Information Value of Two-Prover Games[*]

Mark Braverman[1] and Young Kun Ko[2]

1   **Princeton University, 35 Olden St. Princeton NJ 08540, USA**
    `mbraverm@cs.princeton.edu`
2   **Princeton University, 35 Olden St. Princeton NJ 08540, USA**
    `yko@cs.princeton.edu`

—— Abstract ——

We introduce a generalization of the standard framework for studying the difficulty of two-prover games. Specifically, we study the model where Alice and Bob are allowed to communicate (with information constraints) – in contrast to the usual two-prover game where they are not allowed to communicate after receiving their respective input. We study the trade-off between the information cost of the protocol and the achieved value of the game after the protocol. In particular, we show the connection of this trade-off and the amortized behavior of the game (i.e. repeated value of the game). We show that if one can win the game with at least $(1-\epsilon)$-probability by communicating at most ϵ bits of information, then one can win n copies with probability at least $2^{-O(\epsilon n)}$. This gives an intuitive explanation why Raz's counter-example to strong parallel repetition [16] (the odd cycle game) is a counter-example to strong parallel repetition – one can win the odd-cycle game on a cycle of length m by communicating $O(m^{-2})$-bits where m is the number of vertices.

Conversely, for projection games, we show that if one can win n copies with probability larger than $(1-\epsilon)^n$, then one can win one copy with at least $(1 - O(\epsilon))$-probability by communicating $O(\epsilon)$ bits of information. By showing the equivalence between information value and amortized value, we give an alternative direction for further works in studying amortized behavior of the two-prover games.

The main technical tool is the "Chi-Squared Lemma" which bounds the information cost of the protocol in terms of Chi-Squared distance, instead of usual divergence. This avoids the square loss from using Pinsker's Inequality.

1   Introduction

A *two-prover one-round game* $\mathcal{G}$ on a bipartite graph with (U, V, E) with distribution $\mathcal{D}$ on E is defined as a following process. Referee picks $(u, v) \in E$ according to $\mathcal{D}$, then sends u to Alice and v to Bob. Then Alice gives an assignment to u and Bob to v from alphabet Σ. Referee checks if they are "valid" assignment for the edge (u, v). In this setting, we are interested in what the best response strategy is for Alice and Bob. In particular, we want to find $f : U \to \Sigma$ which denotes Alice's strategy and $g : V \to \Sigma$ which denotes Bob's strategy

---

that maximize the fraction of satisfied edges. In particular, we want to compute the strategy that achieves the *value* of the game which is defined as

$$\text{val}(\mathcal{G}) := \max_{f,g} \Pr_{(u,v)\sim\mathcal{D}} \left[\pi_{(u,v)}(f(u), g(v)) = 1 \right],$$

that is the probability of satisfying a randomly chosen edge according to the best response strategy where $\pi_{(u,v)}$ is the verification function by the referee for the edge (u, v).

In the above setup, it is crucial that Alice's assignment only depends on the input u and Bob's assignment only depends on the input v. In other words, Alice and Bob are assumed to be in separate rooms, leaking zero bits of information about their respective input. Having introduced the amount of information communicated between Alice and Bob into the picture, we could then reformulate the *value* of the game as:

> *Given a game $\mathcal{G}$, Alice and Bob communicate zero bits of information. Then what is the best chance of winning the game?*

Then it is natural to extend to following question: 'If Alice and Bob are allowed to communicate limited information, what is the value of the game?' In particular, we could explicitly ask the following question.

> *If Alice and Bob are allowed to communicate ε bits of information, then what is the value of the game? (in terms of ε)*

First, note that this is a well-defined quantity in a sense that there is a following explicitly bounded curve. Observe that if $|U| = |V| = n$, and they are allowed to communicate $O(\log n)$-bits, the value of the game becomes 1 (if all the edges indeed have at least one satisfying assignment, which can be assumed without loss of generality) due to the following naive strategy. Alice simply sends the hash of her input to Bob, which requires at most $\log n$-bits to do so and vice-versa. Since Alice and Bob both know (u, v), they can simply pick a satisfying assignment (using shared randomness) for (u, v) then answer accordingly. We can further tighten the upper bound (for the amount of information) if we know the structure of the graph. In particular, if the graph were d-regular, given $O(\log d)$-bits, the value of the game again becomes 1, since the entropy (of Alice's input given Bob's input and vice-versa) is at most $\log d$. (We will show this explicitly)

We would like to further investigate *the trade-off between the information vs. the value of the game*. In particular, we initiate the study of *information value of the game*, that is *how much information is necessary* to win the game with say probability $> 1 - \delta$ in terms of δ, which we define more explicitly in Section 2.

Note that this notion can classify "how intrinsically hard" a given two-prover one-round game is. In particular, one could view a game as being "hard" if the value of the game is "resistant" to added information, easy otherwise, providing a better spectrum for analyzing the intrinsic hardness of the game – which is indeed related to the amortized value of the game.

1.1   Our Contribution

We connect the information value of the game with the amortized value of the game – i.e. the value of the repeated game. We note that previous parallel repetition literature can be "translated" as Alice and Bob sharing a common hint provided by the referee. Then we can "remove" the referee from the picture and instead let Alice and Bob sample a common hint. This translation is what we call the "Chi-squared lemma," which is the main technical contribution of this paper.

▶ **Lemma 1** (Chi-squared lemma). *Suppose Alice has access to P and Bob has access to Q which are probability distributions over the universe $\mathcal{U}$. Suppose further that there exists a common distribution R such that $D(R||P) < \varepsilon$ and $D(R||Q) < \varepsilon$. Then there exists a protocol Π that outputs a sample from the distribution $\tilde{P}$ with information cost $\gamma\varepsilon$ with $D(R||\tilde{P}) < O(\varepsilon)$ for some constant $\gamma > 0$.*

where "information cost" refers to the amount of information revealed to each other as in [4].

We also remark that the idea of referring to Chi-Squared distance to bound information cost was also used in proving sharp round complexity of pointer chasing problem [20], avoiding square-loss in the parameter due to the application of Pinsker's inequality.

Via this lemma, we can "translate" previous parallel repetition literature (in particular [3]) as a blackbox to obtain following main theorem.

▶ **Theorem 2** (informal). *If $\mathrm{val}(\mathcal{G}^n) \geq (1 - \varepsilon)^n = 2^{-\Theta(\varepsilon n)}$ and $\mathcal{G}$ is a projection game, that is for any valid pair of assignments, Bob's assignment is a projection of Alice's assignment, then one can win $\mathcal{G}$ with probability $1 - O(\varepsilon)$ by communicating $O(\varepsilon)$-bits of information.*

To show the converse, we need to show that a low-information protocol can be translated to a zero-communication protocol with insignificant loss in success probability. We use results from correlated sampling (in particular we use a lemma from [3]) which yields the converse to the main theorem.

▶ **Theorem 3** (informal). *If one can win $\mathcal{G}$ with probability $1 - \varepsilon$ by communicating ε-bits of information, then $\mathrm{val}(\mathcal{G}^n) \geq 2^{-O(\varepsilon n)}$.*

We also remark that Theorem 3 **does not assume** that $\mathcal{G}$ is a projection game.

1.2 Proof Overview

Proof of Theorem 3

It suffices to show that a low-information cost protocol can be translated to a zero-communication protocol with success parameter depending on the information cost. In particular, we want to show that an $O(I)$ information cost protocol can be simulated by two non-communicating parties with $2^{-O(I)}$ success probability. For our purpose $I = n\varepsilon$, since information cost tensorizes with many copies. Note that if Alice and Bob managed to sample a correct transcript together, the transcript is correct with $(1 - \varepsilon)^{O(n)}$ probability since each coordinates are chosen independently. We show that the zero-communication sampling lemma from [3] indeed gives the range of parameters that we need.

Proof of Theorem 2

In [3], one could view the "common hint" (which actually comes from the multiple copies of the game) as given by the referee to Alice and Bob via sampling some random coordinates with their answers then sending them to Alice and Bob. However, this does not fit in our framework since the referee samples the hint, not Alice and Bob. Using the Chi-squared lemma, we allow Alice and Bob to jointly sample such a hint with $O(\varepsilon)$ information cost. [3] shows that one can win a single copy with probability $1 - O(\varepsilon)$ after running the joint sampling protocol (under some distribution $\mathcal{D}'$ which is $O(\varepsilon)$-away in terms of divergence from the real distribution). We further show that having a strategy for such $\mathcal{D}'$ suffices to win the original game with probability $1 - O(\varepsilon)$ as well.

Chi-squared lemma

Suppose Alice has access to P_1 and Bob has access to P_2 with a guarantee that there exists some common distribution R such that $D(R||P_1), D(R||P_2) < \varepsilon$. In such a setting, if Alice samples from P_1 and tries to transmit the sample to Bob, the naive information cost would be $D(P_1||P_2)$. This could be unbounded however due to the case where P_1 contains an element in the support that is not in the support of P_2. i.e. $D(P_1||P_2) = \infty$. To rule out such scenario, we give Bob an ability to reject. Instead of receiving the full description or index of the sample, Bob receives a stream of hash values and rejects the stream when he "cannot understand" or in other words expects the divergence from the sample to be high. Via this rejection, Bob will not be "too surprised" about Alice's sample. But the main problem with the above simple rejection protocol is that it allows Alice to learn too much about P_2 from Bob's response. For instance, if Bob rejects a sample, then Alice learns that this sample occurs "infrequently" in P_2. In order to "confuse" Alice and prevent her from learning too much about P_2, Bob rejects a valid stream with some constant probability. This suffices to confuse Alice and learn only $O(\varepsilon)$-information about P_2.

2 Preliminaries

2.1 Information Theory

In this section, we provide background on information theory that will be used to prove main results. We remark that throughout the paper, log is of base 2 and ln is of base e. For further references, we refer the reader to [7].

▶ **Definition 4** (Entropy). The *entropy* of a random variable A, denoted by $H(A)$ is defined as

$$\sum_{a \in \mathsf{Supp}(A)} \Pr[A = a] \log \frac{1}{\Pr[A = a]}$$

Intuitively, this quantifies how much uncertainty we have about variable A. With the definition of entropy, we can further define the relation between various variables. For conditional entropy we have $H(A|B) := H(AB) - H(B)$. Then we are ready to define the relation between different random variables.

▶ **Definition 5** (Mutual Information). The *mutual information* between two random variable A and B, denoted by $I(A; B)$ is defined as

$$I(A; B) := H(A) - H(A|B) = H(B) - H(B|A).$$

The *conditional mutual information* between A and B given C, denoted by $I(A; B|C)$, is defined as

$$I(A; B|C) := H(A|C) - H(A|BC) = H(B|C) - H(B|AC).$$

This gives a measure of how much information does B reveal about A and vice-versa. (when one knows C) Mutual Information is further related to the following distance measure, which will be used throughout the proof.

▶ **Definition 6** (Kullback-Leiber Divergence). Given two probability distributions μ_1 and μ_2 on the same sample space Ω such that $(\forall \omega \in \Omega)(\mu_2(\omega) = 0 \Rightarrow \mu_1(\omega) = 0)$, the *Kullback-Leibler*

Divergence or KL-Divergence in short between μ_1 and μ_2 is defined as (also known as relative entropy)

$$D(\mu_1||\mu_2) = \sum_{\omega \in \Omega} \mu_1(\omega) \log \frac{\mu_1(\omega)}{\mu_2(\omega)}.$$

In particular, the following equality holds between KL-divergence and mutual information.

▶ **Fact 7.** *For random variables A, B and C we have*

$$I(A; B|C) = \mathbb{E}_{b,c} \left[D(A_{bc}||A_c) \right].$$

where A_{bc} is the distribution of random variable A conditioned on $B = b, C = c$ and similarly for A_c.

With the definitions in place, we provide useful properties that will be used in the proof. For mutual information, following facts hold.

▶ **Fact 8** (Chain-rule)**.** *If $I(B; D|C) = 0$, then $I(A; B|C) \leq I(A; B|C, D)$.*

▶ **Fact 9** (Super-Additivity of Mutual Information)**.** *Let C_1, C_2, D, B be random variables such that for every fixing of D, C_1 and C_2 are independent. Then*

$$I(C_1; B|D) + I(C_2; B|D) \leq I(C_1 C_2; B|D).$$

For KL-divergence, we use following properties.

▶ **Fact 10.** *Let P and Q be distributions over a universe $\mathcal{U}$. Suppose $\mathcal{V} \subseteq \mathcal{U}$ is such that $P(\mathcal{V}) = 1$. Then $Q(\mathcal{V}) \geq 2^{-D(P||Q)}$.*

Note that Fact 10 immediately implies non-negativity of KL-divergence between any two distributions. For KL-divergence under conditioning, the following property holds.

▶ **Fact 11** (Additivity of KL-Divergence)**.** *Consider two distributions $P(x, y)$ and $Q(x, y)$. Then*

$$D(P(x,y)||Q(x,y)) = D(P(x)||Q(x)) + \mathbb{E}_{x \sim P} D(P(y|x)||Q(y|x))$$

2.2 Previous Work

In this section, we elaborate previous results on parallel repetition and how the state-of-the-art proof technique for parallel repetition is related to the amount of information necessary for Alice and Bob to win the game with probability greater than $(1 - \varepsilon)$.

Recall that r-times parallel repetition of a two-prover game $\mathcal{G}$ denoted as $\mathcal{G}^r$ is defined as: the referee first samples r-tuple of edges i.e. $\vec{x} = (x_1, \ldots x_r) \in U^r$ and $\vec{y} = (y_1, \ldots, y_r) \in V^r$ with $(x_i, y_i) \in E$ for all $i \in [r]$; Alice and Bob give assignments to all r-coordinates say $f : U^r \to \Sigma^r$ for Alice and $g : V^r \to \Sigma^r$ for Bob; then the referee checks all r-coordinates i.e. return $\bigwedge_{i \in [r]} \pi_{x_i, y_i}(f_i(\vec{x}), g_i(\vec{y}))$.

The first parallel repetition theorem (with exponential decay in value) was proved by Raz [15]:

▶ **Theorem 12** ([15])**.** *Let $\mathcal{G}$ be a game with $\mathbf{val}(\mathcal{G}) = 1 - \varepsilon$ and let s be the size of the alphabet ($|\Sigma|$) of the game. Then $\mathbf{val}(\mathcal{G}^n) \leq (1 - \varepsilon^{32}/2)^{\Omega(n/\log(s))}$.*

This was improved (and simplified) by Holenstein [10]:

▶ **Theorem 13** ([10]). *Let $\mathcal{G}$ be a game with $\mathbf{val}(\mathcal{G}) = 1 - \varepsilon$ and let $\log(s)$ be the answer size of the game. Then $\mathbf{val}(\mathcal{G}^n) \leq (1 - \varepsilon^3/2)^{\Omega(n/\log(s))}$.*

For projection games, $\log(s)$ no longer appears in the exponent. In particular, [14] showed improved bound for projection games.

▶ **Theorem 14** ([14]). *Let $\mathcal{G}$ be a projection game with $\mathbf{val}(\mathcal{G}) = 1 - \varepsilon$. Then $\mathbf{val}(\mathcal{G}^n) \leq (1 - \varepsilon^2/2)^{\Omega(n)}$.*

There is an initiation of work from [1], [18] where they study (though not explicitly stated) a quantity $\mathbf{val}_+(\mathcal{G})$ for unique game $\mathcal{G}$ which is an analytic relaxation of $\mathbf{val}(\mathcal{G})$ that tensorizes exactly, that is $\mathbf{val}_+(\mathcal{G}_1 \otimes \mathcal{G}_2) = \mathbf{val}_+(\mathcal{G}_1) \cdot \mathbf{val}_+(\mathcal{G}_2)$ and captures the amortized value of the game. The analysis of $\mathbf{val}_+$ then resulted in analytic proof of parallel repetition for projection games. [8], [19]. In particular, [8] extended [14] to low-value regime using $\mathbf{val}_+$.

▶ **Theorem 15** ([8]). *Let $\mathcal{G}$ be a projection game with $\mathbf{val}(\mathcal{G}) = \beta$. Then $\mathbf{val}(\mathcal{G}^n) \leq (4\beta)^{n/4}$.*

There has been more work on parallel repetition for various special settings: [2] [17] [19] In particular, there also has been a series of works around parallel repetition of games with entanglement [6, 12, 9, 11, 5]. A recent breakthrough in [21] settled a longstanding open problem on whether the value of "any" games with entanglement actually decays to zero as the number of repetition goes to infinity (at a polynomial rate). It would be interesting to see how our framework relates to games with entanglement.

Throughout this paper, we will use the machinery in the latest parallel repetition proof by [3] where they handled the low-value regime. It should be noted that the same proof works in the high-value regime as well, giving an alternate proof for [10], [14] and [8].

▶ **Theorem 16** ([3]). *Let $\mathcal{G}$ be a game with $\mathbf{val}(\mathcal{G}) = \delta$. Then $\mathbf{val}(\mathcal{G}^n) \leq \delta^{\Omega(n\log(1/\delta)/\log(s))}$. Further if $\mathcal{G}$ is a projection game with $\mathbf{val}(\mathcal{G}) = 1 - \varepsilon$, then $\mathbf{val}(\mathcal{G}^n) \leq (1 - \varepsilon^2)^{\Omega(n)}$.*

The main proof technique used to prove parallel repetition in [3] follows the following general roadmap. First one assumes that the value of the repeated game is higher than the desired bound, and focuses on the event where Alice and Bob win the whole copy. Conditioned on winning, one sets up R, the common hint between Alice and Bob, as a subset of question and answer pairs from other coordinates which can be individually sampled by Alice and Bob (approximately), which is the main technical innovation of [3]. Conditioned on successfully sampling R, Alice and Bob's strategy becomes a "too good to be true" strategy for some coordinate, contradicting the original assumption on the value of the game. For the purpose of having a "hint" between Alice and Bob, we mainly focus on sampling R, avoiding technical issues of constructing R correctly via using [3] as a blackbox. One could view the framework in [3] as following explicit model in Protocol 1.

Protocol 1 Non-Protocol Hint.

---

1. Referee picks a random edge $(x, y) \in E$ as a challenge. Referee then samples r from R_{xy} then transmits r to Alice and Bob

2. Alice, depending on r and x provides an assignment a. Bob analogously answers b depending on r and y.

3. Referee accepts if a and b forms a satisfying assignment for (x, y).

Protocol 1 is indeed not a protocol between Alice and Bob, since the referee samples the hint. Converting this to a protocol between Alice and Bob is our main technical contribution.

For the sake of completeness, we describe how the common hint R is constructed below.

Precise construction of R in [3]

They explicitly construct a "combinatorial" hint $R_{S,G,H,I}$ in a following manner which we restate for completeness. Let n be the number of repetitions, that is the number of coordinates for the game. Then let S, G, H be random subsets of $[n]$ distributed as follows: Let s_h and s_g be random numbers from $\{3n/4 + 1, \ldots, n\}$. Let $\sigma : [n] \to [n]$ be a uniformly random permutation. Set $H = \sigma([s_h])$, $G = \sigma(\{n - s_g + 1, \ldots, n\})$. Let I be a uniformly random element of $G \cap H$. Let l be a random number from $[T]$, where $T < n/2$ is a parameter. Let S be a uniformly random subset of $G \cap H \backslash \{I\}$ of size l. Then define $R_{S,G,H,I}$ to denote the random variable $X_{G \backslash \{I\}} Y_{H \backslash \{I\}} A_S B_S$ where s, g, h, i denote instantiations of the random variables S, G, H, I respectively with A_S denoted Alice's assignment on S-coordinates and respectively for B_S. Then Alice can be thought of getting $X_{[n]}$ and $R_{S,G,H,I}$, while Bob gets $Y_{[n]}$ and $R_{S,G,H,I}$ where the input (x, y) is set to (X_I, Y_I).

2.3 Definitions

Recall that the **Information Cost** of a protocol is defined as $I(\Pi; X|Y) + I(\Pi; Y|X)$ where Π is the transcript of the protocol, X and Y are inputs for Alice and Bob respectively. **Information Complexity** of computing f is then defined as infimum over Π that computes f. Inspired by the definitions from information complexity literature, we define information value of the game as following.

▶ **Definition 17.** The **information value** of the game $\mathcal{G} = (X, Y, E)$ with distribution $\mathcal{D}$ over E is

$$\mathbf{IV}_{\mathcal{D}}^{\varepsilon}(\mathcal{G}) := \inf_{\Pi} \left[I(\Pi; X|Y) + I(\Pi; Y|X) \right]$$

where the infimum is taken over the set of transcripts Π between Alice and Bob which wins $\mathcal{G}$ with probability at least $(1 - \varepsilon)$ with $\varepsilon < 1/2$. X and Y represents Alice and Bob's input respectively.

We remark that if $\mathcal{D}$ is not specified, we assume the distribution to be the uniform distribution over the challenges/edges.

As a straightforward exercise, note that $I(\Pi; X|Y) \leq H(X|Y) \leq H(X) \leq \log n$ where n is the number of vertices in the graph, similarly for $I(\Pi; Y|X)$. Thus for any game $\mathcal{G}$ and $\varepsilon \geq 0$, $\mathbf{IV}_{\mathcal{D}}^{\varepsilon}(\mathcal{G}) \leq O(\log n)$. Thereby, this quantity is strictly bounded. Better bound holds for d-regular graphs since $H(X|Y) \leq \log d$ for regular graphs, similarly for $H(Y|X)$. Any better bounds however, requires lower bound on $H(X||\Pi, Y)$.

3 Main Result

First we state the main technical lemma (Chi-Squared lemma) used in proving the main theorem.

▶ **Lemma 18** (Chi-Squared lemma). *Suppose Alice has access to a distribution P and Bob has access to a distribution Q over $\mathcal{U}$. Suppose further that there exists a common distribution R such that $D(R||P) < \varepsilon$ and $D(R||Q) < \varepsilon$. If $\varepsilon < 1/50$, then there exists a protocol Π that outputs a sample from $\tilde{P}$ with information cost $\gamma\varepsilon$ with $D(R||\tilde{P}) < O(\varepsilon)$ for some constant $\gamma > 0$.*

We append the full proof in the full version of the paper. To see why this lemma is interesting, note that the naive information cost is $D(P||Q)$ which could be infinite in some cases.

Technically speaking, the triangle inequality does not hold for divergence. Applying Pinsker's Inequality to have triangle inequality (in total variation distance) leads to a square loss, resulting in information cost $O(\sqrt{\varepsilon})$, instead of $O(\varepsilon)$. We suspect that there are more applications to this technical lemma. The lemma leads to the following theorem.

▶ **Theorem 19.** *Let $\varepsilon < 1/2$. If* $\mathrm{val}(\mathcal{G}^n) \geq (1-\varepsilon)^n = 2^{-\Omega(\varepsilon n)}$ *and $\mathcal{G}$ is a projection game, then there exists constants $\alpha_1, \alpha_2 > 0$ such that* $\mathbf{IV}^{\alpha_1 \varepsilon}(\mathcal{G}) < \alpha_2 \varepsilon$.

The main intuition of the proof of Theorem 19 is to use the common hint used for dependency breaking step of the parallel repetition, which are answers and questions in the other coordinates, as hints between Alice and Bob. However, the hints are not exactly adequate for our application, since they are sampled by the referee. We use Lemma 18 to convert it to a low information cost protocol.

We also show that the converse of Theorem 19.

▶ **Theorem 20.** *If* $\mathbf{IV}^\varepsilon(\mathcal{G}) < \varepsilon$ *and $\varepsilon < 1/2$, then* $\mathrm{val}(\mathcal{G}^n) \geq 2^{-\Omega(\varepsilon n)} = (1-\varepsilon)^{\Omega(n)}$ *with* $n > 1/\varepsilon$.

The main intuition to Theorem 20 is converting a low information cost protocol (for our application $O(n\varepsilon)$) to a zero-communication protocol as seen in [13]. However, the main theorem from [13] does not suffice for our application. Instead, we use a lemma from [3].

As a corollary, we get a complete description of projection games that obey strong parallel repetition in terms of information value of the game.

▶ **Corollary 21.** *If* $\mathbf{IV}^\varepsilon(\mathcal{G}) > \varepsilon$, *then* $\mathrm{val}(\mathcal{G}^n) < (1-\varepsilon)^{O(n)}$ *and vice versa where $\mathcal{G}$ is a projection game.*

Applying previous parallel repetition result, we indeed get a non-trivial lower bound on the information value of any projection game via [8] and [14].

▶ **Corollary 22.** *For any projection game $\mathcal{G}$ with* $\mathrm{val}(\mathcal{G}) \leq 1 - \varepsilon$, $\mathbf{IV}^{O(\varepsilon^2)}(\mathcal{G}) > \Omega(\varepsilon^2)$.

3.1 Proof of Theorem 19

In this section, we prove Theorem 19 via Chi-squared Lemma (Lemma 18). Recall that $R_{s,g,h,i}$ defined in Section 2.2 is a set of challenges and answers on a random set of coordinates. Set $T = n/4$ as the parameter for $R_{s,g,h,i}$. Then we get the following key lemma from [3].

▶ **Lemma 23** (Lemma 5.6 of [3]). *Suppose $2^{-20} \geq \Pr[W] \geq (1-\varepsilon)^n$. Then there exists a fixing of s, g, h, i such that:*
1. $\mathbb{E}_{x,y \sim \mu} D\left(P_{R_{s,g,h,i}|X_i=x,Y_i=y,W} \| P_{R_{s,g,h,i}|X_i=x,W}\right) \leq O(\varepsilon)$.
2. $\mathbb{E}_{x,y \sim \mu} D\left(P_{R_{s,g,h,i}|X_i=x,Y_i=y,W} \| P_{R_{s,g,h,i}|Y_i=y,W}\right) \leq O(\varepsilon)$.
3. $D\left(\mu \| P_{X_i Y_i}\right) \leq O(\varepsilon)$.
4. $\mathbb{E}_{x,y \sim \mu} \mathbb{E}_{r \sim R_{s,g,h,i}|X_i=x,Y_i=y,W}$
   $D\left(P_{A_i,B_i|X_i=x,Y_i=y,R_{s,g,h,i}=r,W} \| P_{A_i|X_i=x,R_{s,g,h,i}=r,W} \otimes P_{B_i|Y_i=y,R_{s,g,h,i}=r}\right) \leq O(\varepsilon)$.
where μ denotes the distribution $P_{X_i,Y_i|W}$ and P_X stands for the probability distribution of random variable X.

We omit the proof

▶ Remark. The last property does not suffice for our application, since we do not get $r \sim R_{s,g,h,i}|X_i = x, Y_i = y, W$ but a distribution that is $O(\varepsilon)$-away from it in divergence at the end of the protocol given by the Chi-squared lemma which we denote as $\tilde{R}_{s,g,h,i}|X_i =$

$x, Y_i = y, W$. However, we remark that the same proof in [3] indeed gives the property that we want. That is

$$\mathbb{E}_{x,y \sim P_{X_i,Y_i|W}} \mathbb{E}_{r \sim \tilde{R}_{s,g,h,i}|X_i=x,Y_i=y,W}$$
$$D\left(P_{A_i,B_i|X_i=x,Y_i=y,R_{s,g,h,i}=r,W} \| P_{A_i|X_i=x,R_{s,g,h,i}=r,W} \otimes P_{B_i|Y_i=y,R_{s,g,h,i}=r}\right) \leq O(\varepsilon). \quad (1)$$

In particular, note that the distribution of s, g, h, i and the permutation remain the same since Alice and Bob can agree (via public randomness prior to running the protocol) on them prior to sampling the actual question and answer sets (by Alice) This suffices for the proof in [3], specifically Lemma 5.2 and Lemma 5.5.

We also need the following lemma to translate a strategy on $X_i, Y_i|W$ to a strategy on actual distribution X_i, Y_i. Due to space constraints, we attach the proof in Section A.

▶ **Lemma 24.** *Suppose $\mathcal{G}$ with μ as the distribution over the edges achieves $val(\mathcal{G}) = 1 - \varepsilon$. Then consider $\tilde{\mu}$ such that $D(\mu\|\tilde{\mu}) < \varepsilon$. Then $\mathcal{G}$ with $\tilde{\mu}$ as distribution over the edges has value $> 1 - O(\varepsilon)$.*

Now we have all the necessary lemmas to prove Theorem 19.

Proof of Theorem 19. First we construct a low information protocol that wins $\mathcal{G}$ with probability $1 - O(\varepsilon)$ under μ.

We write $P_{R_{s,g,h,i}|X_i=x,Y_i=y,W}$ in the above as $R_{x,y}$ and $P_{R_{s,g,h,i}|X_i=x,W}$, $P_{R_{s,g,h,i}|Y_i=y,W}$ respectively as P_x, Q_y. Consider $\mathcal{S} \subset E$ that satisfies all

- $D\left(R_{x,y}\|P_x\right) \leq \frac{1}{10\gamma}$.
- $D\left(R_{x,y}\|Q_y\right) \leq \frac{1}{10\gamma}$.

where γ is the constant from the Chi-Squared Lemma. Then note that $\mu(\mathcal{S}) > 1 - O(\gamma\varepsilon) = 1 - O(\varepsilon)$ by Markov's inequality. Focus pairs in $\mathcal{S}$. Now applying the protocol given by the Chi-Squared Lemma to pairs in $\mathcal{S}$, we obtain a protocol that samples $r \sim \tilde{R}_{s,g,h,i}|X_i = x, Y_i = y, W$ with information cost at most

$$\mathbb{E}_{x,y \sim \mathcal{S}_{X_i,Y_i|W}} \left[D\left(R_{x,y}\|P_x\right) + D\left(R_{x,y}\|Q_y\right)\right] < O(\varepsilon)$$

where $\mathcal{S}_{X_i,Y_i|W}$ is the distribution over the edges further conditioned on $\mathcal{S}$.

Since $\mathcal{S}$ contributes $1 - O(\varepsilon)$-fraction, (1) implies

$$\mathbb{E}_{x,y \sim \mathcal{S}_{X_i,Y_i|W}} \mathbb{E}_{P_{\tilde{R}_{s,g,h,i},X_i=x,Y_i=y|W}}$$
$$D\left(P_{A_i,B_i|X_i=x,Y_i=y,R_{s,g,h,i}=r,W} \| P_{A_i|X_i=x,R_{s,g,h,i}=r,W} \otimes P_{B_i|Y_i=y,R_{s,g,h,i}=r}\right) \leq O(\varepsilon) \quad (2)$$

At the end of the protocol, Alice and Bob obtain same $r \sim \tilde{R}_{s,g,h,i}|X_i = x, Y_i = y, W$. We now construct an explicit answering strategy for Alice and Bob dependent on r when they get edges distributed according to $\mathcal{S}_{X_i,Y_i|W}$. Ideally Alice and Bob would like to answer according to $P_{A_i,B_i|X_i=x,Y_i=y,\tilde{R}_{s,g,h,i}=r,W}$. This would indeed succeed with probability 1. In other words, if we define $\mathcal{G}_{x,y} = \{(a,b)|V(x,y,a,b) = 1\}$, $P_{A_i,B_i|X_i=x,Y_i=y,\tilde{R}_{s,g,h,i}=r,W}(\mathcal{G}_{x,y}) = 1$ for all $(x,y) \in \mathcal{S}$. However, this is not a valid strategy. There is correlation between A_i and B_i, while for any valid strategy they should be independent given respective input.

Instead, they answer according to $P_{A_i|X_i=x,\tilde{R}_{s,g,h,i}=r,W} \otimes P_{B_i|Y_i=y,\tilde{R}_{s,g,h,i}=r}$. Now, we analyze $P_{A_i|X_i=x,\tilde{R}_{s,g,h,i}=r,W} \otimes P_{B_i|Y_i=y,\tilde{R}_{s,g,h,i}=r}(\mathcal{G}_{x,y})$ i.e. the value of such strategy. Applying Fact 10,

$$P_{A_i|X_i=x,\tilde{R}_{s,g,h,i}=r,W} \otimes P_{B_i|Y_i=y,\tilde{R}_{s,g,h,i}=r}(\mathcal{G}_{x,y})$$

$$\geq 2^{-D\left(P_{A_i,B_i|X_i=x,Y_i=y,R_{s,g,h,i}=r,W}||P_{A_i|X_i=x,R_{s,g,h,i}=r,W}\otimes P_{B_i|Y_i=y,R_{s,g,h,i}=r}\right)}$$

By the convexity of 2^{-x} along with (2), we get the desired bound

$$\mathbb{E}_{x,y\sim\mathcal{S}_{X_i,Y_i|W}}\left[P_{A_i|X_i=x,\tilde{R}_{s,g,h,i}=r,W}\otimes P_{B_i|Y_i=y,\tilde{R}_{s,g,h,i}=r}(\mathcal{G}_{x,y})\right]\geq 2^{-O(\varepsilon)}=1-O(\varepsilon)$$

Since $\mathcal{S}$ contributes $1-O(\varepsilon)$-fraction on μ, this strategy wins with $1-O(\varepsilon)$ probability when the edges are distributed according to $\mu=P_{X_i,Y_i|W}$ as well. Finally, applying Lemma 24 to this strategy with $D\left(P_{X_iY_i|W}||P_{X_iY_i}\right)\leq O(\varepsilon)$, we get the desired claim. ◄

3.2 Proof of Theorem 20

In this section, we give a formal proof of Theorem 20. This involves converting a protocol (between Alice and Bob) with $O(n\varepsilon)$-information cost to a zero-communication protocol with success probability $2^{-O(n\varepsilon)}$. We start by stating the following lemma.

▶ **Lemma 25.** *Suppose Alice has access to distribution P and Bob has access to distribution Q over the universe $\mathcal{U}$. They wish to jointly sample from R where $D(R||P)<\delta$ and $D(R||Q)<\delta$. If $\delta>1$, then there exists a zero-communication protocol such that*
1. *There exists an event E such that $\Pr[E]>2^{-\Omega(\delta)}$ and $\Pr[\pi_a=\pi_b|E]=1$, where π_a and π_b refers to the final output of Alice and Bob respectively. Furthermore, E only depends on the public randomness.*
2. *Given E, consider the set of outputs of π, denoted as $\mathcal{S}$. Then $\mathcal{S}\subseteq\mathsf{Supp}(R)$*

▶ **Claim 26.** *Let W be a subset of universe $\mathcal{U}$. Let A and B be a distribution and A_W be a distribution of A conditioned on picking an element from W. Then if $A(W)>\Omega(D(A||B))$, then*

$$D(A_W||B)<\log(1/A(W))+\frac{D(A||B)}{A(W)}+O\left(\frac{1-A(W)}{A(W)}\right)$$

where $A(W)$ corresponds to the probability of picking an element from W under A.

Proof of Lemma 25 and Claim 26 are appended in Section A. Now, we are ready to prove the main lemma of this section which implies Theorem 20.

▶ **Lemma 27.** *If $\mathbf{IV}^\varepsilon(\mathcal{G})<\varepsilon$, then there exists a zero-communication protocol that achieves $\mathsf{val}(\mathcal{G}^n)>2^{-\Omega(\varepsilon n)}$ where $n>1/\varepsilon$.*

Proof. Note that under this model, Alice's strategy and Bob's strategy are dependent on the transcript $\pi_{x,y}$ as well, instead of just their input in zero-communication model. We denote $\Pi_{x,y}$ as the distribution over the transcript that Alice and Bob will have when they are given input x and y respectively. Now Alice and Bob will try to imitate each other by simulating the other party in zero-communication setting. Let Π_x,Π_y denote the simulated transcript with input x and y respectively. More precisely, $\Pi_x:=\mathbb{E}_{y\sim\mu|x}\Pi_{x,y}$ and $\Pi_y:=\mathbb{E}_{x\sim\mu|y}\Pi_{x,y}$. Further, we introduce the notation $\Pi_{x,y}^W:=\Pi_{x,y}|W$, the distribution of $\Pi_{x,y}$ conditioned on referee accepting what Alice and Bob returns as their answer at the end of the protocol.

Note that from our assumption on the information cost of the protocol, we get the following

$$\mathbb{E}_{(x_i,y_i)\sim\mu}\left[D(\Pi_{x_iy_i}||\Pi_{x_i})\right]=\mathbb{E}_{(x_i,y_i)\sim\mu}\left[\mathbb{E}_{\Pi_{x_iy_i}}\left[\log\frac{\Pr_{\Pi_{x_iy_i}}[\pi]}{\Pr_{\Pi_{x_i}}[\pi]}\right]\right]<\varepsilon. \tag{3}$$

The same inequality holds for Bob's side ($D(\Pi_{x_iy_i}||\Pi_{y_i})$) as well. First we define "good" edges. We say π is a good transcript if the referee accepts what Alice and Bob return after following the transcript π. [1] Then edge (x_i, y_i) is good if it satisfies both

- $A_{x_iy_i}(W) := \Pr_{\pi \sim \Pi_{x_iy_i}}[\pi \text{ is a good transcript}] > 1/2$ (i.e. most sampled transcripts are good);
- $D(\Pi_{x_iy_i}||\Pi_{x_i}) < 1/2$.

We argue that most of the edges are good. Due to our assumption on the value of the game that is,

$$\mathsf{val}(\mathcal{G}) = \mathbb{E}_{(x_i,y_i)\sim\mu}[A_{x_iy_i}(W)] > 1 - \varepsilon,$$

at most 2ε-fraction of (x_i, y_i)'s does not satisfy the first condition. Also due to our divergence condition that is,

$$\mathbb{E}_{(x_i,y_i)\sim\mu}[D(\Pi_{x_iy_i}||\Pi_{x_i})] < \varepsilon$$

at most 2ε-fraction of the edges violate the second condition. Thus all but at most 4ε-fraction of the edges are good. Then we can write

$$\mathbb{E}_{x_iy_i\sim\tilde{\mu}}[D(\Pi_{x_iy_i}||\Pi_{x_i})] < O(\varepsilon) \tag{4}$$

where $\tilde{\mu}$ corresponds to μ conditioned on picking an edge that is good. Also note that without loss of generality, in such regime, one can assume that $1 - A_{x_iy_i}(W) > \Omega(D(\Pi_{x_iy_i}||\Pi_{x_i}))$ for all the edges. For edges that do not satisfy such condition, i.e. $1 - A_{x_iy_i}(W) < O(D(\Pi_{x_iy_i}||\Pi_{x_i}))$, the referee can randomly reject with probability $O(D(\Pi_{x_iy_i}||\Pi_{x_i}))$ to satisfy $1 - A_{x_iy_i}(W) > \Omega(D(\Pi_{x_iy_i}||\Pi_{x_i}))$. Indeed it will add up the rejection probability, but by at most $D(\Pi_{x_iy_i}||\Pi_{x_i})$ which indeed is good enough for application in our regime, since it is in expectation at most $O(\varepsilon)$. If (x_i, y_i) is indeed a good edge, applying Claim 26,

$$D(\Pi_{x_iy_i}^W||\Pi_{x_i}) < \log(1/A_{x_iy_i}(W)) + \frac{D(\Pi_{x_iy_i}||\Pi_{x_i})}{A_{x_iy_i}(W)} + O\left(\frac{1 - A_{x_iy_i}(W)}{A_{x_iy_i}(W)}\right)$$

$$< \log(1/A_{x_iy_i}(W)) + 2D(\Pi_{x_iy_i}||\Pi_{x_i}) + O\left(\frac{1 - A_{x_iy_i}(W)}{A_{x_iy_i}(W)}\right)$$

where the second inequality holds since $A_{x_iy_i}(W) > 1/2$, $D(\Pi_{x_iy_i}||\Pi_{x_i}) < 1$, and our assumption that $1 - A_{x_iy_i}(W) > \Omega(D(\Pi_{x_iy_i}||\Pi_{x_i}))$ for all edges. Then

$$\mathbb{E}_{\tilde{\mu}}[D(\Pi_{x_iy_i}^W||\Pi_{x_i})] < \mathbb{E}_{\tilde{\mu}}[\log(1/A_{x_iy_i}(W))] + 2\mathbb{E}_{\tilde{\mu}}[D(\Pi_{x_iy_i}||\Pi_{x_i})] + \mathbb{E}_{\tilde{\mu}}\left[O\left(\frac{1 - A_{x_iy_i}(W)}{A_{x_iy_i}(W)}\right)\right]$$

$$< O(\varepsilon) + O(\varepsilon) + O(\varepsilon) < O(\varepsilon)$$

where the second inequality holds by Jensen's inequality on log and since ε is small enough.

Now consider taking n-copies of the game. In particular, we focus on $(\vec{x}, \vec{y}) \sim \tilde{\mu}^{\otimes n}$, that is all edges are "good" edges. Then observe that

$$\mathbb{E}_{(\vec{x},\vec{y})\sim\tilde{\mu}^{\otimes n}}\left[D(\bigotimes_{i\in[n]}\Pi_{x_iy_i}^W||\bigotimes_{i\in[n]}\Pi_{x_i})\right] = \mathbb{E}_{(\vec{x},\vec{y})\sim\tilde{\mu}^{\otimes n}}\left[\sum_{i\in[n]}D(\Pi_{x_iy_i}^W||\Pi_{x_i})\right] < O(n\varepsilon)$$

---

[1] π does not necessarily depend just on the input. It can depend on private randomness as well. But this is not crucial to the proof as we argue on sampling the transcript conditioned on the edges.

For a randomly picked $(\vec{x}, \vec{y}) \sim \tilde{\mu}^{\otimes n}$, the divergence in consideration is indeed low with high probability by Markov. That is

$$\Pr_{(\vec{x}, \vec{y})} \left[\sum_{i \in [n]} D(\Pi^W_{x_i y_i} || \Pi_{x_i}) > \alpha n \varepsilon \right] \leq O(1/\alpha) \tag{5}$$

Similarly, we get

$$\Pr_{(\vec{x}, \vec{y})} \left[\sum_{i \in [n]} D(\Pi^W_{x_i y_i} || \Pi_{y_i}) > \alpha n \varepsilon \right] \leq O(1/\alpha) \tag{6}$$

Now, we consider the particular set of vectors $(\vec{x}, \vec{y})$ that satisfy
- $\forall i \in [n]$, (x_i, y_i) is a "good" edge.
- $D(\Pi^W_{\vec{x}, \vec{y}} || \Pi_{\vec{x}}) = \sum_{i \in [n]} D(\Pi^W_{x_i y_i} || \Pi_{x_i}) \leq K n \varepsilon$ and $D(\Pi^W_{\vec{x}, \vec{y}} || \Pi_{\vec{y}}) = \sum_{i \in [n]} D(\Pi^W_{x_i y_i} || \Pi_{y_i}) \leq K n \varepsilon$

which we denote as "good" vectors.

If $(\vec{x}, \vec{y}) \sim \mathcal{U}^{\otimes n}$, since (x_i, y_i) is good with probability at least $(1 - 2\varepsilon)$, all the coordinates are good with at least $(1 - 2\varepsilon)^n$ probability. If all the coordinates are good, by (5) and (6) and picking appropriately large K, $\Omega(1)$-fraction of such edges satisfy the second condition as well. In total, $2^{-\Omega(\varepsilon n)}$-fraction of edges satisfies both conditions, since we assume $n > 1/\varepsilon$.

Now we apply Lemma 25 to "good" vectors to complete the proof. In particular, Lemma 25 gives a zero-communication sampling protocol for transcript where Alice and Bob gets a matching transcript from $\mathsf{Supp}(\Pi^W_{\vec{x}, \vec{y}})$ with probability at least $2^{-O(\varepsilon n)}$. Thus for $2^{-O(\varepsilon n)}$-fraction of the edges, we get a zero-communication strategy that wins with probability at least $2^{-O(\varepsilon n)}$, thus $\mathsf{val}(\mathcal{G}^n) > 2^{-O(\varepsilon n)}$. ◀

───── **References** ─────

1 Boaz Barak, Ishay Haviv, Moritz Hardt, Anup Rao, Oded Regev, and David Steurer. Rounding parallel repetitions of unique games. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science.* IEEE Computer Society, 2008.

2 Boaz Barak, Anup Rao, Ran Raz, Ricky Rosen, and Ronen Shaltiel. Strong parallel repetition theorem for free projection games. *RANDOM*, 2009.

3 Mark Braverman and Ankit Garg. Small value parallel repetition for general games. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 335–340. ACM, 2015. doi:10.1145/2746539.2746565.

4 Mark Braverman and Anup Rao. Information equals amortized communication. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 748–757. IEEE, 2011.

5 André Chailloux and Giannicola Scarpa. Parallel repetition of entangled games with exponential decay via the superposed information cost. *41st International Colloquium on Automata, Languages and Programming*, 2014.

6 Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum xor proof systems. *Journal of Computational Complexity*, 17(2):282–299, May 2008.

7   Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley series in telecommunications. J. Wiley and Sons, New York, 1991.

8   Irit Dinur and David Steurer. Analytical approach to parallel repetition. *46th Annual Symposium on the Theory of Computing*, 2014.

9   Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. *IEEE Conference on Computational Complexity*, 2014.

10 Thomas Holenstein. Parallel repetition: Simplifications and the no-signaling case. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, 2007.

11 Rahul Jain, Attila Pereszlényi, and Penghui Yao. A parallel repetition theorem for entangled two-player one-round games under product distributions. *IEEE Conference on Computational Complexity*, 2014.

12 Julia Kempe and Thomas Vidick. Parallel repetition of entangled games. *43rd annual ACM symposium on Theory of computing*, 2011.

13 Iordanis Kerenidis, Sophie Laplante, Virginie Lerays, Jérémie Roland, and David Xiao. Lower bounds on information complexity via zero-communication protocols and applications. *CoRR*, abs/1204.1505, 2012. URL: `http://arxiv.org/abs/1204.1505`.

14 Anup Rao. Parallel repetition in projection games and a concentration bound. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, 2008.

15 Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, jun 1998. Prelim version in STOC'95.

16 Ran Raz. A counterexample to strong parallel repetition. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 2008.

17 Ran Raz and Ricky Rosen. A strong parallel repetition theorem for projection games on ex- panders. *IEEE Conference on Computational Complexity*, pages 247–257, 2012.

18 David Steurer. Improved rounding for parallel repeated unique games. In *Proceedings of the 13th International Conference on Approximation, and 14 the International Conference on Randomization, and Combinatorial Optimization: Algorithms and Techniques*, APPROX/RANDOM'10, pages 724–737, Berlin, Heidelberg, 2010. Springer-Verlag. URL: `http://dl.acm.org/citation.cfm?id=1886521.1886577`.

19 Madhur Tulsiani, John Wright, and Yuan Zhou. Optimal strong parallel repetition for projection games on low threshold rank graphs. *ICALP*, 2014.

20 Amir Yehudayoff. Pointer chasing via triangular discrimination. *TR16-151*, 2016.

21 Henry Yuen. A parallel repetition theorem for all entangled games. *CoRR*, abs/1604.04340, 2016. `arXiv:1604.04340`.

A   Omitted Proof from Section 3

Proof of Lemma 24. Consider the set of edges $\mathcal{S}$ that are satisfied under μ. Denote μ conditioned on being inside $\mathcal{S}$ as ν. We show that $D(\nu||\tilde{\mu}) < O(\varepsilon)$, which indeed implies $\tilde{\mu}(\mathcal{S}) < 1 - O(\varepsilon)$ by Fact 10. But indeed Claim 26 implies $D(\nu||\tilde{\mu}) < O(\varepsilon)$.   ◀

Proof of Lemma 25. The proof follows from Lemma 4.5 of [3]. To argue that the second condition is indeed met, recall that the protocol in [3] is the following:

Protocol 2 Protocol for sampling a transcript (π, q)

- Using shared randomness, get uniformly random samples from $\Pi \times [0, 1]$, which we denote by $\{(\pi_i, q_i)\}_{i=0}^{\infty}$.
- Alice outputs the first π_a that satisfies $q_a < \frac{P(\pi_a)}{\delta}$.
- Bob outputs the first π_b that satisfies $q_b < \frac{Q(\pi_b)}{\delta}$.

Let $\mathcal{A} := \{(\pi, q)|q < P(\pi)/\delta\}$, $\mathcal{B} := \{(\pi, q)|q < Q(\pi)/\delta\}$ and $\mathcal{C} := \{(\pi, q)|q < R(\pi)\}$. We define event E as first dart in $\mathcal{A} \cup \mathcal{B}$ being inside $\mathcal{A} \cap \mathcal{B} \cap \mathcal{C}$. Note that Lemma 4.5. of [3] exactly gives that $\Pr[E] \geq 2^{-\Omega(\delta)}$, therefore the first condition must hold. Furthermore, $\mathcal{S}$ is indeed included in $\mathsf{Supp}(R)$ since for such event it must be the case that $R(\pi) > 0$. Therefore, E satisfies the second condition as well. ◀

Proof of Claim 26. For brevity denote W as the set of x with $W(x) = 1$, $D(A||B) = \delta_0$ and $A(\overline{W}) = \delta_1$. First, we show that $B(\overline{W}) < O(D(A||B)) = O(\delta_0)$. Note that $D(A||B)$ can be written as

$$D(A||B) = \sum_{x \in W} A(x) \log \frac{A(x)}{B(x)} + \sum_{x \notin W} A(x) \log \frac{A(x)}{B(x)} = \delta_0 \tag{7}$$

Applying log-sum, we get

$$A(W) \log \frac{A(W)}{B(W)} + A(\overline{W}) \log \frac{A(\overline{W})}{B(\overline{W})} \leq \delta_0$$

Assume for contradiction that $B(\overline{W}) > KA(\overline{W})$, where K is some parameter that we will setup later for contradiction. And put $B(\overline{W}) = \alpha A(\overline{W})$ for $\alpha > K$. Since $A(W) = 1 - \delta_1$ by our assumption, substituting the terms we get

$$(1 - \delta_1) \log \frac{1 - \delta_1}{1 - \alpha\delta_1} + \delta_1 \log \frac{1}{\alpha} = (1 - \delta_1) \log \frac{1}{1 - \frac{(\alpha-1)\delta_1}{1-\delta_1}} + \delta_1 \log \frac{1}{\alpha}$$

Suppose for now that $\alpha < \frac{1}{2\delta_1} + \frac{1}{2}$. Then note that $\log(1/(1-x)) = \Omega(x)$. Applying this fact to the first term, we get

$$\delta_0 > (1 - \delta_1) \log \frac{1 - \delta_1}{1 - \alpha\delta_1} + \delta_1 \log \frac{1}{\alpha} > \Omega((\alpha - 1 - \log\alpha)\delta_1) > \Omega(\alpha\delta_0).$$

where the last inequality holds due to our assumption that $\delta_1 > \delta_0$. Picking an appropriately large constant $\alpha = O(1)$, we get a contradiction. Instead if $\alpha > \frac{1}{2\delta_1} + \frac{1}{2}$, then

$$(1 - \delta_1) \log \frac{1 - \delta_1}{1 - \alpha\delta_1} + \delta_1 \log \frac{1}{\alpha} > (1 - \delta_1) \log 2(1 - \delta_1) + \delta_1 \log 2\delta_1 = H(\delta_1) + 1 > \delta_0$$

which indeed is a contradiction since $\delta_0 < 1$ and $H(\delta_1) \geq 0$. Thus $B(\overline{W}) < O(\delta_1) = K_0\delta_1$.

To bound $D(A_W||B)$, we first bound $\sum_{x \in W} A(x) \log \frac{A(x)}{B(x)}$. Note that via our bound on $B(\overline{W})$ implies

$$\sum_{x \notin W} A(x) \log \frac{A(x)}{B(x)} > A(\overline{W}) \log \frac{A(\overline{W})}{B(\overline{W})} > \delta_1 \log \frac{1}{K_0}$$

Applying this to (7), we have

$$\sum_{x \in W} A(x) \log \frac{A(x)}{B(x)} = \delta_0 - \sum_{x \notin W} A(x) \log \frac{A(x)}{B(x)} < \delta_0 - \delta_1 \log \frac{1}{K_0}$$

Now we use the above fact to bound $D(A_W||B)$.

$$D(A_W||B) = \sum_{x} a_W(x) \log \frac{a_W(x)}{b(x)} = \sum_{x \in W} a_W(x) \log \frac{a_W(x)}{a(x)} \frac{a(x)}{b(x)}$$

$$= D(A_W||A) + \sum_{x \in W} a_W(x) \log \frac{a(x)}{b(x)} \leq \log(1/A(W)) + \sum_{x \in W} \frac{a(x)}{A(W)} \log \frac{a(x)}{b(x)}$$

$$< \log(1/A(W)) + \frac{D(A||B)}{A(W)} + \frac{1 - A(W)}{A(W)} \log K_0$$

$$< \log(1/A(W)) + \frac{D(A||B)}{A(W)} + O\left(\frac{1 - A(W)}{A(W)}\right)$$

which is indeed the statement of the claim.   ◀