Theorem 1.1 (Newman's Theorem). *Let $M \in \{0,1\}^{m \times n}$ and let Π be any public-coin randomized protocol for M with error ε. Then there is a public-coin randomized protocol for M with error 2ε that uses at most $\log\log(mn) + O(\log(1/\varepsilon))$ random bits.*

Proof. A randomized protocol for M with cost c and error ε is a probability distribution $\mathcal{D}$ over deterministic protocols with depth c, such that

$$\forall x, y\colon \qquad \mathop{\mathbb{P}}_{\boldsymbol{\Pi} \sim \mathcal{D}} [\boldsymbol{\Pi}(x,y) \neq M_{x,y}] \leq \varepsilon. \tag{1}$$

Fix some number k and let $\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_k \sim \mathcal{D}$ be deterministic protocols drawn independently from $\mathcal{D}$. For input pair (x,y), we say that a sequence of protocols $\Pi_1, \dots, \Pi_k$ is *good for* (x,y) if

$$\#\{i \in [k] \mid \Pi_i(x,y) = M_{x,y}\} \geq (1 - 2\varepsilon)k,$$

and otherwise it is *bad for* (x,y).

Claim 1.2. *Suppose there exists a sequence $\Pi_1, \dots, \Pi_k$ that is good for all input pairs (x,y). Then there is a randomized protocol computing M with error 2ε and cost c, using at most $\lceil \log k \rceil$ bits of randomness.*

Proof of claim. The new protocol chooses a random $\boldsymbol{i} \sim [k]$ (which requires $\lceil \log k \rceil$ bits of randomness) and then runs $\Pi_{\boldsymbol{i}}(x,y)$, which has cost c. By definition of goodness, the error of the protocol on (x,y) is

$$\mathop{\mathbb{P}}_{\boldsymbol{i} \sim [k]} [\Pi]_{\boldsymbol{i}}(x,y) \neq M_{x,y} = \frac{1}{k} \cdot \#\{i \in [k] \mid \Pi_i(x,y) \neq M_{x,y}\} \leq 2\varepsilon.$$

Since this holds for all inputs (x,y), the protocol has error 2ε. ■

It remains to show that there exists a sequence of protocols $\Pi_1, \dots, \Pi_k$ that is good for all input pairs (x,y) simultaneously, for appropriately chosen value of k.

Claim 1.3. *For all input pairs (x,y),*

$$\mathbb{P}\left[(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_k) \text{ is bad for } (x,y)\right] \leq e^{-2k\varepsilon^2}.$$

Proof of claim. For a fixed input pair (x, y), let $\boldsymbol{Z}_i = \mathbb{1}[\boldsymbol{\Pi}_i(x, y) \neq M_{x,y}]$ be the random variable that is 1 if and only if $\Pi_i(x, y)$ is incorrect. By Equation (1), $\mathbb{E}[\boldsymbol{Z}_i] \leq \varepsilon$, so

$$\mathbb{E}\left[\sum_{i=1}^{k} \boldsymbol{Z}_i\right] \leq k\varepsilon.$$

The sequence $(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_k)$ is bad for (x, y) if and only if $\sum_i \boldsymbol{Z}_i > 2k\varepsilon$. Since $\boldsymbol{Z}_i$ are independent Bernoulli random variables, by the Hoeffding bound,

$$\mathbb{P}\left[\sum_{i=1}^{k} \boldsymbol{Z}_i > 2k\varepsilon\right] \leq e^{-\frac{2k^2\varepsilon^2}{k}} = e^{-2k\varepsilon^2}. \qquad \blacksquare$$

We may now complete the proof of the theorem by choosing $k = \frac{1}{2\varepsilon^2}\ln(mn)$. Then by the union bound,

$$\begin{aligned}\mathbb{P}[\exists (x, y) \in [m] \times [n] : \quad &(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_k) \text{ is bad for } (x, y)] \\ &\leq mn \cdot \max_{x,y} \mathbb{P}[(\boldsymbol{\Pi}_1, \dots, \boldsymbol{\Pi}_k) \text{ is bad for } (x, y)] \\ &\leq mn \cdot e^{-2k\varepsilon^2} < 1.\end{aligned}$$

Therefore there exists a sequence $\Pi_1, \dots, \Pi_k$ that is good for all (x, y), and we obtain a protocol which uses at most $\lceil \log k \rceil = \log\log(mn) + O(\log(1/\varepsilon))$ random bits. $\blacksquare$

Corollary 1.4. *For all $M \in \{0,1\}^{m \times n}$ and $\varepsilon > 0$, $\mathsf{R}^{\text{priv}}_{2\varepsilon}(M) \leq \mathsf{R}_{\varepsilon}(M) + \log\log(mn) + O(\log(1/\varepsilon))$.*

Proof. Using Newman's theorem, Alice can generate all $\log\log(mn) + O(\log(1/\varepsilon))$ random bits privately and then send them in a message to Bob. Then they can run the public-coin protocol. $\blacksquare$

2 Yao's Principle

Let us now turn our attention to lower bounds on randomized communication. Lower bounds on randomized communication (and in fact, randomized algorithms in general) usually use *Yao's principle*, or *Yao's minimax lemma*.

Definition 2.1 (Distributional Communication Complexity)**.** Let $M \in \{0,1\}^{m \times n}$ be any communication problem and let μ be any probability distribution over the inputs $[m] \times [n]$. Then we write $\mathsf{D}_{\mu,\varepsilon}(M)$ for the minimum depth of a *deterministic* protocol Π such that

$$\mathop{\mathbb{P}}_{(\boldsymbol{x},\boldsymbol{y}) \sim \mu}[\Pi(x, y) \neq M_{x,y}] \leq \varepsilon.$$

In other words, $\mathsf{D}_{\mu,\varepsilon}(M)$ is the optimal cost of a protocol with *average* error ε over μ.

Yao's minimax lets us rewrite randomized communication complexity the *deterministic* communication complexity with small average error:

Theorem 2.2 (Yao's Minimax)**.** *Write $\mathcal{D}(\mathcal{A})$ for the set of probability distributions over any domain $\mathcal{A}$. Let $M \in \{0,1\}^{m\times n}$ be any communication problem.*

$$\mathsf{R}_\varepsilon(M) = \max_{\mu\in\mathcal{D}([m]\times[n])} \mathsf{D}_{\mu,\varepsilon}(M).$$

Proof. We will require a more fundamental minimax theorem known as *von Neumann's minimax theorem*:

Theorem 2.3 (von Neumann Minimax Theorem)**.** *Let $Q \in \mathbb{R}^{\mathcal{A}\times\mathcal{B}}$ be any matrix. Then*

$$\max_{a\in\mathcal{D}(\mathcal{A})} \min_{b\in\mathcal{D}(\mathcal{B})} a^\top Q b = \min_{b\in\mathcal{D}(\mathcal{B})} \max_{a\in\mathcal{D}(\mathcal{A})} a^\top Q b.$$

Equivalently,

$$\max_{a\in\mathcal{D}(\mathcal{A})} \min_{b\in\mathcal{D}(\mathcal{B})} \mathop{\mathbb{E}}_{\boldsymbol{i}\sim a,\boldsymbol{j}\sim b}[Q_{\boldsymbol{i},\boldsymbol{j}}] = \min_{b\in\mathcal{D}(\mathcal{B})} \max_{a\in\mathcal{D}(\mathcal{A})} \mathop{\mathbb{E}}_{\boldsymbol{i}\sim a,\boldsymbol{j}\sim b}[Q_{\boldsymbol{i},\boldsymbol{j}}]$$

Let us rephrase the statement that we wish to prove:

Claim 2.4. *For every $c \in \mathbb{N}$ and $\varepsilon > 0$, $R_\varepsilon(M) \le c$ if and only if, for every distribution μ over inputs, $\mathsf{D}_{\mu,\varepsilon}(M) \le c$.*

To apply von Neumann's theorem, recall that a randomized protocol with cost c is a probability distribution over deterministic protocols Π with cost c. Let $\mathcal{A}_c$ be the set of all deterministic protocols with cost c.

Now define the matrix $Q_c \in \{0,1\}^{([m]\times[n])\times\mathcal{A}_c}$ whose rows are indexed by input pairs $(x,y) \in [m]\times[n]$, and whose columns are indexed by deterministic protocols with cost c. Then define

$$Q_c((x,y),\Pi) = \mathbb{1}[\Pi(x,y) \neq M_{x,y}],$$

i.e. put a 1 in entry $((x,y),A)$ if and only if the cost-c protocol Π makes an error on input (x,y). From von Neumann's minimax, we obtain the following:

$$\begin{aligned}&\min_{A\in\mathcal{D}(\mathcal{A}_c)} \max_{\mu\in\mathcal{D}([m]\times[n])} \mathop{\mathbb{E}}_{\boldsymbol{\Pi}\sim A}\left[\mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[Q_c(\boldsymbol{\Pi},(\boldsymbol{x},\boldsymbol{y}))]\right] \\ &= \max_{\mu\in\mathcal{D}([m]\times[n])} \min_{A\in\mathcal{D}(\mathcal{A}_c)} \mathop{\mathbb{E}}_{\boldsymbol{\Pi}\sim A}\left[\mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[Q_c(\boldsymbol{\Pi},(\boldsymbol{x},\boldsymbol{y}))]\right]\end{aligned} \tag{2}$$

Our goal is to replace this equation with:

$$\min_{A\in\mathcal{D}(\mathcal{A}_c)} \max_{(x,y)\in[m]\times[n]} \mathop{\mathbb{P}}_{\boldsymbol{\Pi}\sim A}[\boldsymbol{\Pi}(x,y) \neq M_{x,y}] = \max_{\mu\in\mathcal{D}([m]\times[n])} \min_{A\in\mathcal{A}_c} \mathop{\mathbb{P}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[\Pi(\boldsymbol{x},\boldsymbol{y}) \neq M_{x,y}]. \tag{3}$$

This will suffice to prove Claim 2.4, hence the whole theorem, because

$$R_\varepsilon(M) = \min\{c \mid \text{ left side of (3) is } \le \varepsilon\}$$

while

$$\max_\mu \mathsf{D}_{\mu,\varepsilon}(M) = \min\{c \mid \text{ right side of (3) is } \le \varepsilon\}$$

Let us prove two statements:

Claim 2.5. *The left side of Equation (2) is*

$$\min_{A\in\mathcal{D}(\mathcal{A}_c)} \max_{\mu\in\mathcal{D}([m]\times[n])} \mathop{\mathbb{E}}_{\mathbf{\Pi}\sim A}\left[\mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[Q_c(\mathbf{\Pi},(\boldsymbol{x},\boldsymbol{y}))]\right] = \min_{A\in\mathcal{D}(\mathcal{A}_c)} \max_{(x,y)\in[m]\times[n]} \mathop{\mathbb{P}}_{\mathbf{\Pi}\sim A}[\mathbf{\Pi}(x,y)\neq M_{x,y}].$$

Proof of claim. It suffices to show that for every distribution $A\in\mathcal{D}(\mathcal{A}_c)$,

$$\max_{\mu\in\mathcal{D}([m]\times[n])} \mathop{\mathbb{E}}_{\mathbf{\Pi}\sim A}\left[\mathop{\mathbb{P}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[\mathbf{\Pi}(\boldsymbol{x},\boldsymbol{y})\neq M_{\boldsymbol{x},\boldsymbol{y}}]\right] = \max_{(x,y)\in[m]\times[n]} \mathop{\mathbb{P}}_{\mathbf{\Pi}\sim A}[\mathbf{\Pi}(x,y)\neq M_{x,y}].$$

It is clear that the left side of this equation is at least the right side, because the max is over a larger set (on the right we are maximizing over distributions μ that put the entire probability mass on a single input pair (x,y).). Therefore we only need to show that the left side is *at most* the right side. Fix any distribution $\mu\in\mathcal{D}([m]\times[n])$ and let

$$\delta_\mu := \mathop{\mathbb{E}}_{\mathbf{\Pi}\sim A}\left[\mathop{\mathbb{P}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[\mathbf{\Pi}(\boldsymbol{x},\boldsymbol{y})\neq M_{\boldsymbol{x},\boldsymbol{y}}]\right] = \mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}\left[\mathop{\mathbb{P}}_{\mathbf{\Pi}\sim A}[\mathbf{\Pi}(\boldsymbol{x},\boldsymbol{y})\neq M_{\boldsymbol{x},\boldsymbol{y}}]\right].$$

Since this is an expectation over $(\boldsymbol{x},\boldsymbol{y})\sim\mu$, there must exist some fixed (x,y) such that

$$\mathop{\mathbb{P}}_{\mathbf{\Pi}\sim A}[\mathbf{\Pi}(x,y)\neq M_{x,y}] \geq \mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}\left[\mathop{\mathbb{P}}_{\mathbf{\Pi}\sim A}[\mathbf{\Pi}(\boldsymbol{x},\boldsymbol{y})\neq M_{\boldsymbol{x},\boldsymbol{y}}]\right] = \delta_\mu.$$

Therefore $\max_\mu \delta_\mu \leq \max_{x,y}\mathbb{P}[\mathbf{\Pi}(x,y)\neq M_{x,y}]$, as desired. ∎

Claim 2.6. *The right side of Equation (2) is*

$$\max_{\mu\in\mathcal{D}([m]\times[n])} \min_{A\in\mathcal{D}(\mathcal{A}_c)} \mathop{\mathbb{E}}_{\mathbf{\Pi}\sim A}\left[\mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[Q_c(\mathbf{\Pi},(\boldsymbol{x},\boldsymbol{y}))]\right] = \max_{\mu\in\mathcal{D}([m]\times[n])} \min_{A\in\mathcal{A}_c} \mathop{\mathbb{P}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[\Pi(\boldsymbol{x},\boldsymbol{y})\neq M_{x,y}].$$

Proof. This proof is essentially the same as the one above. ∎

This concludes the proof of Theorem 2.2. ∎

3 Discrepancy

For deterministic communication complexity, there is an easy proof that, if there is a cost c protocol computing $M\in\{0,1\}^{m\times n}$, then there is a monochromatic rectangle of size $mn\cdot 2^{-c}$. We have no similar proof for randomized communication. Even the simplest version of this is open, for randomized protocols of cost $c=O(1)$:

Open Problem 1 ([CLV19, HHH23])**.** *Let $M\in\{0,1\}^{m\times n}$ have $\mathsf{R}_{1/4}(M)=O(1)$. Then there exists a monochromatic rectangle $X\times Y$ in M such that*

$$|X\times Y| = \Omega(mn).$$

For randomized protocols, instead we have the notion of *discrepancy*, which plays a similar role as monochromatic rectangles.

Definition 3.1 (Discrepancy)**.** Let $M \in \{0,1\}^{m\times n}$, let μ be any probability distribution over entries of the matrix. Let $\mathsf{zeroes}(M) := \{(x,y) \mid M_{x,y} = 0\}$ be the set of 0-entries of M, and $\mathsf{ones}(M) := \{(x,y) \mid M_{x,y} = 1\}$ be the set of 1-entries. For any set $S \subseteq [m] \times [n]$, we define

$$\mu_0(S) := \mu\left(S \cap \mathsf{zeroes}(M)\right)$$
$$\mu_1(S) := \mu\left(S \cap \mathsf{ones}(M)\right).$$

Now we define

$$\mathrm{disc}_\mu(M) := \max_{X\subset[m], Y\subset[n]} |\mu_0(X\times Y) - \mu_1(X \times Y)|$$

Whereas deterministic protocols with cost c must have a monochromatic rectangle with density 2^{-c}, randomized protocols with cost c must have a rectangle with discrepancy $2^{-O(c)}$, as we show in the next theorem.

Theorem 3.2. *Let $M \in \{0,1\}^{m\times n}$ and let $\varepsilon > 0$. Then for all distributions μ of the entries of M,*

$$\mathsf{R}_\varepsilon(M) \geq \log\left(\frac{1-2\varepsilon}{\mathrm{disc}_\mu(M)}\right).$$

Proof. Let $c = \mathsf{R}_\varepsilon(M)$ be the cost of the optimal randomized protocol for M. Fix any distribution μ. First we apply Yao's lemma,

$$c = \mathsf{R}_\varepsilon(M) \geq \mathsf{D}_{\mu,\varepsilon},$$

so there exists a deterministic protocol Π with cost at most c and average error ε over the distribution μ. For each leaf $L \subseteq [m] \times [n]$ of the protocol Π, we may assume without loss of generality that the label of L is 0 if $\mu_0(L) \geq \mu_1(L)$ and 1 if $\mu_1(L) > \mu_0(L)$; otherwise we can swap the label of the leaf and reduce the average error of the protocol. Now we can write the average error as

$$\begin{aligned}
\varepsilon &\geq \sum_{\text{leaves } L} \mathop{\mathbb{P}}_{(\boldsymbol{x},\boldsymbol{y})\sim\mu}[(\boldsymbol{x},\boldsymbol{y}) \in L \wedge \text{ label of } L \text{ is not } M_{x,y}] \\
&= \sum_{\text{leaves } L} \min\{\mu_0(L), \mu_1(L)\} \\
&= \sum_{\text{leaves } L} \frac{1}{2}\left(\mu_0(L) + \mu_1(L) - |\mu_0(L) - \mu_1(L)|\right) \qquad \text{(using } \min(a,b) = \tfrac{1}{2}(a+b-|b-a|)\text{)} \\
&= \frac{1}{2}\left(\sum_{\text{leaves } L} \mu(L)\right) - \frac{1}{2}\left(\sum_{\text{leaves } L} |\mu_0(L) - \mu_1(L)|\right) \\
&= \frac{1}{2} - \frac{1}{2}\left(\sum_{\text{leaves } L} |\mu_0(L) - \mu_1(L)|\right) \\
&\geq \frac{1}{2} - \frac{1}{2} \cdot 2^c \cdot \mathrm{disc}_\mu(M),
\end{aligned}$$

where we used the fact that there are at most 2^c leaves of Π. Rearranging this inequality, we obtain

$$2^c \geq \frac{1-2\varepsilon}{\mathrm{disc}_\mu(M)},$$

and taking the log of both sides gives us the desired conclusion. ■

4 Lower Bound for Inner Product

Definition 4.1 (Inner Product)**.** Let $x, y \in \{0,1\}^n$. We define the inner product mod 2 problem as

$$\mathrm{IP}_n(x,y) := \sum_{i=1}^{n} x_i y_i \mod 2.$$

Theorem 4.2. $\mathsf{R}_{1/4}(\mathrm{IP}_n) = \Theta(n)$.

Proof. We will use the discrepancy lower bound with μ being the uniform distribution over $\{0,1\}^n \times \{0,1\}^n$.

Let $X \times Y$ be any rectangle

$$\begin{aligned}
\mathrm{disc}_\mu(X \times Y) &= |\mu_0(X \times Y) - \mu_1(X \times Y)| \\
&= \left| \mathop{\mathbb{E}}_{\boldsymbol{x},\boldsymbol{y}} [\mathbb{1}[\boldsymbol{x} \in X] \, \mathbb{1}[\boldsymbol{y} \in Y] \, (\mathbb{1}[\langle \boldsymbol{x}, \boldsymbol{y} \rangle \text{ even }] - \mathbb{1}[\langle \boldsymbol{x}, \boldsymbol{y} \rangle \text{ odd }])] \right| \\
&= \left| \mathop{\mathbb{E}}_{\boldsymbol{x},\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{x} \in X] \, \mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right] \right| \\
&= \sqrt{ \mathop{\mathbb{E}}_{\boldsymbol{x},\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{x} \in X] \, \mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right]^2 }
\end{aligned}$$

To advance in the proof, let's introduce an important inequality called Jensen's inequality.

Proposition 4.3 (Jensen's Inequality)**.** *Let $\phi \colon \mathbb{R} \to \mathbb{R}$ be any convex function. Then for any random variable $\boldsymbol{z} \in \mathbb{R}$,*

$$\phi(\mathbb{E}[\phi(\boldsymbol{z})]) \le \mathbb{E}[\phi(\boldsymbol{z})].$$

Since the function $z \mapsto z^2$ is convex,

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{x},\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{x} \in X] \, \mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right]^2 &= \mathop{\mathbb{E}}_{\boldsymbol{x}} \left[\mathbb{1}[\boldsymbol{x} \in X] \mathop{\mathbb{E}}_{\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right] \right]^2 \\
&\le \mathop{\mathbb{E}}_{\boldsymbol{x}} \left[\mathbb{1}[\boldsymbol{x} \in X]^2 \mathop{\mathbb{E}}_{\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right]^2 \right] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{x}} \left[\mathbb{1}[\boldsymbol{x} \in X] \mathop{\mathbb{E}}_{\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right]^2 \right] \\
&\le \mathop{\mathbb{E}}_{\boldsymbol{x}} \left[\mathop{\mathbb{E}}_{\boldsymbol{y}} \left[\mathbb{1}[\boldsymbol{y} \in Y] \, (-1)^{\langle \boldsymbol{x}, \boldsymbol{y} \rangle} \right]^2 \right],
\end{aligned}$$

where the last inequality holds since both terms inside the expectation are positive and $\mathbb{1}[\boldsymbol{x} \in X] \le$

1. Now

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{x}}\left[\mathop{\mathbb{E}}_{\boldsymbol{y}}\left[\mathbb{1}\left[\boldsymbol{y}\in Y\right](-1)^{\langle \boldsymbol{x},\boldsymbol{y}\rangle}\right]^2\right] &= \mathop{\mathbb{E}}_{\boldsymbol{x}}\left[\mathop{\mathbb{E}}_{\boldsymbol{y}}\left[\mathbb{1}\left[\boldsymbol{y}\in Y\right](-1)^{\langle \boldsymbol{x},\boldsymbol{y}\rangle}\right]\mathop{\mathbb{E}}_{\boldsymbol{y}}\left[\mathbb{1}\left[\boldsymbol{y}\in Y\right](-1)^{\langle \boldsymbol{x},\boldsymbol{y}\rangle}\right]\right] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{x}}\left[\mathop{\mathbb{E}}_{\boldsymbol{y},\boldsymbol{y}'}\left[\mathbb{1}\left[\boldsymbol{y},\boldsymbol{y}'\in Y\right](-1)^{\langle \boldsymbol{x},\boldsymbol{y}\rangle}(-1)^{\langle \boldsymbol{x},\boldsymbol{y}'\rangle}\right]\right] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{x}}\left[\mathop{\mathbb{E}}_{\boldsymbol{y},\boldsymbol{y}'}\left[\mathbb{1}\left[\boldsymbol{y},\boldsymbol{y}'\in Y\right](-1)^{\langle \boldsymbol{x},\boldsymbol{y}+\boldsymbol{y}'\rangle}\right]\right] \\
&= \mathop{\mathbb{E}}_{\boldsymbol{y},\boldsymbol{y}'}\left[\mathbb{1}\left[\boldsymbol{y},\boldsymbol{y}'\in Y\right]\mathop{\mathbb{E}}_{\boldsymbol{x}}\left[(-1)^{\langle \boldsymbol{x},\boldsymbol{y}+\boldsymbol{y}'\rangle}\right]\right]
\end{aligned}$$

For any fixed $y, y' \in \{0,1\}^n$, if $y + y'$ has any coordinate $i \in [n]$ where $y_i + y'_i = 1$, then

$$\mathop{\mathbb{E}}_{\boldsymbol{x}}\left[(-1)^{\langle \boldsymbol{x},y+y'\rangle}\right] = \mathop{\mathbb{E}}_{\boldsymbol{x}}\left[(-1)^{\langle \boldsymbol{x}_{[n]\setminus i},(y+y')_{[n]\setminus i}\rangle}(-1)^{\boldsymbol{x}_i}\right] = \mathop{\mathbb{E}}_{\boldsymbol{x}_{[n]\setminus i}}\left[(-1)^{\langle \boldsymbol{x}_{[n]\setminus i},(y+y')_{[n]\setminus i}\rangle}\right]\mathop{\mathbb{E}}_{\boldsymbol{x}_i}\left[(-1)^{\boldsymbol{x}_i}\right] = 0,$$

where $x_{[n]\setminus i}$ is the substring obtained by deleting coordinate i. On the other hand, if $y + y'$ has no coordinate i where $y_i + y'_i = 1$, then $y = y'$. So

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{y},\boldsymbol{y}'}\left[\mathbb{1}\left[\boldsymbol{y},\boldsymbol{y}'\in Y\right]\mathop{\mathbb{E}}_{\boldsymbol{x}}\left[(-1)^{\langle \boldsymbol{x},\boldsymbol{y}+\boldsymbol{y}'\rangle}\right]\right] &= \mathbb{P}\left[\boldsymbol{y}=\boldsymbol{y}'\right]\cdot\mathop{\mathbb{E}}_{\boldsymbol{y},\boldsymbol{y}'}\left[\mathbb{1}\left[\boldsymbol{y}\in Y\right]\mathop{\mathbb{E}}_{\boldsymbol{x}}\left[(-1)^{\langle \boldsymbol{x},2y\rangle}\right] \;\middle|\; \boldsymbol{y}=\boldsymbol{y}'\right] \\
&= \frac{1}{2^n}\cdot\mathop{\mathbb{P}}_{\boldsymbol{y}}\left[\boldsymbol{y}\in Y\right] \le \frac{1}{2^n},
\end{aligned}$$

since $(-1)^{\langle x,2y\rangle} = 1$ always, since $\sum_i 2x_iy_i$ is even. Therefore

$$\mathrm{disc}_\mu(X\times Y) \le \sqrt{2^{-n}} = 2^{-n/2},$$

so

$$\mathsf{R}_{1/4}(\mathrm{IP}_n) \ge \Omega\left(\log\left(\frac{1}{\mathrm{disc}_\mu(\mathrm{IP}_n)}\right)\right) = \Omega(n). \qquad \blacksquare$$

References

[CLV19] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. In *34th Computational Complexity Conference (CCC 2019)*, pages 14–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019.

[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.