

Bit-Security Preserving Hardness Amplification

Shun Watanabe¹ and Kenji Yasunaga²

¹ Tokyo University of Agriculture and Technology, Japan shunwata@cc.tuat.ac.jp

² Institute of Science Tokyo, Japan yasunaga@c.titech.ac.jp

Abstract. Hardness amplification is one of the important reduction techniques in cryptography, and it has been extensively studied in the literature. The standard XOR lemma known in the literature evaluates the hardness in terms of the probability of correct prediction; the hardness is amplified from mildly hard (close to 1) to very hard $1/2+\varepsilon$ by inducing ε^2 multiplicative decrease of the circuit size. Translating such a statement in terms of the bit-security framework introduced by Micciancio-Walter (EUROCRYPT 2018) and Watanabe-Yasunaga (ASIACRYPT 2021), it may cause a bit-security loss of $\log(1/\varepsilon)$. To resolve this issue, we derive a new variant of the XOR lemma in terms of the Rényi advantage, which directly characterizes the bit security. In the course of proving this result, we prove a new variant of the hardcore lemma in terms of the conditional squared advantage; our proof uses a boosting algorithm that may output the \perp symbol in addition to 0 and 1, which may be of independent interest.

1 Introduction

In modern cryptography, cryptographic primitives are usually proposed with security proofs. When proving the security of a primitive under some hardness assumption, we show a *reduction* that solves a hard problem by assuming the existence of an adversary attacking the primitive. If the reduction requires much more computational cost than the assumed adversary, we need a stronger hardness assumption to achieve a target level of security. Thus, *tight* reductions of security proofs are desirable for the efficient use of cryptographic primitives.

A recent approach of *concrete security* reveals quantities related to security reductions. Suppose we want to prove the security of primitive Q by assuming the security of primitive P (or the hardness of some problem). Typically, we show that for any adversary B of primitive Q with running time $t_B(n)$ and advantage $\varepsilon_B(n)$, there is an adversary A of primitive P such that the running time $t_A(n)$ and the advantage $\varepsilon_A(n)$ satisfy $t_A(n) \leq \phi(t_B(n))$ and $\varepsilon_A(n) \geq \psi(\varepsilon_B(n))$ for some functions ϕ and ψ . Here, n is a security parameter, and a reduction is a construction of A out of B . We may understand the tightness of the reduction by specifying two functions, ϕ and ψ . We prefer smaller $t_A(n)$ and larger $\varepsilon_A(n)$ for tight reductions. Thus, it is tempted to combine the two quantities as $t_A(n)/\varepsilon_A(n)$ and achieve the value as small as $t_B(n)/\varepsilon_B(n)$ of

adversary B . Namely, we want the loss function

$$L(n) = \frac{t_A(n)}{\varepsilon_A(n)} \cdot \frac{\varepsilon_B(n)}{t_B(n)}$$

to be as small as possible.

The above way of quantifying security loss has been used in cryptographic literature. In [29], the quantity of $t_A(n)/\varepsilon_A(n)$ was used to define the security of primitives. The same treatment has been employed in the literature [4, 2, 33, 6]. For *search* primitives such as one-way functions and signature schemes, the advantage $\varepsilon_A(n)$ is simply defined as the adversary's success probability. For *decision* primitives such as pseudorandom generators and encryption schemes, it is usually defined as the gap between two probabilities, which we want to make as small as possible. This treatment of defining advantages has been standard in the cryptography community. In the literature listed above, the advantage $\varepsilon_A(n)$ was defined in this way for analyzing the quantity $t_A(n)/\varepsilon_A(n)$.

In [12], Goldreich noted that Levin suggested using another quantity $t_A(n)/\varepsilon_A^2(n)$, called *work*, for decision primitives. The reason is that if the gap of two probabilities is $\varepsilon_A(n)$, we need to repeat the experiment (security game) for $\mathcal{O}(1/\varepsilon_A(n)^2)$ times to amplify it to a constant, say $2/3$. This was also suggested in [18]. However, the use of this quantity was not justified well at that time.

1.1 Bit security

Micciancio and Walter [32] initiated a theoretical study for quantifying the security level of primitives, referring to it as *bit security*. They proposed using another notion of advantage, which we call *conditional squared (CS) advantage*, for evaluating the decision primitives. It is defined as

$$\text{Adv}_A^{\text{CS}}(n) = \Pr(Y \neq \perp)(2\Pr(Y = U|Y \neq \perp) - 1)^2,$$

where Y is the random variable representing the adversary A 's output and U is a (randomly chosen) secret bit of the decision game. They defined the bit security as the quantity of

$$\min_A \log_2 \frac{t_A(n)}{\text{Adv}_A^{\text{CS}}(n)}.$$

Their notion elegantly resolved paradoxical situations in pseudorandom generators and approximate samplers (For the former case, we elaborate later in this section). In [28], the notion of [32] was extended for capturing both computational and statistical parameters. Watanabe and Yasunaga [39] defined bit security with an *operational meaning* to justify formalizing the security level of primitives. Roughly speaking, for a given adversary of a primitive, the bit security in [39] is defined as the number of invocations of that adversary to break the primitive with probability close to 1; it was shown in [39] that the *Rényi advantage* characterizes the number of invocations. Specifically, it is given by

$$\text{Adv}_A^{\text{Rényi}}(n) = -2 \ln \sum_a \sqrt{\Pr(Y = a|U = 0) \Pr(Y = a|U = 1)},$$

(See Appendix A for more explanations of the two frameworks [32, 39].) The follow-up work [40] demonstrated that the two advantages of [32, 39] are essentially equivalent;³ owing to this result, we can use the CS advantage and the Rényi advantage almost interchangeably, and we will use both the advantages in this paper as well.

1.2 Hardness amplification

In this work, based on the recent advances in the notion of bit security (or the quantity $t_A(n)/\varepsilon_A(n)$), we focus on a basic problem of *hardness amplification* [14] of Boolean functions. A Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is said to be *mildly hard (unpredictable)* if every polynomial-time algorithm fails to compute f on a δ -fraction of input $x \in \{0, 1\}^n$ for a noticeable δ . The task of hardness amplification is to convert f into another function f' so that f' is *strongly hard* in the sense that every polynomial-time algorithm fails to compute f' on a $(1/2 - \varepsilon)$ -fraction of input. The most well-known technique is Yao's XOR lemma; $f'(x_1, \dots, x_k) = f(x_1) \oplus \dots \oplus f(x_k)$ for $x_i \in \{0, 1\}^n$. In cryptography, Yao's XOR lemma (and its variants) have been used for amplifying the security of pseudorandom generators (PRGs) and pseudorandom functions (PRFs) [34, 7, 31]. In the framework of bit security, hardness amplification is to reduce the advantage $\text{Adv}_{A,f}(n)(= 1/2 - \delta)$ to $\text{Adv}_{B,f'}(n)(= \varepsilon)$, where $\text{Adv}_{A,f}(n)$ is the advantage of adversary A predicting f over random guessing.

In the two bit-security frameworks of [32, 39], a decision game is formalized such that an adversary tries to guess the secret bit $u \in \{0, 1\}$ by playing the game. Thus, we can write a decision game as $G = (G_0, G_1)$, where a secret bit u is initially chosen uniformly at random, and the adversary plays G_u for guessing u . The hardness of predicting a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be captured by the game G_f as follows; first, $x \in \{0, 1\}^n$ and $\sigma \in \{0, 1\}$ are chosen uniformly at random. Then, the adversary receives $(x, f(x))$ when $u = 0$, and (x, σ) when $u = 1$, and outputs a symbol in $\{0, 1, \perp\}$.⁴

When employing the framework of [39, 40], the bit security of game $G_f = (G_0, G_1)$ against adversaries with computational cost s can be approximated as

$$\text{BS}_s(G_f) = \log \min_{A \text{ with cost } s} \frac{s}{\text{Adv}_{A,G_f}^{\text{Rényi}}}, \quad (1)$$

where $\text{Adv}_{A,G_f}^{\text{Rényi}} = D_{1/2}(A_0 \| A_1)$ is the Rényi advantage of A in game G_f , $D_{1/2}(\cdot \| \cdot)$ is the Rényi divergence of order $1/2$, and A_u is the output distribution of A when playing G_u .

With the notions of bit security, hardness amplification is the task of converting f into f' such that $\max_B \text{Adv}_{B,G_{f'}}^{\text{Rényi}}$ is much smaller than $\max_A \text{Adv}_{A,G_f}^{\text{Rényi}}$,

³ Generally, the CS advantage is bounded above by the Rényi advantage. While the CS advantage may take a much smaller value in some cases, the CS advantage can be increased to the level of the Rényi advantage by modifying adversaries appropriately.

⁴ The symbol \perp indicates that the adversary gives up the prediction.

where A and B are taken over adversaries with costs s and s' , respectively. We want the following loss function

$$L^{\text{amp}}(n) = \frac{s \cdot \max_B \text{Adv}_{B, G_{f'}}^{\text{Renyi}}}{s' \cdot \max_A \text{Adv}_{A, G_f}^{\text{Renyi}}}$$

to be as small as possible. Ideally, we want to achieve $L^{\text{amp}}(n) = \mathcal{O}(1)$.

Balanced/unbalanced adversaries. The most efficient reductions of the XOR lemma until now were given in [25, 3] using boosting versions of *hardcore lemmas* [23]. They guarantee $s' = \Omega(\varepsilon^2 / \log(1/\delta)) \cdot s$, where $\max_A \Pr(A(x) = f(x)) = 1 - \delta$ and $\max_B \Pr(B(x) = f'(x)) = 1/2 + \varepsilon$, where the maxima are taken over algorithms with cost s and s' , respectively. Such a *predictor* A with $\Pr(A(x) = f(x)) = 1 - \delta$ can be easily converted to a *distinguisher* A' with the same cost such that $\text{Adv}_{A', G_f}^{\text{TV}} = d_{\text{TV}}(A'_0, A'_1) = (1 - \delta) - 1/2 = 1/2 - \delta$, where $d_{\text{TV}}(\cdot, \cdot)$ is the total variation distance. For any adversary A of game $G_f = (G_0, G_1)$, it holds that

$$\left(\text{Adv}_{A, G_f}^{\text{TV}}\right)^2 \leq \text{Adv}_{A, G_f}^{\text{Renyi}} \leq \mathcal{O}\left(\text{Adv}_{A, G_f}^{\text{TV}}\right).$$

Thus, if $\text{Adv}_{A, G_f}^{\text{Renyi}} \approx \left(\text{Adv}_{A, G_f}^{\text{TV}}\right)^2$ holds for every adversary A , the XOR lemma in [25, 3] gives

$$L^{\text{amp}}(n) = \mathcal{O}\left(\frac{\log(1/\delta)}{\varepsilon^2} \cdot \frac{\varepsilon^2}{(1/2 - \delta)^2}\right) = \mathcal{O}(\log(1/\delta)),$$

meaning that the reduction seems to be optimal. Indeed, as observed in [39], $\text{Adv}_{A, G_f}^{\text{Renyi}} \approx \left(\text{Adv}_{A, G_f}^{\text{TV}}\right)^2$ holds for *balanced* adversaries, who output every value with probability $\Omega(1)$. However, generally, we have $\text{Adv}_{A, G_f}^{\text{Renyi}} = \mathcal{O}\left(\text{Adv}_{A, G_f}^{\text{TV}}\right)$. Thus, the reductions in [25, 3] imply that

$$L^{\text{amp}}(n) = \mathcal{O}\left(\frac{\log(1/\delta)}{\varepsilon^2} \cdot \frac{\varepsilon}{1/2 - \delta}\right) = \mathcal{O}\left(\frac{\log(1/\delta)}{\varepsilon}\right), \quad (2)$$

which does not seem to be optimal.

It should be emphasized that the above-mentioned difference between the balanced and unbalanced adversaries is crucial in the bit security frameworks of [32, 39]; the difference can be well explained by the following example first posed in [32] and further elaborated in [40]. Let us consider the decision game to distinguish between the pseudorandom number generator (PRG) and the true random number generator (TRG): the outcome (y, z) of PRG consists of the image $y = F(r)$ of a one-way permutation F over $\{0, 1\}^n$ and its hard-core predicate $z = h(r)$; the outcome (y, z) of TRG consists of $y = F(r)$ and a random bit $z = \sigma$ that is independent of the seed r . For this game, we can consider the following two possible attacks:

Table 1. Advantages for the linear test attack and the inversion attack

Attacks	Adv^{TV}	Adv^{CS}	$\text{Adv}^{\text{Rényi}}$
Linear test attack (balanced)			
$A_0 = (1/2 + \varepsilon, 1/2 - \varepsilon)$	ε	ε^2	$\Theta(\varepsilon^2)$
$A_1 = (1/2, 1/2)$			
Inversion attack (unbalanced)			
$A_0 = (\varepsilon, 0, 1 - \varepsilon)$	$\varepsilon/2$	$\varepsilon/2$	$\Theta(\varepsilon)$
$A_1 = (\varepsilon/2, \varepsilon/2, 1 - \varepsilon)$			

1. *Linear test attack:* For a prescribed binary vector v of length $n + 1$, the adversary computes the inner product of v and (y, z) ; if the outcome is 0, the adversary outputs 0 (PRG); and outputs 1 (TRG) otherwise. For such an attack, the output distribution A_u of the adversary A given $u \in \{0, 1\}$ ($u = 0$ represent PRG and $u = 1$ represent TRG) are $A_0 = (1/2 + \varepsilon, 1/2 - \varepsilon)$ and $A_1 = (1/2, 1/2)$ for some bias ε , where $A_u = (p_0, p_1)$ means that $\Pr[A = 0 \mid u] = p_0$ and $\Pr[A = 1 \mid u] = p_1$.
2. *Inversion attack:* First, the adversary tries to invert the one-way permutation, which will succeed with probability ε . If the inversion is unsuccessful, the adversary outputs \perp ; if the inversion is successful and $h(r)$ coincides with z , then the adversary outputs 0 (PRG); otherwise (if the inversion is successful but $h(r) \neq z$), then the adversary outputs 1 (TRG). For such an attack, the output distribution of the adversary A given $u \in \{0, 1\}$ consists of $A_0 = (\varepsilon, 0, 1 - \varepsilon)$ and $A_1 = (\varepsilon/2, \varepsilon/2, 1 - \varepsilon)$, where the third coordinate represents the probability that $\Pr[A = \perp \mid u]$.

The standard advantage Adv^{TV} , the CS advantage Adv^{CS} , and the Rényi advantage for the linear test attack and the inversion attack are summarized in Table 1. Note that, in the inversion attack, the adversary outputs \perp with significant probability $1 - \varepsilon$, and it has a completely different nature compared to the linear test attack. However, the standard advantages of the two attacks are almost the same, and it does not capture the different nature of the attacks. One of the notable features of the bit security in [32, 39] is that the difference between these attacks is reflected as different orders of the advantages. Thus, it is necessary to extend the hardness amplification to the CS advantage or the Rényi advantage, which involves several challenges mentioned in the next section.

1.3 Our results

In this work, in order to evaluate the bit security of hardness amplification directly, we derive a new variant of the XOR lemma in terms of the Rényi advantage. Roughly, our XOR lemma (Theorem 1) claims that if a function is mildly hard in the sense that $\Pr(A(x) = f(x)) \leq 1 - \delta$ for any adversary A of size s , then the Rényi advantage of the XOR function f' satisfies $\text{Adv}_{B, G_{f'}}^{\text{Rényi}} \leq \varepsilon$

for any adversary B of size $s' = \Omega(\varepsilon/\ln(1/\delta)) \cdot s$.⁵ This implies that the loss of the reduction is

$$L^{\text{amp}}(n) = \mathcal{O}\left(\frac{\log(1/\delta)}{\varepsilon} \cdot \frac{\varepsilon}{1/2 - \delta}\right) = \mathcal{O}(\log(1/\delta)),$$

which improves upon the loss in (2) by the factor of $1/\varepsilon$ for general adversaries.

To derive our XOR lemma for the Rényi advantage, we prove a new variant of the hardcore lemma, originally proved by Impagliazzo [23]. Our hardcore lemma is stated in terms of the CS advantage. Then, by using the connection between the CS advantage and the Rényi advantage in [40], we prove the XOR lemma via the hardcore lemma.

To prove our hardcore lemma, we analyze the performance of the boosting algorithm such that weak learners may output the \perp symbol in addition to 0 and 1. Our main technical contribution in this paper is characterizing the performance of the boosting algorithm with \perp in terms of the CS advantage.

Ideas and techniques. Roughly, our hardcore lemma (Lemma 1) claims that if a function is mildly hard in the sense that $\Pr(A(x) = f(x)) \leq 1 - \delta$ for any adversary A of size s , then there exists a hardcore distribution H with density δ such that the CS advantage of predicting f for inputs distributed according to H satisfies $\text{Adv}_{B,f|H}^{\text{CS}} \leq \varepsilon$ for any adversary B of size $s' = \Omega(\varepsilon/\ln(1/\delta)) \cdot s$. Compared to the standard hardcore lemma known in the literature, which involves the decrease of circuit size by the factor of $\Omega(\tilde{\varepsilon}^2/\ln(1/\delta))$ for the standard advantage $\tilde{\varepsilon}$, the decrease of circuit size in our hardcore lemma is by the factor of $\Omega(\varepsilon/\ln(1/\delta))$ for the CS advantage ε . For balanced adversaries, since the CS advantage behaves as $\varepsilon \simeq \tilde{\varepsilon}^2$, our hardcore lemma is essentially the same as the standard hardcore lemma; however, for unbalanced adversaries, since the CS advantage behaves as $\varepsilon \simeq \tilde{\varepsilon}$, our hardcore lemma improves the decrease of circuit size upon the standard hardcore lemma.

A utility of the CS advantage in the context of the Goldreich-Levin (GL) algorithm has been reported by Hast in [17]; he proposed a modified version of the GL algorithm by taking into account adversaries that may output \perp when predicting the hardcore bit and characterized the performance of such a GL algorithm in terms of the CS advantage. His algorithm was used in [40] to prove the tightness of the GL theorem. In this work, we use the utility of outputting \perp in a hardcore lemma to provide a tight reduction of hardness amplification. Since the performance of the GL algorithm can be improved by taking into account the effect of \perp , it is natural to seek a hardcore lemma that takes into account the effect of \perp as well. However, such a hardcore lemma has not been proposed in the literature. In fact, it is not immediately clear how \perp can be incorporated into a hardcore lemma. In the case of the GL algorithm, the algorithm can

⁵ More precisely, this statement assumes that the cost of a weighted majority gate can be ignored. Indeed, if $s = \omega(\log(1/\delta)/\varepsilon^2)$, the cost is negligible; See Remark 1 for discussion.

be interpreted as a list decoding of the Hadamard code. Then, it is natural to consider \perp as an erasure of error correcting code. On the other hand, the hardcore lemma is interpreted as the boosting algorithm in the learning theory. Unlike the error correcting code, the role of \perp in the boosting is not so clear. In fact, as we will discuss below, the usage of \perp in our boosting algorithm is subtle, and the effect of \perp only shows up after a judicious analysis of the algorithm.

Readers might wonder why the XOR lemma is stated and proved in terms of the Rényi advantage while the hardcore lemma is stated and proved in terms of the CS advantage. When we prove the standard XOR lemma using the standard hardcore lemma, we decompose the uniform distribution of inputs into a weighted sum of the “easy part” and the “hardcore part.” Since the probability of correct prediction is an affine function of the distribution of inputs, it can also be decomposed as the corresponding weighted sum of the probability of correct prediction for the easy and the hardcore parts. A main obstacle to using such a decomposition argument in our setting is that neither the CS advantage nor the Rényi advantage are affine functions of input distributions. Fortunately, by using the joint convexity of the Rényi divergence in a judicious manner, we can go through the decomposition argument for the Rényi advantage; it is not clear if a similar trick can be applied for the CS advantage, which is the reason why we consider the XOR lemma for the Rényi advantage.

On the other hand, while the labeling of the adversary’s output symbols is irrelevant to the Rényi advantage, the abort symbol \perp has a special meaning in the CS advantage. Thus, the CS advantage is more suitable to be used when we prove the hardcore lemma by using the boosting with \perp .

Heuristically, a gist of the standard boosting algorithm (i.e., weak learners without \perp) is to update the input distribution iteratively as follows: for a given weak learner of the current round, we update the distribution of inputs so that the probabilities of symbols correctly predicted by the current weak learner are decreased, and the probabilities of symbols erroneously predicted by the current weak learner are increased; we choose the multiplicative weight (for decrease/increase) of the update so that the advantage of the weak learner of the previous round is zero for the updated distribution. In our setting, weak learners may output \perp in addition to the binary outputs. Then, a natural question is whether we should increase or decrease the probabilities of symbols when the current weak learner output \perp . Perhaps surprisingly, our boosting algorithm (Algorithm 1) neither increases nor decreases the probabilities of those symbols,⁶ and an appropriate choice of the multiplicative weight provides the desired strong learner. In fact, we will argue that our update rule is an approximate version of the rule such that the CS advantage of the weak learner of the previous round is zero for the updated distribution (Appendix C).

Future perspectives. It has been known that the hardcore lemma and its variants play a pivotal role in the construction of cryptographic primitives: for

⁶ Note that since there is a renormalization procedure, the updated probabilities of those symbols may be changed.

instance, the construction of strong PRGs from weak ones [30] and the constructions of PRGs from one-way functions [38, 16]. In order to discuss the bit security of such constructions, we need to take into account unbalanced adversaries, as discussed above. Our hardcore lemma for the CS advantage may have a potential use in that direction of research.

1.4 Related work

The study of hardness amplification has a long history, and there are several proofs of the XOR lemma; see [10] for a thorough review. As mentioned in Section 1.3, in this paper, we prove our XOR lemma along the lines of the proof by Impagliazzo using the hardcore lemma [23].

Another line of studies on hardness amplification is the direct-product constructions; it aims to construct strongly hard (search type) functions from weak ones [13, 11, 26]. See [26] and the literature therein for recent related work. In this work, we focus on amplifying the hardness of Boolean (decision type) functions.

In the original paper [23], Impagliazzo provided two proofs of the hardcore lemma, a constructive one and one based on the min-max theorem.⁷ Later, it was pointed out that the constructive proof can be interpreted as the boosting algorithm in learning theory [25]. Based on such identification, there have been several improvements and applications of the hardcore lemma [22, 24, 3, 30, 37].

In contrast to the standard hardcore lemma stated in terms of the probability of correct prediction, our hardcore lemma is stated in terms of the CS advantage. The main difficulty of handling the CS advantage is that it may not be affine with respect to either the input distribution or (stochastic) circuit. Thus, it is unclear if the min-max theorem is applicable to prove the hardcore lemma for the CS advantage. To overcome this difficulty, we devise a modified version of the boosting algorithm in [24, 3] by considering that the adversary (weak learner in the context of learning) may output \perp in addition to 0 and 1. In the context of learning theory, by considering the asymmetry of weak learners' confidence for each output, we can improve the standard AdaBoost, which is known as the confidence-rated AdaBoost or the infoBoosting [36, 1, 19, 20]. Our boosting algorithm is closely related in spirit to those algorithms in the sense that the symbol \perp signifies that weak learners' confidence is zero. Since the CS advantage is a criterion initiated in cryptography, we believe it is an interesting contribution to characterize the boosting algorithm with \perp in terms of the CS advantages of weak learners; perhaps, it may have certain applications in learning theory.

Exploring a single quantity for measuring the security of primitives dates back to [21, 29], where the time-success ratio was proposed. Micciancio and Walter [32] revisited the notion of bit security, especially for decision primitives, and justified the CS advantage. Watanabe and Yasunaga [39] introduced the Rényi advantage to characterize an operational definition of bit security. The two advantages turned out to be (essentially) equivalent when adversaries are allowed to output \perp [40]. Recently, Lee [27] introduced another notion of bit security that captures search and decision games in a single framework.

⁷ The latter was attributed to Nisan.

1.5 Paper organization

We present the formulation of the hardness amplification and the XOR lemma for the Rényi advantage in Section 2. In Section 3, we present the hardcore lemma for the CS advantage and its proof using the boosting algorithm with \perp . Section 4 presents the proof of the XOR lemma by using the hardcore lemma. Other than the fact of approximation as in (1), we do not use the knowledge of bit security frameworks [32, 39]. For readers' convenience, we review the bit-security frameworks in Appendix A.

2 Hardness amplification for Rényi advantage

For $0 \leq \rho \leq 1$ and a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, ρ -hardness $H_{\text{avg}}^\rho(f)$ of function f is the largest integer s such that any circuit $C : \{0, 1\}^n \rightarrow \{0, 1, \perp\}$ of size at most s satisfies

$$\Pr_{x \sim U_n} (C(x) = f(x)) \leq \rho,$$

where U_n is the uniform distribution on $\{0, 1\}^n$. For a prescribed (typically small) margin $\delta > 0$, a function f is regarded as *mildly hard* if the value of $H_{\text{avg}}^{1-\delta}(f)$ is sufficiently large. By using the function f as a building block, we are interested in constructing another function that is much harder than f itself. A typically used construction is the so-called XOR construction: for a given integer $k \geq 2$, let $f^{\oplus k} : \{0, 1\}^{nk} \rightarrow \{0, 1\}$ be the function defined by

$$f^{\oplus k}(x_1, \dots, x_k) := f(x_1) \oplus \dots \oplus f(x_k),$$

where $x_1, \dots, x_k \in \{0, 1\}^n$. The standard XOR lemma of Yao claims that $f^{\oplus k}$ is hard in the sense that $H_{\text{avg}}^{1/2+\varepsilon}(f^{\oplus k})$ is as large as $\frac{\varepsilon^2}{\ln(1/\delta)} H_{\text{avg}}^{1-\delta}(f)$ for $\varepsilon \geq 2(1-\delta)^k$; this means that even though the circuit size is decreased by the factor of $\frac{\varepsilon^2}{\ln(1/\delta)}$, we can guarantee that the adversary's success probability of predicting the value of function $f^{\oplus k}$ is at most $\frac{1}{2} + \varepsilon$. More precisely, the following holds:

Proposition 1 (XOR lemma). *For $\varepsilon \geq 2(1-\delta)^k$, it holds that*

$$\Pr_{x_1, \dots, x_k \sim U_n} (C(x_1, \dots, x_k) = f^{\oplus k}(x_1, \dots, x_k)) \leq \frac{1}{2} + \varepsilon$$

for every circuit C of size at most s for $s = \Omega\left(\frac{\varepsilon^2}{\ln(1/\delta)}\right) \cdot H_{\text{avg}}^{1-\delta}(f)$.

In order to discuss the bit security of the XOR function, let us consider the distinguishing game between $u = 0$ instance $(x_1, \dots, x_k, f^{\oplus k}(x_1, \dots, x_k))$ and $u = 1$ instance $(x_1, \dots, x_k, \sigma)$, where σ is a random bit that is independent of (x_1, \dots, x_k) . Proposition 1 implies that (by the standard argument of converting a distinguisher to a predictor), for every circuit of size at most s , the standard distinguishing advantage (in terms of the total variation distance) is less than

ε . However, as discussed in [32, 39] (see also [40] for more detail), the standard advantage is not suitable for evaluating bit security. Thus, the above-mentioned XOR lemma does not guarantee that the bit security is preserved during the process of constructing $f^{\oplus k}$ from f . To resolve this issue, we derive an alternative version of the XOR lemma in terms of Rényi advantage

$$\text{Adv}_{A, f^{\oplus k}}^{\text{Rényi}} = D_{1/2}(A_0 \| A_1) = -2 \ln \sum_a \sqrt{A_0(a)A_1(a)},$$

where A_u is the distribution of the output by adversary when the instance is u . As we mentioned in Section 1 (see also Appendix A), the bit security can be approximated by the Rényi advantage up to a constant. To that end, it is desirable to derive a trade-off between the adversary's Rényi advantage and the circuit size. We use the weighted majority gate once in the reduction proof of the following theorem. To avoid the effect of how the weighted majority gate is implemented, we first assume that the weighted majority gate is available for free in the evaluation of the initial hardness $H_{\text{avg}}^{1-\delta}(f)$; in Remark 1, we will provide an estimate for the effect of implementing the weighted majority gate.

Theorem 1 (XOR lemma for Rényi advantage). *For $\varepsilon \geq 2(1 - \delta)^k$, it holds that*

$$\text{Adv}_{A, f^{\oplus k}}^{\text{Rényi}} \leq \varepsilon$$

for every circuit A of size $s' \leq \frac{\varepsilon}{48 \ln(1/\delta)} H_{\text{avg}}^{1-\delta}(f)$, where the initial hardness $H_{\text{avg}}^{1-\delta}(f)$ is evaluated under the assumption that the weighted majority gate is available for free.

Remark 1. The assumption of the availability of the weighted majority gate comes from the fact that it is used in the proof of the hardcore lemma of Lemma 1. As we discuss in Remark 2, the effect of implementing the weighted majority gate can be estimated. More specifically, the statement of Theorem 1 holds for circuit size $s' \leq \frac{\varepsilon}{64 \ln(1/\delta)} H_{\text{avg}}^{1-\delta}(f) - \frac{c}{\varepsilon}$ for some constant c . Thus, when the initial hardness is $H_{\text{avg}}^{1-\delta}(f) = \omega(\log(1/\delta)/\varepsilon^2)$, then the effect of the weighted majority is negligible.

We shall discuss an implication of Theorem 1. For a given integer s , the bit security against adversaries with cost s is evaluated as (1). In fact, in the setting of this section, the initial hardness $s = H_{\text{avg}}^{1-\delta}(f)$ means $\text{BS}_s(G_f) = \log s + \mathcal{O}(1)$. For the function f itself, since circuits of much smaller size s' may have the same success probability $1 - \delta$, we cannot guarantee $\text{BS}_{s'}(G_f) \geq \text{BS}_s(G_f)$. However, Theorem 1 implies that the XOR function $f^{\oplus k}$ satisfies

$$\text{BS}_{s'}(G_{f^{\oplus k}}) \geq \text{BS}_s(G_f) - \mathcal{O}(\log \ln(1/\delta))$$

for $s' = \frac{\varepsilon}{48 \ln(1/\delta)} H_{\text{avg}}^{1-\delta}(f)$. In this sense, the bit security is preserved in the hardness amplification.

3 Hardcore lemma for CS advantage

We shall prove Theorem 1 along the lines of the proof by Impagliazzo using the hardcore lemma [23]. To that end, we develop a new variant of the hardcore lemma in this section.

By the definition of hardness, any circuit C of size $s \leq H_{\text{avg}}^{1-\delta}(f)$ must satisfy

$$\Pr_{x \sim U_n} (C(x) = f(x)) \leq 1 - \delta. \quad (3)$$

This means that there exists a set $\mathcal{H}_C \subset \{0, 1\}^n$ of hard inputs such that $|\mathcal{H}_C| \geq \delta 2^n$ and the circuit C fails to compute $f(x)$ for every $x \in \mathcal{H}_C$; however, the hard sets may differ for different circuits. Impagliazzo's hardcore lemma claims that there exists a set of inputs that are universally hard for every circuit with a smaller size. It is more convenient to consider probability distributions, rather than sets, having density δ ; a distribution P on $\{0, 1\}^n$ is said to have density δ if $P(x) \leq \frac{1}{\delta 2^n}$ for every $x \in \{0, 1\}^n$, or equivalently, the min-entropy satisfies $H_{\min}(P) \geq n - \log(1/\delta)$. The standard hardcore lemma is a statement as follows:

Proposition 2 (Hardcore lemma). *There exists a hardcore distribution H having density δ such that*

$$\Pr_{x \sim H} (C(x) = f(x)) \leq \frac{1}{2} + \varepsilon \quad (4)$$

for every circuit C of size at most s for $s = \Omega\left(\frac{\varepsilon^2}{\ln(1/\delta)}\right) \cdot H_{\text{avg}}^{1-\delta}(f)$.

Since the standard hardcore lemma, Lemma 2, is insufficient to prove Theorem 1, we derive the following variant of the hardcore lemma in terms of the conditional squared (CS) advantage. For a given distribution P on $\{0, 1\}^n$ and a circuit $C : \{0, 1\}^n \rightarrow \{0, 1, \perp\}$, the CS advantage of predicting f is defined as⁸

$$\text{Adv}_{C,f|P}^{\text{CS}} := \frac{4 \left(\Pr(C(x) = f(x)) - \frac{1}{2} \Pr(C(x) \neq \perp) \right)^2}{\Pr(C(x) \neq \perp)} \quad (5)$$

$$= \frac{\left(\Pr(C(x) = f(x)) - \Pr(C(x) = \overline{f(x)}) \right)^2}{\Pr(C(x) \neq \perp)} \quad (6)$$

where the probability is with respect to $x \sim P$ and $\overline{f(x)} = f(x) \oplus 1$.

Lemma 1 (Hardcore lemma for CS advantage). *There exists a hardcore distribution H having density δ such that*

$$\text{Adv}_{C,f|H}^{\text{CS}} \leq \varepsilon \quad (7)$$

for every circuit C of size at most $s' := \frac{\varepsilon}{8 \ln(1/\delta)} H_{\text{avg}}^{1-\delta}(f)$.

⁸ More precisely, the CS advantage in (5) is for a predictor; on the other hand, when we define the bit security, we consider the CS advantage for a distinguisher (cf. (27)). The CS advantage for a predictor was first introduced in the context of the Goldreich-Levin algorithm [17].

Algorithm 1: Boosting

Input: The number $T \in \mathbb{N}$ of iteration and a circuit C_P satisfying (8) for each P with density δ

Output: A sequence of circuits $C_{P^{(1)}}, \dots, C_{P^{(T)}}$

- 1: Initialize $P^{(1)}$ as the uniform distribution on $\{0, 1\}^n$;
- 2: Repeat Step 3 and Step 4 for $1 \leq t \leq T$;
- 3: For a circuit $C_{P^{(t)}}$ that satisfies (8) for $P^{(t)}$, set $\gamma_t = \frac{\Delta_t}{4\alpha_t}$ for

$$\alpha_t := \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp),$$

$$\Delta_t := \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}),$$

and set

$$\hat{P}^{(t+1)}(x) = \frac{P^{(t)}(x) \exp(-\gamma_t \{\mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}]\})}{Z_{P^{(t)}}},$$

where $\overline{f(x)} = f(x) \oplus 1$ and

$$Z_{P^{(t)}} = \sum_x P^{(t)}(x) \exp(-\gamma_t \{\mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}]\})$$

is the normalizer;

- 4: For the set \mathcal{P}_δ of all distributions with density δ , set

$$P^{(t+1)} = \operatorname{argmin}_{P \in \mathcal{P}_\delta} D(P \parallel \hat{P}^{(t+1)}),$$

where D is the KL-divergence.

For a circuit that does not output \perp , i.e., $\Pr(C(x) \neq \perp) = 1$, we can rewrite (5) as

$$\Pr(C(x) = f(x)) = \frac{1}{2} + \frac{\sqrt{\operatorname{Adv}_{C,f|P}^{\text{CS}}}}{2}.$$

For such a circuit, the bounds (4) and (7) are the same up to a constant (ε^2 in Proposition 2 corresponds to ε in Lemma 1). A main new feature of Lemma 1 is that it can be applied to circuits that may output \perp with significant probability.

In contrast to the standard correct probability (the left-hand side of (4)), the CS advantage is not affine with respect to either the input distribution or (stochastic) circuit. Thus, it is unclear if the min-max theorem is applicable to prove Lemma 1. Instead, we consider a modified version of the boosting algorithm by taking into account the fact that the adversary (weak learner in the context of learning) may output \perp in addition to 0 and 1.

To prove Lemma 1 via a contradiction, suppose that for each distribution P having density δ , there exists a circuit C_P of size at most s' such that

$$\operatorname{Adv}_{C_P, f|P}^{\text{CS}} > \varepsilon. \tag{8}$$

Starting from the uniform distribution $P^{(1)}$ and a circuit that satisfies (8) for $P^{(1)}$, we are going to sequentially update distributions and corresponding circuits that satisfy (8); then, by combining those circuits, we eventually construct a circuit that violates the assumption (3) on the hardness of f . As we mentioned above, this procedure is essentially the same as the boosting algorithm in learning theory, in which we construct a strong learner from weak learners. Our algorithm for boosting is described in Algorithm 1. Perhaps surprisingly, the update rule of Algorithm 1 does not alter the distribution when the output of a circuit is \perp ; a rationale for the algorithm will be discussed in Appendix C.

The sequence of distributions generated by the algorithm satisfies the following.

Lemma 2. *The distributions $P^{(1)}, \dots, P^{(T)}$ generated by Algorithm 1 satisfy*

$$\begin{aligned} & \sum_{t=1}^T \frac{1}{T} \frac{\text{Adv}_{C_P^{(t)}, f|P^{(t)}}^{\text{CS}}}{8} \\ & \leq \mathbb{E}_{x \sim P} \left[\sum_{t=1}^T \frac{1}{T} \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \right] + \frac{D(P \| P^{(1)})}{T} \end{aligned} \quad (9)$$

for every $P \in \mathcal{P}_\delta$, where \mathcal{P}_δ is the set of all distributions with density δ .

Lemma 2 is a counterpart of technical results that appeared in the literature for the boosting without \perp (e.g. [24, Theorem 1] or [3, Lemma 4.1]). Since our boosting algorithm invokes circuits that may output \perp , the performance is evaluated in terms of the CS advantage instead of the standard advantage. A main technical difficulty of proving Lemma 2 is connecting Algorithm 1 to the CS advantage since the output \perp does not appear explicitly in the algorithm. The following lemma is a crucial step in proving Lemma 2, which enables us to evaluate the performance of Algorithm 1 in terms of the CS advantage. The lemma will be proved in Section 3.1.

Lemma 3. *For every $1 \leq t \leq T$, the distributions $P^{(t)}$ and $\hat{P}^{(t+1)}$ in Algorithm 1 satisfy*

$$D(P^{(t)} \| \hat{P}^{(t+1)}) \leq 2\gamma_t^2 \Pr_{x \sim P^{(t)}}(C_{P^{(t)}}(x) \neq \perp). \quad (10)$$

Since we set $\gamma_t = \frac{\Delta_t}{4\alpha_t}$ in Algorithm 1, note that the right-hand side of (10) is $\frac{\Delta_t^2}{8\alpha_t} = \frac{1}{8} \text{Adv}_{C_P^{(t)}, f|P^{(t)}}^{\text{CS}}$. In this manner, the CS advantage shows up in the analysis even though the output \perp does not appear in the algorithm explicitly. The details of the proof of Lemma 2 are presented in Section 3.2.

Once we prove Lemma 2, we can prove the hardcore lemma for the CS advantage (Lemma 1) via a contrapositive argument: we assume that for each distribution P having density δ , there exists a circuit C_P of size at most s' satisfying (8). This assumption enables us to find a weak learner in each iteration

of Algorithm 1, and we can construct a strong learner from a sequence of circuits found by the algorithm. One key difference from the proof of the standard hardcore lemma is that, instead of the standard majority, we take the weighted majority to construct the strong learner from weak learners since reliabilities of weak learners vary depending on the probability of output \perp . The detailed proof of Lemma 1 is provided in Section 3.3.

3.1 Proof of Lemma 3

From the update rule of Algorithm 1, we have

$$\begin{aligned}
& D(P^{(t)} \parallel \hat{P}^{(t+1)}) \\
&= \sum_x P^{(t)}(x) \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} + \ln Z_{P^{(t)}} \\
&= \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} + \ln Z_{P^{(t)}}. \quad (11)
\end{aligned}$$

Now, we evaluate $\ln Z_{P^{(t)}}$ as

$$\begin{aligned}
\ln Z_{P^{(t)}} &= \ln \sum_x P^{(t)}(x) \exp \left(-\gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \right) \\
&\leq \ln \sum_x P^{(t)}(x) (1 - \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \\
&\quad + 2\gamma_t^2 \mathbf{1}[C_{P^{(t)}}(x) \neq \perp]) \\
&= \ln \left(1 - \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \right. \\
&\quad \left. + 2\gamma_t^2 \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp) \right) \\
&\leq -\gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \\
&\quad + 2\gamma_t^2 \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp), \quad (12)
\end{aligned}$$

where the first inequality follows from⁹ $e^{-\theta} \leq 1 - \theta + 2\theta^2$ for $\theta \in [-1, 1]$ and that

$$\exp \left(-\gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \right) = 1$$

when $C_{P^{(t)}}(x) = \perp$; and the second inequality follows from $\ln(1 - \theta) \leq -\theta$ for $\theta < 1$. By combining (11) and (12), we have the desired bound. \square

⁹ By the Taylor approximation, we have

$$e^{-\theta} \leq 1 - \theta + \sup_{-1 \leq \tau \leq 1} \frac{e^{-\tau}}{2} \theta^2 \leq 1 - \theta + \frac{e}{2} \theta^2 \leq 1 - \theta + 2\theta^2.$$

3.2 Proof of Lemma 2

By the cosine law of the KL-divergence, we have

$$\begin{aligned}
& D(P\|P^{(t)}) - D(P\|\hat{P}^{(t+1)}) \\
&= \sum_x (P^{(t)}(x) - P(x)) (\ln P^{(t)}(x) - \ln \hat{P}^{(t+1)}(x)) - D(P^{(t)}\|\hat{P}^{(t+1)}) \\
&= \sum_x (P^{(t)}(x) - P(x)) \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \\
&\quad - D(P^{(t)}\|\hat{P}^{(t+1)}) \\
&= \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \\
&\quad - \gamma_t \mathbb{E}_{x \sim P} \left[\left\{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \right\} \right] - D(P^{(t)}\|\hat{P}^{(t+1)}) \\
&\geq \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \\
&\quad - \gamma_t \mathbb{E}_{x \sim P} \left[\left\{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \right\} \right] \\
&\quad - 2\gamma_t^2 \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp), \tag{13}
\end{aligned}$$

where, in the second equality, we used the fact that $\ln Z_{P^{(t)}}$ does not depend on x and $\sum_x (P^{(t)}(x) - P(x)) = 0$ to eliminate $\ln Z_{P^{(t)}}$; and we used Lemma 3 in the final inequality. Here, we apply the Pythagorean inequality by noting that \mathcal{P}_δ is a closed convex set: it holds that (e.g., see [5, Theorem 3.1])

$$D(P\|P^{(t+1)}) + D(P^{(t+1)}\|\hat{P}^{(t+1)}) \leq D(P\|\hat{P}^{(t+1)})$$

for any $P \in \mathcal{P}_\delta$, which implies

$$D(P\|P^{(t+1)}) \leq D(P\|\hat{P}^{(t+1)}). \tag{14}$$

Thus, (13) and (14) imply

$$\begin{aligned}
& D(P\|P^{(t)}) - D(P\|P^{(t+1)}) \\
&\geq \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \\
&\quad - \gamma_t \mathbb{E}_{x \sim P} \left[\left\{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \right\} \right] \\
&\quad - 2\gamma_t^2 \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp).
\end{aligned}$$

By taking the summation of both sides for $t = 1$ through T , we have

$$\begin{aligned}
& D(P\|P^{(1)}) - D(P\|P^{(T+1)}) \\
& \geq \sum_{t=1}^T \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \\
& \quad - \mathbb{E}_{x \sim P} \left[\sum_{t=1}^T \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \right] \\
& \quad - \sum_{t=1}^T 2\gamma_t^2 \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp). \tag{15}
\end{aligned}$$

Since $D(P\|P^{(T+1)}) \geq 0$, by substituting α_t, Δ_t and $\gamma_t = \frac{\Delta_t}{4\alpha_t}$ defined in Algorithm 1, and by rearranging terms, we have

$$\begin{aligned}
& \sum_{t=1}^T \frac{\Delta_t^2}{8\alpha_t} \\
& \leq \mathbb{E}_{x \sim P} \left[\sum_{t=1}^T \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \right] + D(P\|P^{(1)}).
\end{aligned}$$

Finally, by noting that $\text{Adv}_{C_P^{(t)}, f|P^{(t)}}^{\text{CS}} = \frac{\Delta_t^2}{\alpha_t}$ and by dividing by T , we have (9). \square

3.3 Proof of Lemma 1

To prove via a contradiction, suppose that for each distribution P having density δ , there exists a circuit C_P of size at most s' satisfying (8). We shall prove that there exists a circuit C^* of size at most $s := H_{\text{avg}}^{1-\delta}(f)$ such that

$$\Pr_{x \sim U_n} (C^*(x) = f(x)) > 1 - \delta. \tag{16}$$

We construct C^* as follows. For $T = \left\lceil \frac{8 \ln(1/\delta)}{\epsilon} \right\rceil$, let $C_{P^{(1)}}, \dots, C_{P^{(T)}}$ be the circuits obtained by Algorithm 1. For a given input $x \in \{0, 1\}^n$, by invoking the weighted majority oracle, C^* outputs $a \in \{0, 1\}$ if (the tie can be decided arbitrarily)

$$\sum_{t=1}^T \frac{1}{T} \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = a] - \mathbf{1}[C_{P^{(t)}}(x) = \bar{a}] \} > 0, \tag{17}$$

where $\bar{a} = a \oplus 1$. Note that the size of C^* is Ts' . Note also that C^* makes an error for input x , i.e., $C^*(x) \neq f(x)$ only if

$$\sum_{t=1}^T \frac{1}{T} \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \leq 0. \tag{18}$$

Let

$$\mathcal{E} = \left\{ x \in \{0, 1\}^n : \sum_{t=1}^T \frac{1}{T} \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \leq 0 \right\}.$$

be the set of all inputs such that C^* may make an error. If we prove $\frac{|\mathcal{E}|}{2^n} < \delta$, we are done, i.e., (16) holds. To prove via a contradiction, assume that $\frac{|\mathcal{E}|}{2^n} \geq \delta$, which implies that the uniform distribution $P_{\mathcal{E}}$ on \mathcal{E} has density δ . Since $P^{(1)}$ is the uniform distribution, we have

$$D(P_{\mathcal{E}} \| P^{(1)}) = \sum_x P_{\mathcal{E}}(x) \ln 2^n P_{\mathcal{E}}(x) \leq \sum_x P_{\mathcal{E}}(x) \ln(1/\delta) = \ln(1/\delta).$$

By applying Lemma 2 for $P_{\mathcal{E}}$, by noting (8) for each $C_{P^{(t)}}$, and by noting that (18) with probability 1 for $x \sim P_{\mathcal{E}}$, we have

$$\begin{aligned} \frac{\varepsilon}{8} &< \mathbb{E}_{x \sim P_{\mathcal{E}}} \left[\sum_{t=1}^T \frac{1}{T} \gamma_t \{ \mathbf{1}[C_{P^{(t)}}(x) = f(x)] - \mathbf{1}[C_{P^{(t)}}(x) = \overline{f(x)}] \} \right] + \frac{D(P_{\mathcal{E}} \| P^{(1)})}{T} \\ &\leq \frac{D(P_{\mathcal{E}} \| P^{(1)})}{T} \\ &\leq \frac{\varepsilon}{8}, \end{aligned}$$

which is a contradiction. \square

Remark 2. Even though we did not take into account the precision of the weight γ_t in the proof of Lemma 1, it can be evaluated as follows. Note that α_t and Δ_t in Algorithm 1 satisfy

$$\varepsilon < \frac{\Delta_t^2}{\alpha_t} \leq \frac{\Delta_t}{\alpha_t} \leq 1.$$

By setting $\tau = \lceil \log(1/\varepsilon) \rceil$ so that $\frac{1}{2^\tau} \leq \varepsilon$, we divide the interval $(1/2^\tau, 1]$ into τ parts

$$\left(\frac{1}{2^\tau}, \frac{1}{2^{\tau-1}} \right], \left(\frac{1}{2^{\tau-1}}, \frac{1}{2^{\tau-2}} \right], \dots, \left(\frac{1}{2}, 1 \right].$$

Then, each $\frac{\Delta_t}{\alpha_t}$ satisfies

$$\frac{\Delta_t}{\alpha_t} \in \left(\frac{1}{2^{\ell_t}}, \frac{1}{2^{\ell_t-1}} \right]$$

for some ℓ_t ; if we set $\gamma_t = \frac{1}{2^{\ell_t+2}}$, then we have

$$\frac{\Delta_t}{8\alpha_t} \leq \gamma_t < \frac{\Delta_t}{4\alpha_t}. \quad (19)$$

If we use $\gamma_t = \frac{1}{2^{t+2}}$ instead of $\gamma_t = \frac{\Delta_t}{4\alpha_t}$ in Algorithm 1, the last two terms of (15) is lower bounded as

$$\begin{aligned}
& \sum_{t=1}^T \gamma_t \left\{ \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = f(x)) - \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) = \overline{f(x)}) \right\} \\
& \quad - \sum_{t=1}^T 2\gamma_t^2 \Pr_{x \sim P^{(t)}} (C_{P^{(t)}}(x) \neq \perp) \\
& = \sum_{t=1}^T \frac{\Delta_t^2}{\alpha_t} \gamma_t \frac{\alpha_t}{\Delta_t} \left(1 - 2\gamma_t \frac{\alpha_t}{\Delta_t} \right) \\
& \geq \sum_{t=1}^T \frac{\Delta_t^2}{\alpha_t} \frac{1}{8} \left(1 - \frac{2}{8} \right) = \sum_{t=1}^T \frac{3\Delta_t^2}{32\alpha_t},
\end{aligned}$$

where the inequality follows from that the function $g(\theta) = \theta(1 - 2\theta)$ is lower bounded by $g(1/8)$ for $1/8 \leq \theta \leq 1/4$ and (19). Thus, even if we take into account the precision of the weight γ_t , we have the same claim as Lemma 1 except that the factor $\frac{1}{8}$ in s' is replaced by $\frac{3}{32}$.

Furthermore, since each weight γ_t takes a value between $1/2^{\tau+2}$ and $1/8$, we can implement the weighted majority by creating $2^{\tau+2}\gamma_t \leq 1/\varepsilon$ copies of each input and by using the majority.¹⁰ If we take into account the cost of the weighted majority, the proof goes through as long as

$$Ts' + \frac{cT}{\varepsilon} = T \left(s' + \frac{c}{\varepsilon} \right) \leq s$$

for some constant c .¹¹ Thus, the claim of Lemma 1 still holds for every circuit of size at most $s' = \frac{3\varepsilon}{32 \ln(1/\delta)} \mathbf{H}_{\text{avg}}^{1-\delta}(f) - \frac{c}{\varepsilon}$; when the initial hardness is $\mathbf{H}_{\text{avg}}^{1-\delta}(f) = \omega(\log(1/\delta)/\varepsilon^2)$, then the cost of the weighted majority is negligible.

4 Proof of Theorem 1

For notational simplicity, we prove the case of $k = 2$; general $k \geq 2$ can be proved similarly. Toward deriving a contradiction, suppose that there exists A with size $s' \leq \frac{\varepsilon}{48 \ln(1/\delta)} \mathbf{H}_{\text{avg}}^{1-\delta}(f)$ such that $\text{Adv}_{A, f^{\oplus k}}^{\text{Renyi}} > \varepsilon$. By Lemma 1, there exists a hardcore distribution H with density δ such that

$$\text{Adv}_{C, f|H}^{\text{CS}} \leq \frac{\varepsilon}{6} \tag{20}$$

¹⁰ Such a naive implementation of the weighted majority has been studied in the circuit complexity [9].

¹¹ It comes from the fact that the majority can be realized by a circuit of linear size of inputs [41].

for every circuit C of size $s' \leq \frac{\varepsilon}{48 \ln(1/\delta)} H_{\text{avg}}^{1-\delta}(f)$. Let G be the distribution on $\{0, 1\}^n$ given by $G(x) = \frac{1/2^n - \delta H(x)}{1-\delta}$. Then, the uniform distribution U_n on $\{0, 1\}^n$ can be decomposed as

$$U_n(x) = (1 - \delta)G(x) + \delta H(x). \quad (21)$$

When (x_1, x_2) are distributed according to distribution Q , we denote the output distribution of adversary A by A_u^Q for $u = 0, 1$ (note that, when $u = 1$, σ is generated independently of $(x_1, x_2) \sim Q$). Then, the output distribution A_u of adversary A under the uniform distribution can be decomposed as

$$A_u = (1 - \delta)^2 A_u^{GG} + (1 - \delta)\delta A_u^{GH} + \delta(1 - \delta)A_u^{HG} + \delta^2 A_u^{HH}. \quad (22)$$

By the joint convexity of the Rényi divergence of order $1/2$ (e.g., see [8, Theorem 11]), we have

$$\begin{aligned} \varepsilon &< \text{Adv}_{A, f \oplus k}^{\text{Rényi}} \\ &= D_{1/2}(A_0 \| A_1) \\ &\leq (1 - \delta)^2 D_{1/2}(A_0^{GG} \| A_1^{GG}) + (1 - \delta)\delta D_{1/2}(A_0^{GH} \| A_1^{GH}) \\ &\quad + \delta(1 - \delta)D_{1/2}(A_0^{HG} \| A_1^{HG}) + \delta^2 D_{1/2}(A_0^{HH} \| A_1^{HH}). \end{aligned}$$

Since $(1 - \delta)^2 < \frac{\varepsilon}{2}$ by the assumption and $D_{1/2}(A_0^{GG} \| A_1^{GG}) \leq 1$ (cf. [40, Proposition 1]), we have

$$\begin{aligned} \frac{\varepsilon}{2} &< (1 - \delta)\delta D_{1/2}(A_0^{GH} \| A_1^{GH}) \\ &\quad + \delta(1 - \delta)D_{1/2}(A_0^{HG} \| A_1^{HG}) + \delta^2 D_{1/2}(A_0^{HH} \| A_1^{HH}). \end{aligned}$$

Since the summation of $(1 - \delta)\delta$, $\delta(1 - \delta)$, and δ^2 is less than 1, by the averaging argument, at least one of $D_{1/2}(A_0^{GH} \| A_1^{GH})$, $D_{1/2}(A_0^{HG} \| A_1^{HG})$, or $D_{1/2}(A_0^{HH} \| A_1^{HH})$ is larger than $\frac{\varepsilon}{2}$. For instance, suppose that $D_{1/2}(A_0^{GH} \| A_1^{GH}) > \frac{\varepsilon}{2}$ (other cases are similar). We can write

$$\begin{aligned} A_0^{GH}(a) &= \Pr_{\substack{x \sim G \\ y \sim H}} (A(x, y, f(x) \oplus f(y)) = a) \\ &= \mathbb{E}_{x \sim G} \left[\Pr_{y \sim H} (A(x, y, f(x) \oplus f(y)) = a) \right] \end{aligned}$$

and

$$\begin{aligned} A_1^{GH}(a) &= \Pr_{\substack{x \sim G \\ y \sim H}} (A(x, y, \sigma) = a) \\ &= \Pr_{\substack{x \sim G \\ y \sim H}} (A(x, y, f(x) \oplus \sigma) = a) \\ &= \mathbb{E}_{x \sim G} \left[\Pr_{y \sim H} (A(x, y, f(x) \oplus \sigma) = a) \right], \end{aligned} \quad (23)$$

where the identity (23) holds since σ being a random bit independent of (x, y) implies that $f(x) \oplus \sigma$ is a random bit independent of (x, y) . For each $x \in \{0, 1\}^n$, let us consider an adversary A^x for distinguishing between $(y, f(y))$ and (y, σ) for $y \sim H$; given input (y, z) , A^x runs $A(x, y, f(x) \oplus z)$ (since we consider non-uniform complexity, $f(x)$ can be precomputed and provided to the circuit, and the size of A^x is s'). By applying the joint convexity of the Rényi divergence of order $1/2$ once more, we have

$$\frac{\varepsilon}{2} < D_{1/2}(A_0^{GH} \| A_1^{GH}) \quad (24)$$

$$\leq \mathbb{E}_{x \sim G} \left[D_{1/2}(A_0^x \| A_1^x) \right]. \quad (25)$$

Thus, there exists $x \in \{0, 1\}^n$ such that

$$\frac{\varepsilon}{2} < D_{1/2}(A_0^x \| A_1^x).$$

By the same argument as [40, Theorem 3] (for completeness, we provide a proof in Appendix B), there exists a predictor $C : \{0, 1\}^n \rightarrow \{0, 1, \perp\}$ that invokes A^x once and satisfies

$$\text{Adv}_{C, f|H}^{\text{CS}} \geq \frac{1}{3} D_{1/2}(A_0^x \| A_1^x) > \frac{\varepsilon}{6}, \quad (26)$$

which contradicts (20). \square

Acknowledgements This work was supported in part by JSPS KAKENHI Grant Numbers 23H00468, 23K17455, and 24H00071.

A Bit-security frameworks

As we mentioned in Section 1, the bit security for decision game G was first introduced in [32], and an operational framework was later introduced in [39]. For readers' convenience, in this appendix, we review the bit-security frameworks of [32] and [39] and discuss their equivalence shown in [40]. Since the main focus of this paper is decision games, we only present the formulation of the decision game in the following; see [32, 39] for the formulation of the search game.

Let U be the random variable describing the choice of a decision game; for instance, in the indistinguishability game of pseudorandom number generator (PRG) from the true random number generator (TRG), $U = 0$ corresponds to the game played with PRG and $U = 1$ corresponds to the game played with TRG. In the framework of [32], we consider an adversary A that outputs \perp in addition to 0 and 1; the symbol \perp signifies that the adversary has difficulty predicting the value of U and gives up the prediction. For the random variable Y describing the adversary's output, let

$$\begin{aligned} \alpha_A &:= \Pr(Y \neq \perp), \\ \beta_A &:= \Pr(Y = U | Y \neq \perp), \end{aligned}$$

and define the conditional squared (CS) advantage

$$\text{Adv}_{A,G}^{\text{CS}} := \alpha_A(2\beta_A - 1)^2. \quad (27)$$

Then, the bit security of [32] is defined as¹²

$$\min_A \left\{ \log_2 \left(\frac{s_A}{\text{Adv}_{A,G}^{\text{CS}}} \right) \right\}, \quad (28)$$

where s_A is the cost of the adversary.¹³

In the bit-security framework of [39], in order to define bit-security operationally, we consider an outer adversary B in addition to the inner adversary A that plays the given security game. In the framework, the outer adversary B seeks to increase the success probability of predicting U by invoking the inner adversary A for $N_{A,B}$ times. Then, by integrating the outputs from the invocations of the inner adversary, the outer adversary outputs the predicted value Z . Then, for a prescribed success probability $1 - \mu$ (say 0.99), the bit security of decision game G is defined as

$$\text{BS}_G^\mu := \min_{A,B} \left\{ \log_2(N_{A,B} \cdot s_A) : \Pr(Z = U) \geq 1 - \mu \right\}. \quad (29)$$

Furthermore, it was shown in [39] that the bit security is characterized by the Rényi advantage up to a constant, i.e.,

$$\text{BS}_G^\mu = \min_A \left\{ \log_2 s_A + \log_2 \left[\frac{1}{\text{Adv}_{A,G}^{\text{Rényi}}} \right] \right\} + \mathcal{O}(1), \quad (30)$$

where the Rényi advantage is given by the Rényi divergence

$$\text{Adv}_{A,G}^{\text{Rényi}} := D_{1/2}(A_0 \| A_1)$$

for the distributions A_0 and A_1 of adversary for $U = 0$ and $U = 1$, respectively.

At first glance, the bit security defined in (28) and that in (30) (defined via operational formula (29)) are different quantities. However, it was shown in [40] that the two notions of bit security are equivalent in the following sense.

Proposition 3. *For an arbitrary adversary A for a decision game G , it holds that $\text{Adv}_{A,G}^{\text{CS}} \leq 8\text{Adv}_{A,G}^{\text{Rényi}}$. On the other hand, for an arbitrary adversary A satisfying $\text{Adv}_{A,G}^{\text{Rényi}} \leq 1$, there exists an adversary \tilde{A} having the same cost as A such that $\text{Adv}_{A,G}^{\text{Rényi}} \leq 12\text{Adv}_{\tilde{A},G}^{\text{CS}}$.*

¹² In [32], the authors first introduced an advantage using the Shannon entropy and the mutual information; then, in order to justify the definition (28), they discussed that that advantage is approximated by the CS advantage.

¹³ In this paper, we focus on the circuit size.

By using the conversion of two advantages in Proposition 3, we can argue that the two notions of bit security coincide up to a constant; for more detail, see [40, Section 4]. Since the two notions are equivalent, in the main body of the paper, we focus on the bit security characterized by the Rényi advantage, (30). However, the CS advantage adapted for predictors also plays an important role when we prove the hardcore lemma in Section 3. From a technical perspective, it seems that the CS advantage is more suitable for analyzing the performance of algorithms; on the other hand, the Rényi advantage is more convenient for analysis in a certain situation since it satisfies (joint) convexity with respect to the distributions, which is used in the proof of Theorem 1.

Here, we should note that the standard advantage defined by the total variation distance between A_0 and A_1 is unsuitable for evaluating bit security. Notably, as was pointed out in [32] (see also [40, Section 1.3] further discussion), the standard advantage cannot resolve the paradoxical nature of the linear test for the PRG.

B Proof of (26)

In this appendix, we prove that there exists a predictor $C : \{0, 1\}^n \rightarrow \{0, 1, \perp\}$ that invokes A^x once and satisfies (26). We use the following technical lemma.

Lemma 4. *For given distributions P and Q with $P \ll Q$, we have*

$$D_{1/2}(P\|Q) \leq D(P\|Q) \leq \sum_{x \in \mathcal{X}^+} \frac{(P(x) - Q(x))^2}{Q(x)}.$$

where $\mathcal{X}^+ = \{x : Q(x) > 0\}$, and $D(P\|Q) = \sum_x P(x) \log(P(x)/Q(x))$ is the KL-divergence.

Proof. The former inequality follows from the fact that the Rényi divergence is monotonically non-decreasing with respect to α and $D(P\|Q) = \lim_{\alpha \rightarrow 1} D_\alpha(P\|Q)$. The latter inequality appears in the middle of the proof of [15, Lemma 4.1]. \square

Note that, for $y \sim H$, the distribution of the adversary A^x for $u = 0$ instance $(y, f(y))$ and $u = 1$ instance (y, σ) are given by

$$\begin{aligned} A_0^x(a) &= \Pr_{y \sim H} (A^x(y, f(y)) = a), \\ A_1^x(a) &= \Pr_{y \sim H} (A^x(y, \sigma) = a) \end{aligned}$$

for $a \in \{0, 1, \perp\}$. Note that the support of $(y, f(y))$ is included in the support of (y, σ) .¹⁴ Thus, if the adversary A^x outputs a symbol a with positive probability under $u = 0$, then A^x must output a with positive probability under $u = 1$ as well, i.e., $A_0^x \ll A_1^x$.

¹⁴ Here, the support is the set of realizations that occur with positive probability.

Let $a^* \in \{0, 1, \perp\}$ be such that $A_1^x(a^*) > 0$ and

$$\max_{\substack{a \in \{0, 1, \perp\}: \\ A_1^x(a) > 0}} \frac{(A_0^x(a) - A_1^x(a))^2}{A_1^x(a)} = \frac{(A_0^x(a^*) - A_1^x(a^*))^2}{A_1^x(a^*)}.$$

Then, by Lemma 4, we have

$$D_{1/2}(A_0^x \| A_1^x) \leq 3 \frac{(A_0^x(a^*) - A_1^x(a^*))^2}{A_1^x(a^*)}. \quad (31)$$

We consider two cases separately.

When $A_0^x(a^*) \geq A_1^x(a^*)$ In this case, we consider the following predictor C . First, we sample the uniform random bit σ . Second,

- If $A^x(y, \sigma) = a^*$, then C outputs σ ;
- If $A^x(y, \sigma) \neq a^*$, then C outputs \perp .

For this predictor, we have

$$\begin{aligned} \Pr_{y \sim H} (C(y) \neq \perp) &= \Pr_{y \sim H} (A^x(y, \sigma) = a^*) \\ &= A_1^x(a^*) \end{aligned}$$

and

$$\begin{aligned} \Pr_{y \sim H} (C(y) = f(y)) &= \frac{1}{2} \Pr_{y \sim H} (A^x(y, \sigma) = a^* | \sigma = f(y)) \\ &= \frac{1}{2} \Pr_{y \sim H} (A^x(y, f(y)) = a^*) \\ &= \frac{A_0^x(a^*)}{2}, \end{aligned}$$

which implies

$$\Pr_{y \sim H} (C(y) = f(y)) - \frac{1}{2} \Pr_{y \sim H} (C(y) \neq \perp) = \frac{A_0^x(a^*) - A_1^x(a^*)}{2}.$$

Thus, the CS advantage of C satisfies (cf. (5))

$$\begin{aligned} \text{Adv}_{C, f|H}^{\text{CS}} &= \frac{(A_0^x(a^*) - A_1^x(a^*))^2}{A_1^x(a^*)} \\ &\geq \frac{1}{3} D_{1/2}(A_0^x \| A_1^x), \end{aligned}$$

where the last inequality follows from (31).

When $A_0^x(a^*) < A_1^x(a^*)$ In this case, we consider the following predictor. First, we sample the uniform random bit σ . Second,

- If $A^x(y, \sigma) = a^*$, then C outputs $\sigma \oplus 1$;
- If $A^x(y, \sigma) \neq a^*$, then C outputs \perp .

For this predictor, we have

$$\begin{aligned} \Pr_{y \sim H} (C(y) \neq \perp) &= \Pr_{y \sim H} (A^x(y, \sigma) = a^*) \\ &= A_1^x(a^*) \end{aligned}$$

and

$$\begin{aligned} \Pr_{y \sim H} (C(y) = f(y)) &= \Pr_{y \sim H} (\sigma = f(y) \oplus 1, A^x(y, \sigma) = a^*) \\ &= \Pr_{y \sim H} (A^x(y, \sigma) = a^*) - \Pr_{y \sim H} (\sigma = f(y), A^x(y, \sigma) = a^*) \\ &= A_1^x(a^*) - \frac{1}{2} \Pr_{y \sim H} (A^x(y, \sigma) = a^* | \sigma = f(y)) \\ &= A_1^x(a^*) - \frac{1}{2} \Pr_{y \sim H} (A^x(y, f(y)) = a^*) \\ &= A_1^x(a^*) - \frac{A_0^x(a^*)}{2}, \end{aligned}$$

which implies

$$\Pr_{y \sim H} (C(y) = f(y)) - \frac{1}{2} \Pr_{y \sim H} (C(y) \neq \perp) = \frac{A_1^x(a^*) - A_0^x(a^*)}{2}.$$

Thus, the CS advantage of C again satisfies

$$\begin{aligned} \text{Adv}_{C, f|H}^{\text{CS}} &= \frac{(A_0^x(a^*) - A_1^x(a^*))^2}{A_1^x(a^*)} \\ &\geq \frac{1}{3} D_{1/2}(A_0^x \| A_1^x). \end{aligned}$$

□

C Justification of update rule of Algorithm 1

Since we only discuss one iteration of the update, we omit the round number t . For the convenience of notations, in this appendix, we assume that the outputs of function f and circuits are ± 1 or 0 with correspondence $0 \leftrightarrow +1$, $1 \leftrightarrow -1$, and $\perp \leftrightarrow 0$.

Let us start with a review of the update rule of the standard boosting in our terminology (e.g. see [35]). For given distribution P and circuit C_P , the standard boosting algorithm creates the updated distribution as

$$\hat{P}_\gamma(x) = \frac{\exp(-\gamma C_P(x) f(x))}{Z(\gamma)} \quad (32)$$

for some $\gamma \in \mathbb{R}$, where

$$Z(\gamma) = \sum_x P(x) \exp(-\gamma C_P(x) f(x))$$

is the normalizer. The advantage of C_P for $x \sim P$ can be written as

$$\Pr_{x \sim P}(C_P(x) = f(x)) - \Pr_{x \sim P}(C_P(x) = \overline{f(x)}) = \sum_x P(x) C_P(x) f(x),$$

where $\overline{f(x)} = -f(x)$. By denoting $q_{+1} = \Pr_{x \sim P}(C_P(x) = f(x))$ and $q_{-1} = \Pr_{x \sim P}(C_P(x) = \overline{f(x)})$, we can rewrite the normalizer $Z(\gamma)$ as

$$Z(\gamma) = q_{+1} e^{-\gamma} + q_{-1} e^{\gamma}. \quad (33)$$

Then, if we set $\gamma^* = \frac{1}{2} \ln \frac{q_{+1}}{q_{-1}}$, then we have

$$\left. \frac{dZ(\gamma)}{d\gamma} \right|_{\gamma=\gamma^*} = \left[-q_{+1} e^{-\gamma} + q_{-1} e^{\gamma} \right] \Big|_{\gamma=\gamma^*} = 0.$$

By noting these facts, we can compute the advantage of C_P for $x \sim \hat{P}_\gamma$ as

$$\begin{aligned} \sum_x \hat{P}_\gamma(x) C_P(x) f(x) &= \frac{1}{Z(\gamma)} \sum_x P(x) \exp(-\gamma C_P(x) f(x)) C_P(x) f(x) \\ &= -\frac{1}{Z(\gamma)} \frac{dZ(\gamma)}{d\gamma}. \end{aligned}$$

Thus, if we set $\gamma = \gamma^*$ in the update rule, the advantage of C_P for $x \sim \hat{P}_{\gamma^*}$ is 0. Let us denote $\Delta = q_{+1} - q_{-1}$. Then, we have

$$\gamma^* = \frac{1}{2} \ln(1 + \Delta) - \frac{1}{2} \ln(1 - \Delta) \simeq \Delta$$

for small Δ . In fact, the approximate value Δ has been used in the update of the standard boosting algorithm in the literature (e.g. [24, 3]).

Next, let us consider the boosting with \perp . We use the same notations as above, and the update rule itself is also given by (32). However, note that C_P may also output 0 (corresponding to \perp). Let $q_0 = \Pr_{x \sim P}(C_P(x) = 0)$. Note that the CS advantage of C_P for $x \sim P$ can be written as

$$\text{Adv}_{C_P, f|P}^{\text{CS}} = \frac{(q_{+1} - q_{-1})^2}{1 - q_0}.$$

The normalizer can be written as

$$\begin{aligned} Z(\gamma) &= \sum_x P(x) \exp(-\gamma C_P(x) f(x)) \\ &= q_{+1} e^{-\gamma} + q_{-1} e^{\gamma} + q_0. \end{aligned}$$

Thus, if we set $\gamma^* = \frac{1}{2} \ln \frac{q_{+1}}{q_{-1}}$, then we again have

$$\left. \frac{dZ(\gamma)}{d\gamma} \right|_{\gamma=\gamma^*} = [-q_{+1}e^{-\gamma} + q_{-1}e^{\gamma}] \Big|_{\gamma=\gamma^*} = 0.$$

By denoting

$$\begin{aligned} \hat{\alpha}_\gamma &= \Pr_{x \sim \hat{P}_\gamma} (C_P(x) \neq 0), \\ \hat{\Delta}_\gamma &= \Pr_{x \sim \hat{P}_\gamma} (C_P(x) = f(x)) - \Pr_{x \sim \hat{P}_\gamma} (C_P(x) = \overline{f(x)}), \end{aligned}$$

the CS advantage of C_P for $x \sim \hat{P}_\gamma$ can be written as $\text{Adv}_{C_P, f | \hat{P}_\gamma}^{\text{CS}} = \frac{\hat{\Delta}_\gamma}{\hat{\alpha}_\gamma}$. Note that

$$\begin{aligned} \hat{\Delta}_\gamma &= \sum_x \hat{P}_\gamma(x) C_P(x) f(x) \\ &= \frac{1}{Z(\gamma)} \sum_x P(x) \exp(-\gamma C_P(x) f(x)) C_P(x) f(x) \\ &= -\frac{1}{Z(\gamma)} \frac{dZ(\gamma)}{d\gamma}. \end{aligned}$$

Thus, if we set $\gamma = \gamma^*$ in the update rule, the CS advantage of C_P for $x \sim \hat{P}_{\gamma^*}$ satisfies $\text{Adv}_{C_P, f | \hat{P}_{\gamma^*}}^{\text{CS}} = 0$. By denoting $\alpha = 1 - q_0$ and $\Delta = q_{+1} - q_{-1}$, we have

$$\begin{aligned} \gamma^* &= \frac{1}{2} \ln \left(\frac{1}{2} (1 - q_0) + \frac{\Delta}{2} \right) - \frac{1}{2} \ln \left(\frac{1}{2} (1 - q_0) - \frac{\Delta}{2} \right) \\ &= \frac{1}{2} \ln \left(1 + \frac{\Delta}{\alpha} \right) - \frac{1}{2} \ln \left(1 - \frac{\Delta}{\alpha} \right) \\ &\simeq \frac{\Delta}{\alpha}. \end{aligned}$$

for small $\frac{\Delta}{\alpha}$. Note that $\frac{\Delta}{4\alpha}$ is the multiplicative weight used in Algorithm 1. Thus, the update rule in Algorithm 1 can be regarded as an approximate version of the rule such that the CS advantage of the current weak learner for the updated distribution is zero.

References

1. Aslam, J.A.: Improving algorithms for boosting. In: Proceedings of the 13th Annual Conference on Computational Learning Theory. pp. 200–207 (2000)
2. Bader, C., Jager, T., Li, Y., Schäge, S.: On the impossibility of tight cryptographic reductions. In: Fischlin, M., Coron, J. (eds.) Advances in Cryptology - EUROCRYPT 2016. Lecture Notes in Computer Science, vol. 9666, pp. 273–304. Springer (2016). https://doi.org/10.1007/978-3-662-49896-5_10, https://doi.org/10.1007/978-3-662-49896-5_10

3. Barak, B., Hardt, M., Kale, S.: The uniform hardcore lemma via approximate Bregman projections. In: Proceedings of the 2009 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 1193–1200 (2009)
4. Bellare, M., Ristenpart, T.: Simulation without the artificial abort: Simplified proof and improved concrete security for waters’ ibe scheme. In: Advances in Cryptology - EUROCRYPT 2009, 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques. Lecture Notes in Computer Science, vol. 5479, pp. 407–424. Springer (2009). https://doi.org/10.1007/978-3-642-01001-9_24
5. Csiszár, I., Shields, P.C.: Information theory and statistics: A tutorial. Found. Trends Commun. Inf. Theory **1**(4) (2004). <https://doi.org/10.1561/0100000004>, <https://doi.org/10.1561/0100000004>
6. Diemert, D., Jager, T.: On the tight security of TLS 1.3: Theoretically sound cryptographic parameters for real-world deployments. J. Cryptol. **34**(3), 30 (2021). <https://doi.org/10.1007/s00145-021-09388-x>, <https://doi.org/10.1007/s00145-021-09388-x>
7. Dodis, Y., Impagliazzo, R., Jaiswal, R., Kabanets, V.: Security amplification for interactivecryptographic primitives. In: Reingold, O. (ed.) Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5444, pp. 128–145. Springer (2009). https://doi.org/10.1007/978-3-642-00457-5_9, https://doi.org/10.1007/978-3-642-00457-5_9
8. van Erven, T., Harremoës, P.: Rényi divergence and Kullback-Leibler divergence. IEEE Trans. Inform. Theory **60**(7), 3797–3820 (July 2014)
9. Goldmann, M., Håstad, J., Razborov, A.: Majority gates vs. general weighted threshold gates. Computational Complexity **2**, 277–300 (1992)
10. Goldreich, O., Nisan, N., Wigderson, A.: On Yao’s XOR lemma. Electronic Colloquium on Computational Complexity (March 1995)
11. Goldreich, O.: The Foundations of Cryptography - Volume 1: Basic Techniques. Cambridge University Press (2001). <https://doi.org/10.1017/CBO9780511546891>, <http://www.wisdom.weizmann.ac.il/%7Eoded/foc-vol1.html>
12. Goldreich, O.: On security preserving reductions - revised terminology. In: Goldreich, O. (ed.) Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation, Lecture Notes in Computer Science, vol. 6650, pp. 540–546. Springer (2011)
13. Goldreich, O., Impagliazzo, R., Levin, L.A., Venkatesan, R., Zuckerman, D.: Security preserving amplification of hardness. In: 31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume I. pp. 318–326. IEEE Computer Society (1990). <https://doi.org/10.1109/FSCS.1990.89550>, <https://doi.org/10.1109/FSCS.1990.89550>
14. Goldreich, O., Nisan, N., Wigderson, A.: On yao’s xor-lemma. In: Goldreich, O. (ed.) Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation, Lecture Notes in Computer Science, vol. 6650, pp. 273–301. Springer (2011). https://doi.org/10.1007/978-3-642-22670-0_23, https://doi.org/10.1007/978-3-642-22670-0_23
15. Götze, F., Sambale, H., Sinulis, A.: Higher order concentration for functions of weakly dependent random variables. Electronic Journal of Probability **24**(85), 1–19 (2019)

16. Haitner, I., Reingold, O., Vadhan, S.P.: Efficiency improvements in constructing pseudorandom generators from one-way functions. *SIAM J. Comput.* **42**(3), 1405–1430 (2013). <https://doi.org/10.1137/100814421>, <https://doi.org/10.1137/100814421>
17. Hast, G.: Nearly one-sided tests and the Goldreich-Levin predicate. *Journal of Cryptology* **17**, 209–229 (2004)
18. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM J. Comput.* **28**(4), 1364–1396 (1999). <https://doi.org/10.1137/S0097539793244708>, <https://doi.org/10.1137/S0097539793244708>
19. Hatano, K., Warmuth, M.K.: Boosting versus covering. In: *Advances in Neural Information Processing Systems*. vol. 16 (2003)
20. Hatano, K., Watanabe, O.: Learning r -of- k functions by boosting. In: *Proceedings of the 15th International Conference on Algorithmic Learning Theory*. pp. 114–126 (2004)
21. Herzberg, A., Luby, M.: Pseudorandomness in cryptography. In: Brickell, E.F. (ed.) *Advances in Cryptology - CRYPTO '92*, 12th Annual International Cryptology Conference, Santa Barbara, California, USA, August 16–20, 1992, Proceedings. *Lecture Notes in Computer Science*, vol. 740, pp. 421–432. Springer (1992). https://doi.org/10.1007/3-540-48071-4_29, https://doi.org/10.1007/3-540-48071-4_29
22. Holenstein, T.: Key agreement from weak bit agreement. In: *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC '05)*. pp. 664–673. ACM Press (2005)
23. Impagliazzo, R.: Hard-core distribution for somewhat hard problems. In: *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science (FOCS '95)*. pp. 538–545 (1995)
24. Kale, S.: Boosting and hard-core set constructions: a simplified approach (2007), *Electronic Colloquium on Computational Complexity (ECCC)*, Report No. 131
25. Klivans, A.R., Servedio, R.A.: Boosting and hard-core set construction. *Machine Learning* **51**, 217–238 (2003)
26. Lanzenberger, D., Maurer, U.: Direct product hardness amplification. In: Nissim, K., Waters, B. (eds.) *Theory of Cryptography - 19th International Conference, TCC 2021*, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II. *Lecture Notes in Computer Science*, vol. 13043, pp. 605–625. Springer (2021)
27. Lee, K.: Bit security as cost to demonstrate advantage. *IACR Communications in Cryptology* **1**(1) (2024). <https://doi.org/10.62056/an5txol7>
28. Li, B., Micciancio, D., Schultz, M., Sorrell, J.: Securing approximate homomorphic encryption using differential privacy. In: Dodis, Y., Shrimpton, T. (eds.) *Advances in Cryptology - CRYPTO 2022*. *Lecture Notes in Computer Science*, vol. 13507, pp. 560–589. Springer (2022). https://doi.org/10.1007/978-3-031-15802-5_20, https://doi.org/10.1007/978-3-031-15802-5_20
29. Luby, M.: Pseudorandomness and cryptographic applications. *Princeton computer science notes*, Princeton University Press (1996)
30. Maurer, U., Tessaro, S.: A hardcore lemma for computational indistinguishability: Security amplification for arbitrary weak PRGs with optimal stretch. In: *TCC 2010*. *Lecture Notes in Computer Science*, vol. 5078, pp. 237–254. Springer (2010)
31. Maurer, U.M., Tessaro, S.: Computational indistinguishability amplification: Tight product theorems for system composition. In: Halevi, S. (ed.) *Advances in Cryptology - CRYPTO 2009*, 29th Annual International Cryptology Conference, Santa

- Barbara, CA, USA, August 16-20, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5677, pp. 355–373. Springer (2009). https://doi.org/10.1007/978-3-642-03356-8_21, https://doi.org/10.1007/978-3-642-03356-8_21
32. Micciancio, D., Walter, M.: On the bit security of cryptographic primitives. In: Nielsen, J.B., Rijmen, V. (eds.) *Advances in Cryptology - EUROCRYPT 2018*. Lecture Notes in Computer Science, vol. 10820, pp. 3–28. Springer (2018). https://doi.org/10.1007/978-3-319-78381-9_1
 33. Morgan, A., Pass, R.: On the security loss of unique signatures. In: *Theory of Cryptography*. *Theory of Cryptography*, vol. 11239, pp. 507–536. Springer (2018). https://doi.org/10.1007/978-3-030-03807-6_19
 34. Myers, S.A.: Efficient amplification of the security of weak pseudo-random function generators. *J. Cryptol.* **16**(1), 1–24 (2003). <https://doi.org/10.1007/s00145-002-0007-1>, <https://doi.org/10.1007/s00145-002-0007-1>
 35. Schapire, R.E., Freund, Y.: *Boosting: Foundations and Algorithms*. MIT Press (2012)
 36. Schapire, R.E., Singer, Y.: Improved boosting algorithm using confidence-rated predictions. *Machine Learning* **37**(3), 297–336 (1999)
 37. Vadhan, S., Zheng, C.J.: A uniform min-max theorem with applications in cryptography. In: *Advances in Cryptology–CRYPTO ’13*. Lecture Notes in Computer Science, vol. 8042, pp. 93–110. Springer (2013)
 38. Vadhan, S., Zheng, C.J.: Characterizing pseudentropy and simplifying pseudorandom generator constructions. *STOC ’12*, Association for Computing Machinery, New York, NY, USA (2012). <https://doi.org/10.1145/2213977.2214051>, <https://doi.org/10.1145/2213977.2214051>
 39. Watanabe, S., Yasunaga, K.: Bit security as computational cost for winning games with high probability. In: Tibouchi, M., Wang, H. (eds.) *Advances in Cryptology - ASIACRYPT 2021*. Lecture Notes in Computer Science, vol. 13092, pp. 161–188. Springer (2021)
 40. Watanabe, S., Yasunaga, K.: Unified view for notions of bit security. In: Guo, J., Steinfeld, R. (eds.) *Advances in Cryptology - ASIACRYPT 2023*. Lecture Notes in Computer Science, vol. 14443, pp. 361–389. Springer (2023)
 41. Wegener, I.: *The Complexity of Boolean Functions*. Wiley (1991)