IACR Communications in Cryptology
ISSN 3006-5496, Vol. 1, No. 1, 24 pages.

https://doi.org/10.62056/an5txol7

Bit Security as Cost to Demonstrate Advantage

Keewoo Lee

UC Berkeley, USA

Abstract. We revisit the question of what the definition of bit security should be, previously answered by Micciancio-Walter (Eurocrypt 2018) and Watanabe-Yasunaga (Asiacrypt 2021). Our new definition is simple, but (i) captures both search and decision primitives in a single framework like Micciancio-Walter, and (ii) has a firm operational meaning like Watanabe-Yasunaga. It also matches intuitive expectations and can be well-formulated regarding Hellinger distance. To support and justify the new definition, we prove several classic security reductions with respect to our bit security. We also provide pathological examples that indicate the ill-definedness of bit security defined in Micciancio-Walter and Watanabe-Yasunaga.

Keywords: bit security · security definition · Hellinger distance

1 Introduction

Bit security (a.k.a. security level) is a central concept in cryptography, which bridges asymptotic and concrete regimes. Bit security summarizes complex security descriptions of a concrete instantiation of a cryptographic scheme in a single number, being a simple enough measure for the level of security. Whereas the asymptotic approach does not provide any guidance on concrete parameter selection, bit security helps us choose an appropriate set of parameters to guarantee a certain level of security when deploying cryptographic schemes. When we say a scheme has λ-bit security, we roughly expect that it costs more than 2^λ resources to *break* the scheme or that the scheme is as secure as its *idealized version* with a λ-bit secret key.[1] However, despite its importance, we still do not have a well-accepted formal definition of bit security.

Conventional Definition

The most common definition of bit security is $\min \log(T/\varepsilon)$. Here, the minimum is taken over all possible adversaries $\mathcal{A}$, T is the cost (e.g. runtime) of $\mathcal{A}$, and ε is the advantage of $\mathcal{A}$. The definition captures trade-offs between cost and advantage for an idealized primitive with a λ-bit secret key. Two trivial extreme attacks are (i) brute-force search with $T = 2^\lambda$ and $\varepsilon = 1$ and (ii) guessing at random with $T = 1$ and $\varepsilon = 1/2^\lambda$.

Another intuition behind the conventional definition is the following. When $\mathcal{A}$ (with cost T and advantage ε) is given, we can run $\mathcal{A}$ for $N \approx 1/\varepsilon$ times to obtain an *amplified* adversary with cost $N \cdot T \approx T/\varepsilon$ and advantage $1 - (1-\varepsilon)^N \approx N \cdot \varepsilon \approx 1$. That is, when such an adversary $\mathcal{A}$ is given, we can *break* the scheme with a cost of roughly T/ε.

However, the above intuitions work only for (certain) search primitives (e.g., one-way functions). In particular, brute-force search or probability amplification is not allowed for decision primitives (e.g., pseudorandom generators). Moreover, we quantify the advantage differently for decision and search primitives, namely $\varepsilon = |P - 1/2|$ for decision primitives and $\varepsilon = P$ for search primitives, where P is the success probability. Thus, using the

E-mail: keewoo.lee@berkeley.edu (Keewoo Lee)

[1] The ambiguity here is how we define the terms *break* and *idealized version*.

Received: 2023-12-29 Accepted: 2024-03-05

same definition of bit security $\min \log(T/\varepsilon)$ for both types of primitives sounds already problematic. However, this is the widely used definition in the literature. Rather obviously, this conventional definition led to several paradoxical situations, such as the following.

Peculiar Case of Linear Test against PRG

It is folklore, which goes back at least to [AGHP90], that there is a non-uniform[2] attack (linear tests) against pseudorandom number generators (PRG) with λ-bit seed, which achieves advantage $\Omega(2^{-\lambda/2})$ in time $O(\lambda)$. Thus, according to the conventional definition of bit security, a PRG with λ-bit seed can guarantee not much more than $\lambda/2$-bit security. This contradicts our expectation that λ-bit security of a PRG reflects the security of the ideal PRG with λ-bit seed.

Peculiar Case of Distribution Approximation

When constructing cryptographic schemes (especially in lattice-based cryptography [Reg05, Pei16]), we often make use of certain distributions (e.g., discrete Gaussian). That is, sampling from a particular distribution is often an essential part of executing cryptographic schemes. And their security proofs assume an ideal situation where we can sample the distributions exactly. In actual implementations, however, we can only sample from an approximate distribution due to limited resources.

The question is how these approximations affect the security of schemes. In terms of the statistical distance (a.k.a. total variation distance), the standard measure in cryptography, it is an easy fact that λ-bit precision is sufficient to maintain λ-bit security. While this sounds quite natural already, ambitious researchers have proved that it is enough to achieve $\lambda/2$-bit closeness with respect to other *nice* divergences (e.g., Rényi [PDG14, BLL$^+$15], max-log [MW17]), yielding much better parameters for practical uses.

However, all mentioned results apply only to search primitives,[3] and a corresponding result for decision primitives has eluded researchers. The paradox is that it is generally believed that the security of encryption schemes (which is a decision primitive) is more robust against approximation errors than that of signature schemes (which is a search primitive).[4]

Ad Hoc Definition

Observing the peculiar cases, the definition of bit security for decision primitives seems *right* to be *doubled*, i.e., $\min \log(T/\varepsilon^2)$. We remark that this *ad hoc* definition was considered by classic works [GL89, HILL99] and is often used in the community without satisfactory understanding. That is, the questions remain: What is the source of this *quadratic gap* between decision and search primitives? What is the *right* definition for bit security?

Previous Approaches

Micciancio-Walter [MW18] explicitly pointed out these situations for the first time and provided a general formal definition of bit security; their definition resolves the above peculiar cases and captures both search and decision primitives in a single framework. The approach of Micciancio-Walter was to consider a general cryptographic game in which an adversary has to guess an n-bit string and define a general advantage of an adversary that captures both search (for large n) and decision (for $n = 1$) games, building on

[2] We note that there are debates on whether (e.g. [KM13]) and how (e.g. [BL13]) non-uniform adversaries should be considered in cryptography.

[3] We excluded [BLL$^+$15], which also considers decision primitives, since the result requires decision problems to satisfy a specific property called *public sampleability*.

[4] See, e.g., Section 4 of [ADPS16].

concepts from information theory. However, in the definition, they introduce a hypothetical random variable that lacks intuitive meaning without a satisfactory explanation. (Refer to Section 5.3 for details.)

Watanabe-Yasunaga [WY21] pointed out this weakness of Micciancio-Walter as a lack of *operational meaning* and provided another definition as cost for winning certain games with high probability — which also resolves the peculiar cases mentioned above and has an operational meaning by nature. However, they defined the games *qualitatively* differently for search and decision primitives, losing the generality that Micciancio-Walter sought. (Refer to Section 5.4 for details.)

1.1 Our Contributions

Our main result is a new definition of bit security (Def. 11). Our definition is so simple that we can put it in plain language: We define bit security as the cost to *demonstrate* advantage of adversaries. That is, we measure the total work done by an adversary to allow an observer to distinguish it from a *dummy* adversary by observing wins and loses while repeating games. Our simple definition (i) captures both search and decision primitives in a single framework like Micciancio-Walter [MW18] and (ii) has a firm operational meaning like Watanabe-Yasunaga [WY21]. Indeed, our definition also resolves the peculiar cases introduced above, matching the intuitive expectations (Remark 10 and Thm. 2). Moreover, our bit security can be well-formulated in terms of Hellinger distance, supporting the practical usability of our definition (Thm. 1 and Def. 13).

Besides, to support and justify our new definition of bit security, we:

- prove several security reductions with respect to our definition. Our proofs are arguably simpler and more intuitive than the previous proofs of [MW18, WY21]. Namely, we prove:
  - a theorem on distribution approximation in cryptographic schemes. Our theorem states that, with respect to the Hellinger distance, $\lambda/2$-bit precision is sufficient to maintain λ-bit security for *any* security games. This resolves the peculiar case introduced above. Our proof is much shorter than the previous proofs, leveraging nice properties of the Hellinger distance. (Section 4.1)
  - the hybrid argument. Our proof is essentially the same as the conventional proof, whereas [MW18, WY21] had to develop new techniques. This is due to the structural similarity between the conventional bit security and ours. In particular, like the conventional definition, our definition of bit security depends only on the success probability of adversaries and is independent of other information. (Section 4.2)
  - natural decision-to-search reductions. This includes the reductions from PRG to OWF, from DDH to CDH, and from IND-CPA to OW-CPA. Our reductions are *tight*, like the conventional proofs. (Section 4.3)
- point out several weaknesses of the previous definitions by [MW18, WY21].
  - Their definitions only cover security games with specific structures. In particular, their definitions do not even cover the EUF-CMA game for digital signature schemes. In contrast, our definition of security games is as inclusive as possible: Our framework captures all security games covered by the definitions of falsifiable assumptions of Gentry-Wichs [GW11]. (Section 5.1)
  - According to their definitions, search primitives always have *finite* bits of security. This circumstance is arguably counter-intuitive considering the presence of unconditionally secure search primitives, e.g., information-theoretic MAC. In

contrast, according to our definition, cryptographic schemes with information-theoretic security always satisfy ∞-bit security. (Section 5.5)

- Under their definitions, there are pathological examples where two security games, G and G', are essentially the same in common sense, but G is defined as a decision game and G' is defined as a search game. Moreover, the bit security of G and G' differ in their definitions. (Section 5.2 and 5.5)

1.2 Overview

The main body of this paper consists of three parts: Definitions (Section 3), Theorems (Section 4), and Comparisons (Section 5). The essence of this paper is Section 3, where our bit security is defined. Other sections are to support and justify the new definition.

In Section 3, we first formally define security game (Def. 3) to clarify the scope of our new definition of bit security. Then, bit security is formally defined as the cost to *demonstrate* advantage (Def. 11), leveraging a *meta*-game which we call *advantage observation game* (Def. 10). We also show that our bit security can be tightly estimated in terms of Hellinger distance (Thm. 1). We make several remarks on this new definition, including how it relates to the conventional bit security (Remark 9) and how it embraces the quadratic gap between decision and search primitives (Remark 10).

In Section 4, we prove several security reductions concerning our bit security: a distribution approximation theorem (Section 4.1), the hybrid argument (Section 4.2), and decision-to-search reductions (Section 4.3). In Section 5, we review the previous definitions of bit security proposed by Micciancio-Walter [MW18] and Watanabe-Yasunaga [WY21]. Then, we compare their definitions with our new definition. This section also points out several weaknesses of the previous definitions.

2 Preliminaries

2.1 Notations and Terminologies

We denote the logarithm to the base 2 by $\log(\cdot)$ and the one to the base e by $\ln(\cdot)$. We use standard arithmetic over extended non-negative real numbers $[0, \infty]$, i.e. $\frac{a}{0} = \infty$ for $a \in (0, \infty)$. The minimum of the empty set is defined to be ∞. We denote the total variation distance and Hellinger distance by $d_{\mathrm{TV}}(,)$ and $d_{\mathrm{H}}(,)$, respectively (Section 2.2). For notational convenience, we often identify the Bernoulli distribution $\mathcal{B}(p)$ with the probability p itself. For example, we use $d_{\mathrm{TV}}(p, q)$ and $d_{\mathrm{H}}(p, q)$ in place of $d_{\mathrm{TV}}\big(\mathcal{B}(p), \mathcal{B}(q)\big)$ and $d_{\mathrm{H}}\big(\mathcal{B}(p), \mathcal{B}(q)\big)$. We do not strictly distinguish the terms *hardness* and *security* and often use them interchangeably. For an algorithm $\mathcal{A}$, we denote its *cost* by $T_{\mathcal{A}}$.[5] We denote the complement of a relation R by $\bar{R}$.

2.2 Statistical Distances

In this section, we recall definitions and a few properties of two statistical distances: *total variation distance* and *Hellinger distance*. For proofs and detailed discussions, please refer to, e.g., [PW].

Definition 1 (Total Variation Distance)**.** For two discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ on the same domain $\mathcal{X}$, we denote and define their *total variation distance* (a.k.a. *the* statistical

[5] The primary cost model considered in this work is the time complexity. However, as long as it grows linearly with *repetitions* of the algorithm, any cost model can be used. For example, our framework (Def. 11) also makes sense with query complexity.

distance) as follows.

$$d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1) = \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} \left| \mathcal{D}_0(x) - \mathcal{D}_1(x) \right|$$

Definition 2 (Hellinger Distance)**.** For two discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ on the same domain $\mathcal{X}$, we denote and define their *Hellinger distance* as follows.

$$d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1) = \frac{1}{\sqrt{2}} \cdot \sqrt{\sum_{x \in \mathcal{X}} \left(\sqrt{\mathcal{D}_0(x)} - \sqrt{\mathcal{D}_1(x)} \right)^2}$$

Proposition 1 (Properties of Hellinger Distance)**.** *Let $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2$ be discrete distributions on the same domain $\mathcal{X}$.*

(a) Triangle Inequality: *The Hellinger distance is a metric. In particular, the following inequality holds.*

$$d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_2) \le d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1) + d_{\mathrm{H}}(\mathcal{D}_1, \mathcal{D}_2)$$

(b) Data-Processing Inequality: *Squared Hellinger distance is an f-divergence. In particular, the following inequality holds for any function $g : \mathcal{X} \to \mathcal{Y}$.*

$$d_{\mathrm{H}}\left(\mathcal{D}_0 \circ g^{-1}, \mathcal{D}_1 \circ g^{-1}\right) \le d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)$$

(c) Strong Decomposition Property on Product Distributions: *For any positive integer N, the following equality holds.*

$$1 - d_{\mathrm{H}}(\mathcal{D}_0^N, \mathcal{D}_1^N)^2 = \left(1 - d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2\right)^N$$

(d) Relation with Total Variation Distance: *The following inequalities hold.*

$$1 - \sqrt{1 - d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1)^2} \le d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2 \le d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1)$$

3 Definitions

3.1 General Security Game

We first formally define the *security game* to clarify the scope of our new definition of bit security. Our framework is abstract enough to capture every *game-based* security definition in the cryptography literature. The definition has already implicitly appeared in the definition of *falsifiable assumption* of Gentry-Wichs [GW11] (See also [Nao03]).

Definition 3 (Security Game)**.** A *security game* $G = (X, L)$ consists of an interactive *challenger* X and a decidable *winning condition* $L \subset \{0,1\}^*$.[6] The game is played by an adversary $\mathcal{A}$ interacting with X. During the game, if $\mathtt{view}_X \in L$, X outputs a special symbol `win` and we say $\mathcal{A}$ wins G. Here, $\mathtt{view}_X \in \{0,1\}^*$ is the *view* of the game from the perspective of X, i.e. the transcript and randomness used by X.

Remark 1 (Comparison). Our definition of security games is as inclusive as possible. We do not put any restrictions on the structure of the games. Our framework captures all security games covered by the definitions of *falsifiable assumptions* of Gentry-Wichs [GW11] and thus *complexity assumptions* of Goldwasser-Kalai [GK16]. On the other hand, previous frameworks of [MW18] and [WY21] only capture certain types of games, as they assert specific structures to the games. In particular, they do not include the very basic EUF-CMA game for signature schemes (Example 5). For a detailed discussion, refer to Section 5.1.

[6] We may include the winning condition L into the description of challenger X as in [GW11]. However, for easier comparison with frameworks of [MW18, WY21], we separate the description of L from that of X.

3.1.1 Examples

We give examples of security games regarding our framework. We also define specific classes of games for later discussions. Readers may skip the examples and come back when needed.

Definition 4 (Decision Game)**.** A *decision game* is a security game $G = (X, L)$ which has a certain structure on X and L as follows:

1. (Challenge) At the beginning of the game, the challenger X chooses a uniform random challenge bit $b \in \{0, 1\}$.
2. (Query) The adversary $\mathcal{A}$ is allowed to send certain queries to X. Whenever X receives a legitimate query, it sends a corresponding response to $\mathcal{A}$.
3. (Answer) The game ends when $\mathcal{A}$ sends its answer $b' \in \{0, 1\}$ to X.
4. (Winning Condition) $\mathcal{A}$ wins the game if $b = b'$.

Example 1 (Pseudorandomness)**.** For a *pseudorandom generator (PRG)* $f : \{0, 1\}^\ell \to \{0, 1\}^m$, we define its *pseudorandomness* by a decision game where the only allowed query for an adversary is to send a special symbol `sample` to the challenger. Whenever the challenger receives `sample`, it responds with $y = f(x)$ for a uniform random $x \in \{0, 1\}^\ell$ when $b = 0$ and a uniform random $y \in \{0, 1\}^m$ when $b = 1$.

Example 2 (IND-CPA[7])**.** For a public-key encryption scheme $(\mathtt{Gen}, \mathtt{Enc}, \mathtt{Dec})$, we define its *IND-CPA* security by the following decision game: Before the first query, the challenger runs `Gen` and sends the public key to the adversary. The only allowed query for an adversary is to send a special symbol `LR` to the challenger together with messages m_0 and m_1. Whenever the challenger receives $(\mathtt{LR}, m_0, m_1)$, it responds with $\mathtt{Enc}(m_b)$.

Example 3 (DDH)**.** For a cyclic group $\mathbb{G}$ and its generator g, we define the decisional Diffie-Hellman (DDH) game on $(\mathbb{G}, g)$ as the following decision game: The only allowed query for an adversary is to send a special symbol `sample` to the challenger. Whenever the challenger receives `sample`, it responds with (g^x, g^y, g^z) where $z = xy$ with uniform random x, y when $b = 0$ and x, y, z are all uniform random when $b = 1$.

Definition 5 (Distribution Distinguishing Game)**.** The *distribution distinguishing game* for distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ is a decision game, where the only allowed query for an adversary $\mathcal{A}$ is to send a special symbol `sample` to X. Whenever X receives `sample`, it draws a sample from the distribution $\mathcal{D}_b$ and sends the result to $\mathcal{A}$.

The following examples belong to a class so-called *search games*, although we do not precisely define search games in this paper. (See Remark 3.)

Example 4 (One-wayness)**.** For a *one-way function (OWF)* $f : \{0, 1\}^\ell \to \{0, 1\}^m$, we define its *one-wayness* by the following security game: At the beginning of the game, the challenger chooses uniform random $x \in \{0, 1\}^\ell$ and sends $y = f(x)$ to the adversary. The adversary sends an answer $x' \in \{0, 1\}^\ell$ and it wins the game if $y = f(x')$.

Example 5 (EUF-CMA)**.** For a digital signature scheme $(\mathtt{Gen}, \mathtt{Sign}, \mathtt{Verify})$, we define its *EUF-CMA* security by the following game: At the beginning of the game, the challenger runs `Gen` and sends the public key `pk` to the adversary. The only allowed query for an adversary is to send a special symbol `Sign` to the challenger together with a message m. When the challenger receives $(\mathtt{Sign}, m)$, it responds with a signature of m. The adversary sends a pair (m', σ') and it wins the game if $\mathtt{Verify}_{\mathtt{pk}}(m', \sigma')$ is `true` and m' was never queried by the adversary.

[7] We adopt *Left-Or-Right* formalization of IND-CPA, following Bellare-Rogaway [BR05] and Rosulek [Ros21]. (See also Katz-Lindell [KL14] and Goldreich [Gol04].) The formalization corresponds to FTG-CPA (Find-Then-Guess) of [BDJR97].

Example 6 (CDH)**.** For a cyclic group $\mathbb{G}$ and its generator g, we define the computational Diffie-Hellman (CDH) game on $(\mathbb{G}, g)$ as follows: At the beginning of the game, the challenger sends (g^x, g^y) to the adversary where x, y are uniform random. The adversary sends an answer $g^z \in \mathbb{G}$ and it wins the game if $g^z = g^{xy}$.

Example 7 (OW-CPA)**.** For a public-key encryption scheme $(\mathtt{Gen}, \mathtt{Enc}, \mathtt{Dec})$, we define its *OW-CPA* security by the following game: At the beginning of the game, the challenger runs `Gen` and chooses a message m randomly. Then, it sends $\mathtt{Enc}(m)$ and the public key to the adversary. The adversary sends an answer m' and it wins the game if $m' = m$.

3.2 Baseline Probability

Next, we define *baseline probability* of a game that plays an important role in refining the definition of (conventional) advantage and describing our new definition of bit security. The baseline probability of a game is the maximal success probability of *dummy* adversaries who do not learn anything while playing the game with the challenger.

Definition 6 (Success Probability)**.** Let G be a security game. We denote and define the *success probability* of an adversary $\mathcal{A}$ against G as the following.

$$P_{\mathcal{A}}^{G} = \Pr\left[\ \mathcal{A} \text{ wins } G\ \right]$$

Definition 7 (Dummy Adversary)**.** An adversary $\mathcal{A}$ against a security game G is called *dummy* if messages sent from $\mathcal{A}$ to the challenger do not depend on any previous messages that $\mathcal{A}$ has received. That is, the outputs of a dummy adversary are independent of the randomness of the challenger.

Remark 2. We note that the concept of dummy adversaries is not intended to capture every *trivial* attack. For an extreme example, consider the following game: At the beginning of the game, the challenger chooses uniform random $s \in \{0,1\}^{128}$ and sends s itself to the adversary. The adversary sends an answer $s' \in \{0,1\}^{128}$ and it wins the game if $s = s'$. Then, an adversary has a trivial strategy of just forwarding the received message. However, such a *trivial* adversary is not *dummy* according to our definition. Looking ahead, a security game will have small bits of security if there are (trivial) attacks performing much better than dummy adversaries. (See Remark 5.)

Definition 8 (Baseline Probability)**.** We denote and define the *baseline probability* of a security game G as the following.

$$P_{\emptyset}^{G} = \max_{\mathcal{A}:\text{dummy}} P_{\mathcal{A}}^{G}$$

We call a dummy adversary $\mathcal{A}$ against G a *baseline adversary* if $P_{\mathcal{A}}^{G} = P_{\emptyset}^{G}$ holds.

We also define baseline probability for cases where available resources are limited.

Definition 9 (Bounded Baseline Probability)**.** An adversary $\mathcal{A}$ is called *T-bounded* if $T_{\mathcal{A}} \leq T$ holds. We denote and define *T-bounded baseline probability* of G as the following.

$$P_{\emptyset}^{G}[T] = \max_{\substack{\mathcal{A}:\text{dummy} \\ \&\ T\text{-bounded}}} P_{\mathcal{A}}^{G}$$

We call T-bounded dummy adversary $\mathcal{A}$ against G a *T-baseline adversary* if $P_{\mathcal{A}}^{G} = P_{\emptyset}^{G}[T]$.

3.2.1 Examples

To describe how our definitions apply to security games, we compute some baseline probabilities. They will be referred to in later discussions.

Example 8 (Decision Game)**.** In a decision game (Def. 4), an adversary wins the game if it correctly guesses the challenge bit $b \in \{0,1\}$ chosen by the challenger during the game. However, by definition, dummy adversaries do not learn anything during the game. Thus, a dummy adversary cannot do better than a random guess on $\{0,1\}$, and the baseline probability $P^G_\emptyset[T]$ of a decision game G is $1/2$ for any T.

Example 9 (One-wayness)**.** In the one-wayness game for $f : \{0,1\}^\ell \to \{0,1\}^m$ (Example 4), an adversary wins the game if it correctly outputs x' such that $f(x') = y$, where y is chosen by the challenger during the game. Thus, a dummy adversary cannot do much better than a random guess on $\{0,1\}^\ell$ under reasonable assumptions. For example, if f maps at most k inputs to an output (i.e., at most k-to-one), the baseline probability of the game is not greater than $k/2^\ell$.

Example 10 (EUF-CMA)**.** In the EUF-CMA game for a signature scheme (Example 5) with *signature space* $\mathcal{S}$, an adversary wins the game only if it correctly outputs a pair of message and signature (m', σ') such that $\texttt{Verify}_{\texttt{pk}}(m', \sigma')$ is $\texttt{true}$, where $\texttt{pk}$ is chosen by the challenger during the game. Thus, a dummy adversary cannot do much better than a random guess on $\mathcal{S}$ for σ' under reasonable assumptions. For example, if there are at most k valid signatures for a pair of a public key and a message, the baseline probability of the EUF-CMA game is not greater than $k/|\mathcal{S}|$.

Example 11 (CDH)**.** In the CDH game on $(\mathbb{G}, g)$ (Example 6), an adversary wins the game if it correctly outputs the element g^{xy} where x, y are chosen by the challenger during the game. Thus, the baseline probability of the CDH game is $1/|\mathbb{G}|$.

Example 12 (OW-CPA)**.** In the OW-CPA game for a public-key encryption scheme (Example 7) with message space $\mathcal{M}$, an adversary wins the game if it correctly outputs a message m' such that $m' = m$, where m is chosen by the challenger during the game. Thus, the baseline probability of the CDH game is $1/|\mathcal{M}|$.

Remark 3 (Search Game)*.* Unlike decision games (Def. 4), we do not precisely define search games in this work. Previous definitions of search games in [MW18, WY21] assert specific structures to security games, leading to several problematic situations. (For detailed discussions, refer to Section 5.2.) Meanwhile, Example 9 - 12 demonstrate that search primitives are expected to have negligible baseline probabilities. We take this extremely small baseline probability as a fuzzy characterization of search games. Looking ahead, our new definition of bit security only depends on baseline probability and is independent of any other features of games. Thus, characterizing search games with extremely small baseline probability would suffice to see how our definition applies to search games.

3.2.2 Conventional Advantage

As said, under our framework (Def. 8), we can refine and unite conventional definitions of *advantage*, where $|P^G_\mathcal{A} - 1/2|$ is used for decision games and $P^G_\mathcal{A}$ is used for search games. We define the (conventional) advantage of an adversary as the (rectified) difference between the success probability of the adversary and the baseline probability.

Example 13 (Conventional Advantage)**.** We denote and define the *(conventional) advantage* of $\mathcal{A}$ against G as the following.

$$\mathrm{adv}^G(\mathcal{A}) = \max\left\{ P^G_\mathcal{A} - P^G_\emptyset[T_\mathcal{A}],\ 0 \right\}$$

For a decision game, where the baseline probability is 1/2 (Example 8), the advantage is $\max\left\{P^G_{\mathcal{A}} - 1/2,\ 0\right\}$ according to our definition. Thus, our definition matches with the conventional definition for decision games $|P^G_{\mathcal{A}} - 1/2|$ (when $P^G_{\mathcal{A}} \geq P^G_{\emptyset}[T_{\mathcal{A}}]$ holds[8]).

For a search game, note that the baseline probability is expected to be extremely small (Remark 3). Then, we can easily see that our definition approximately matches the conventional definition for search games $P^G_{\mathcal{A}}$ (when $P^G_{\emptyset}[T_{\mathcal{A}}]$ is sufficiently small compared with $P^G_{\mathcal{A}}$ to be approximated as zero).

Remark 4 (Reformulation of Conventional Advantage). We note that our definition of conventional advantage can be reformulated in terms of total variation distance (a.k.a. statistical distance) between two Bernoulli distributions as the following. (Refer to Section 2 for notations.) This reformulation will later be used to see how our new definitions of advantage and bit security relate to conventional definitions (Remark 9).

$$\mathrm{adv}^G(\mathcal{A}) = \begin{cases} d_{\mathrm{TV}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right) & \text{if } P^G_{\mathcal{A}} \geq P^G_{\emptyset}[T_{\mathcal{A}}] \\ 0 & \text{if } P^G_{\mathcal{A}} < P^G_{\emptyset}[T_{\mathcal{A}}] \end{cases}$$

3.2.3 Information-Theoretic Security

Under our framework of baseline probability, we can obtain a natural and simple characterization of information-theoretic security.

Example 14 (Information-Theoretic Security)**.** By definition, information-theoretic security guarantees that adversaries gain no information at all during the security game. In terms of our framework, this corresponds to the situation where all adversaries are essentially dummies. That is, we can characterize (or even define) information-theoretic security for game G as the condition where $P^G_{\mathcal{A}} \leq P^G_{\emptyset}[T_{\mathcal{A}}]$ and thus $\mathrm{adv}^G(\mathcal{A}) = 0$ hold for any adversary $\mathcal{A}$ (Example 13). This natural and simple characterization of information-theoretic security cannot be obtained if we adopt $P^G_{\mathcal{A}}$ as the definition of advantage for search primitives. Using such a definition, search primitives always yield negligible but positive advantages for some adversaries.

3.3 Bit Security as Cost to Demonstrate Advantage

In this section, we propose a new definition of bit security. Intuitively, we define bit security as the cost to demonstrate that an adversary $\mathcal{A}$ indeed has an advantage over dummy adversaries for game G, i.e., $P^G_{\mathcal{A}} > P^G_{\emptyset}[T_{\mathcal{A}}]$ holds,[9] when $\mathcal{A}$ is given in a black-box manner. In other words, it can also be described as the cost to enable a third party to empirically experience or *observe* the advantage of an adversary.

For a formal presentation, we first introduce a *meta*-game, which we call *advantage observation game.* The advantage observation game for adversary $\mathcal{A}$ against game G is a decision game where an adversary $\mathcal{B}$ tries to distinguish $\mathcal{A}$ from dummy adversaries. If $\mathcal{B}$ wins the game, then we may say it *observed* the advantage of $\mathcal{A}$ during the game.

[8] The intuition of defining advantage as the absolute value $|P^G_{\mathcal{A}} - 1/2|$ for decision games is that we can always transform an adversary with success probability $P^G_{\mathcal{A}}$ into one with $1 - P^G_{\mathcal{A}}$ by switching the output. However, this transformation does not apply to search games. We consider this issue as a roadblock to a unified definition and thus use a *rectified* version rather than an absolute value. This issue and treatment were previously considered by Bernstein-Hülsing when defining Decisional Second-Preimage Resistance (DSPR) in [BH19]. DSPR is defined by an *unbalanced* decision game whose baseline probability is not 1/2.

[9] This condition captures the idea that there are two sources of *advantage*, namely success probability and cost. For example, if $P^G_{\mathcal{A}_1} = P^G_{\mathcal{A}_2}$ but $T_{\mathcal{A}_1} < T_{\mathcal{A}_2}$, we should admit that $\mathcal{A}_1$ has advantage over $\mathcal{A}_2$. Note that the conventional definition of advantage ($P^G_{\mathcal{A}}$ for search primitives and $|P^G_{\mathcal{A}} - 1/2|$ for decision primitives) only considers the success probability.

Definition 10 (Advantage Observation Game)**.** Let $\mathcal{A}$ be an adversary against a security game $G = (X, L)$. The *advantage observation game* for $\mathcal{A}$ against G is denoted as $\hat{G}_{\mathcal{A}}$ and defined as follows (See also Figure 1.):

1. (Setting) Let $\mathcal{A}_0$ be $\mathcal{A}$ and $\mathcal{A}_1$ be a $T_{\mathcal{A}}$-baseline adversary against G (Def. 9).
2. (Challenge) At the beginning of the game, the challenger $\hat{X}$ chooses a uniform random challenge bit $b \in \{0, 1\}$.
3. (Query) The only allowed query for an adversary $\mathcal{B}$ is to send a special symbol `sample` to $\hat{X}$. Whenever $\hat{X}$ receives `sample`, it invokes the game G with $\mathcal{A}_b$ and sends the result (`win` or `lose`) to $\mathcal{B}$.
4. (Answer) The game ends when $\mathcal{B}$ sends its answer $b' \in \{0, 1\}$ to $\hat{X}$.
5. (Winning Condition) $\mathcal{B}$ wins the game if $b = b'$ and $P^G_{\mathcal{A}} \geq P^G_{\emptyset}[T_{\mathcal{A}}]$.

We denote the query complexity of $\mathcal{B}$ as $N_{\mathcal{B}}$.

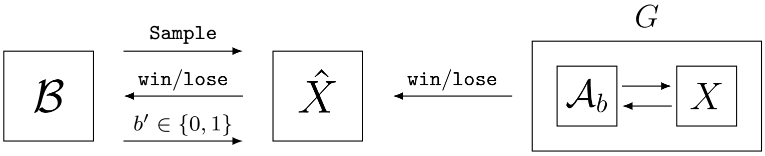

Figure 1: A Description of the Advantage Observation Game

Remark 5. We note that the advantage observation game is defined with respect to a *dummy* adversary, which is not intended to capture every *trivial* attack. (See Remark 2.)

Remark 6. We note that an adversary can never win the advantage observation game $\hat{G}_{\mathcal{A}}$ when $P^G_{\mathcal{A}} < P^G_{\emptyset}[T_{\mathcal{A}}]$. This additional restriction captures the fact that there is no advantage to observe in such a case. (See also Footnote 8.)

We now define bit security in terms of the advantage observation game. As said, we define bit security as the cost to demonstrate the advantage of the adversary $\mathcal{A}$ against the game G, i.e., the total cost of invoking $\mathcal{A}$ multiple times[10] to allow adversary $\mathcal{B}$ to win the advantage observation game $\hat{G}_{\mathcal{A}}$ with high probability. To elaborate, we measure the total cost as the cost $T_{\mathcal{A}}$ of $\mathcal{A}$ multiplied by the query complexity $N_{\mathcal{B}}$ of $\mathcal{B}$ against the game $\hat{G}_{\mathcal{A}}$.

Definition 11 (Bit Security as Cost to Demonstrate Advantage)**.** For any security game G, we denote and define its *(demonstration) bit security* (with respect to error probability $0 < \delta < 1/2$) as the following.

$$BS^{\delta}_{\mathrm{Dem}}(G) = \min_{\mathcal{A},\mathcal{B}} \left\{ \log\left(T_{\mathcal{A}} \cdot N_{\mathcal{B}}\right) : P^{\hat{G}_{\mathcal{A}}}_{\mathcal{B}} \geq 1 - \delta \right\}$$

Remark 7 (Comparison). Our unified definition of bit security is independent of the structures of security games (e.g., decision or search), unlike Watanabe-Yasunaga [WY21]. At the same time, our definition enjoys a clear and firm operational meaning from the advantage observation game, unlike Micciancio-Walter [MW18]. Nonetheless, we will see how our definition still resolves hiccups of the conventional definition of bit security (Remark 10, Thm. 2). For a more detailed discussion, refer to Section 5.3 and 5.4.

[10] Recall that $\mathcal{A}$ is called multiple times by challenger $\hat{X}$ on demand of adversary $\mathcal{B}$ during the game $\hat{G}_{\mathcal{A}}$.

3.3.1 Estimation

Here, we examine demonstration bit security (Def. 11) in a quantitative manner. The goal is to find a more *computable* characterization of our definition. We first review a folklore fact in learning theory: the sample complexity of distinguishing discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ is $\Theta\big(1/d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2\big)$, where d_{H} is the Hellinger distance (Section 2.2) between two distributions. The statement and proof are more or less verbatim of [BY02, Can17] with extra care on constants behind the asymptotic expressions.

Proposition 2 (Sample Complexity Bounds)**.** *Let $N_\delta(\mathcal{D}_0, \mathcal{D}_1)$ denote the sample complexity of distinguishing discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ with probability at least $1-\delta$. That is, $N_\delta(\mathcal{D}_0, \mathcal{D}_1)$ is the minimum among query complexity of adversaries against the distribution distinguishing game on $\mathcal{D}_0$ and $\mathcal{D}_1$ (Def. 5) with success probability at least $1-\delta$. For $0 < \delta < 1/2$, we have the following bounds. The upper bound holds always, and the lower bound holds if $d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2 \le 1/2$.*

$$\frac{1}{4\ln 2} \cdot \frac{\ln(\frac{1}{4\delta(1-\delta)})}{d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2} \;\le\; N_\delta(\mathcal{D}_0, \mathcal{D}_1) \;\le\; \frac{\ln(\frac{1}{2\delta})}{d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2}$$

Proof. Refer to Appendix A. □

We now give an estimation of demonstration bit security (Def. 11) leveraging Prop. 2. The following theorem suggests that we can well-estimate the security in terms of Hellinger distance. (Refer to Section 2 for notations.)

Theorem 1 (Estimation of Bit Security)**.** *For $0 < \delta \le (2-\sqrt{3})/4$, we have the following estimation of the demonstration bit security, up to a small additive error α, satisfying $0 \le \alpha \le 1 + \log\ln(\frac{1}{2\delta})$.*

$$BS^{\delta}_{\mathrm{Dem}}(G) = \min_{\mathcal{A}:\ P^G_{\mathcal{A}} \ge P^G_{\emptyset}[T_{\mathcal{A}}]} \log\left(\ln(\frac{1}{2\delta}) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}\right) - \alpha$$

Proof. Let $\mathcal{B}^*$ be an adversary against the advantage observation game $\hat{G}_{\mathcal{A}^*}$, which has the minimal query complexity $N_{\mathcal{B}^*}$ among adversaries with success probability at least $1-\delta$. The theorem is easy to prove after dividing it into two cases.

Case 1: Suppose $d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2 \le 1/2$. By Prop. 2, we have the following bounds.

$$\frac{1}{4\ln(2)} \cdot \frac{\ln(\frac{1}{4\delta(1-\delta)})}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2} \;\le\; N_{\mathcal{B}^*} \;\le\; \frac{\ln(\frac{1}{2\delta})}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}$$

That is, the ratio of $\ln(\frac{1}{2\delta})/d_{\mathrm{H}}\left(P^G_{\mathcal{A}^*}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2$ to $N_{\mathcal{B}^*}$ is less than or equal to $4\ln(2) \cdot \ln(\frac{1}{2\delta})/\ln(\frac{1}{4\delta(1-\delta)})$, which is again not greater than $2\ln(\frac{1}{2\delta})$ for $0 < \delta \le (2-\sqrt{3})/4$.

Case 2: Suppose $d_{\mathrm{H}}\left(P^G_{\mathcal{A}^*}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2 > 1/2$. In this case, we can only use the upper bound of Prop. 2. Nonetheless, we also have a trivial lower bound: $1 \le N_{\mathcal{B}^*}$. Then, the ratio of $\ln(\frac{1}{2\delta})/d_{\mathrm{H}}\left(P^G_{\mathcal{A}^*}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2$ to $N_{\mathcal{B}^*}$ is not greater than $2\ln(\frac{1}{2\delta})$. □

Remark 8 (Tightness)*.* We note that the estimation of Thm. 1 is *tight* in the sense that the additive error is bounded by a double-logarithm of $1/\delta$. In particular, when $\delta = 2^{-128}$, the additive error is smaller than 7.5.

3.4 Bit Security in terms of Hellinger Distance

Although our definition of bit security based on *demonstration of advantage* (Def. 11) provides a nice operational meaning, there remain issues with its practical usability. In particular, the definition is parameterized by statistical significance δ, which might hinder the adoption of the definition. In this respect, we redefine bit security, dropping the dependency on the choice of δ and directly exploiting the Hellinger distance.

Definition 12 (Hellinger-Advantage)**.** For any security game G and adversary $\mathcal{A}$ against G, we denote and define the *Hellinger-advantage* of $\mathcal{A}$ against G as the following.

$$\mathrm{adv}^G_{\mathrm{H}^2}(\mathcal{A}) = \begin{cases} d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2 & \text{if } P^G_{\mathcal{A}} \geq P^G_{\emptyset}[T_{\mathcal{A}}] \\ 0 & \text{if } P^G_{\mathcal{A}} < P^G_{\emptyset}[T_{\mathcal{A}}] \end{cases}$$

Definition 13 (Hellinger-Bit Security)**.** For any security game G, we denote and define its *Hellinger-bit security* as the following.

$$BS_{\mathrm{H}^2}(G) = \min_{\mathcal{A}} \log\left(\frac{T_{\mathcal{A}}}{\mathrm{adv}^G_{\mathrm{H}^2}(\mathcal{A})}\right)$$

Regarding Thm. 1, our Hellinger-based definition can be understood as a *normalized* version of Def. 11 deleting the $\log(1/\delta)$ term. Thus, Hellinger-bit security (BS_{H^2}) is a more conservative measure compared to Def. 11 ($BS^{\delta}_{\mathrm{Dem}}$), when δ is reasonably small.

Remark 9 (Comparison with the Conventional Definition). Our Hellinger-based definitions share the same structure as the conventional definitions of advantage and bit security (Remark 4). That is, our definition only differs from the conventional definition in the choice of how to measure the difference between the success probability of an adversary and baseline probability. The conventional definition utilizes the total variation distance (a.k.a. statistical distance), whereas our definition utilizes the square of the Hellinger distance. While our new definition with an operational meaning (Section 3.3) resolves hiccups of the conventional definition (Remark 10, Thm. 2), this similarity in structures does not seriously harm existing proof outlines and techniques (Section 4).

Next, as examples, we compute the Hellinger-advantage of several security games. In particular, we apply our definition to decision and search primitives to demonstrate how our definition embraces the *quadratic gap* (See Introduction).

Example 15 (Decision Primitives)**.** Let G be a decision game (Def. 4). Then, we have $P^G_{\emptyset}[T] = 1/2$ (Example 8). In addition, let $\mathcal{A}$ be an adversary against G with *conventional* advantage $\varepsilon > 0$ (Example 13). That is, we have $P^G_{\mathcal{A}} = 1/2 + \varepsilon$. Then, we can compute the Hellinger-advantage of $\mathcal{A}$ against G as the following. The last equality is obtained from a Taylor approximation.

$$\begin{aligned} \mathrm{adv}^G_{\mathrm{H}^2}(\mathcal{A}) &= d_{\mathrm{H}}\left(\frac{1}{2} + \varepsilon,\ \frac{1}{2}\right)^2 \\ &= 1 - \frac{\sqrt{1+2\varepsilon}}{2} - \frac{\sqrt{1-2\varepsilon}}{2} \\ &= \frac{1}{2} \cdot \varepsilon^2 + O(\varepsilon^4) \end{aligned}$$

Example 16 (Search Primitives)**.** Let G be a *search* game. That is, we may assume $P^G_{\emptyset}[T]$ to be 0 (Remark 3). In addition, let $\mathcal{A}$ be an adversary against G with *conventional* advantage $\varepsilon > 0$ (Example 13). That is, we have $P^G_{\mathcal{A}} = \varepsilon$. Then, we can compute

Hellinger-advantage of $\mathcal{A}$ against G as the following. The last equality is obtained from a Taylor approximation.

$$\begin{aligned}\mathrm{adv}_{\mathrm{H}^2}^{G}(\mathcal{A}) &= d_{\mathrm{H}}(\ \varepsilon\ ,\ \ 0\ )^2 \\ &= 1-\sqrt{1-\varepsilon} \\ &= \frac{1}{2}\cdot\varepsilon + O(\varepsilon^2)\end{aligned}$$

Remark 10 (Decision/Search Primitives)*.* Through Example 15 and 16, we checked how our Hellinger-advantage embraces the quadratic gap of conventional bit security between decision and search primitives. That is, according to our definition, bit security of a decision (resp. search) primitive is roughly $\min\log(T/\varepsilon^2)$ (resp. $\min\log(T/\varepsilon)$), where ε is *conventional* advantage (Example 13). As the quadratic gap is the central question in the line of works, the definitions of previous works [MW18, WY21] also capture this gap. However, (i) unlike Watanabe-Yasunaga [WY21], we capture the gap in a single unified framework, and (ii) unlike Micciancio-Walter [MW18], our definition has a clear and firm operational meaning. For a more detailed discussion, refer to Section 5.3 and 5.4.

Example 17 (Information-Theoretic Security)**.** Let G be a security game with information-theoretic security. That is, for any adversary $\mathcal{A}$ against G, we have $P_{\mathcal{A}}^{G} \leq P_{\emptyset}^{G}[T_{\mathcal{A}}]$ (Example 14). Therefore, $\mathrm{adv}_{\mathrm{H}^2}^{G}(\mathcal{A}) = 0$ for all $\mathcal{A}$ and thus $BS_{\mathrm{H}^2}(G) = \infty$. This is something one might expect from definitions of advantage and bit security. However, according to previous definitions [MW18, WY21], some unconditionally secure primitives do not enjoy infinite bit security. For a more detailed discussion, refer to Section 5.5.

4 Theorems

In this section, we state and prove several security reductions with respect to Hellinger-bit security (Def. 13). We show how our new definition (i) resolves the peculiar situation of the conventional definition regarding distribution approximations (Section 4.1) and (ii) still provides the classic results proven under the conventional definition (Section 4.2, 4.3). In addition, throughout this section, we demonstrate that Hellinger-bit security is not too difficult to use compared to the conventional definition and often provides simpler and more intuitive proofs compared to previous definitions of [MW18, WY21]. These support the suitability of our new definition.

4.1 Distribution Approximation

We first prove that we can replace distributions in λ-bit secure games with a $\lambda/2$-bit close distribution (with respect to Hellinger distance) while preserving security. Our result holds regardless of the structures of the game (e.g., decision/search). This resolves the peculiar situation of the conventional definition, where the theorem was proved only for search primitives (See Introduction).

The proof outline is very intuitive as we use a standard trick with triangle inequality. Other parts easily follow from nice properties of Hellinger distance. We emphasize that we prove the theorem as a whole and do not divide cases into decision and search primitives.

Theorem 2 (Distribution Approximation)**.** *Let $G_0 = (X_0, L)$ and $G_1 = (X_1, L)$ be identical security games except that challenger X_0 uses distribution $\mathcal{D}_0$ at points where X_1 uses $\mathcal{D}_1$. Assume that the challengers sample from the distributions at most $c^2 \cdot T$ times when playing with adversaries with cost T.*[11] *If G_1 is λ-bit secure and $d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1) \leq 2^{-\lambda/2}$, then G_0 is $\big(\lambda - 2\log(2c+1)\big)$-bit secure, with respect to Hellinger-bit security.*

[11] For example, we can choose c^2 as the maximum number of accesses between a query and its response. In a reasonable situation, we can choose c^2 to be a small constant (e.g., 1 or 2).

Proof. We first note that any adversary $\mathcal{A}$ with cost T satisfies the following inequalities. The first inequality is from the data-processing inequality (Prop. 1(b)) together with the fact that $\mathcal{A}$ can learn information on at most $c^2 \cdot T$ samples.

$$\begin{aligned} d_{\mathrm{H}}(P_{\mathcal{A}}^{G_0}, P_{\mathcal{A}}^{G_1})^2 &\leq d_{\mathrm{H}}(\mathcal{D}_0^{c^2T}, \mathcal{D}_1^{c^2T})^2 && \text{(Prop. 1(b))} \\ &= 1 - \left(1 - d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2\right)^{c^2T} && \text{(Prop. 1(c))} \\ &\leq 1 - (1 - 2^{-\lambda})^{c^2T} \\ &\leq c^2 \cdot T \cdot 2^{-\lambda} \end{aligned}$$

Thus, we have the following inequality.

$$d_{\mathrm{H}}(P_{\mathcal{A}}^{G_0}, P_{\mathcal{A}}^{G_1}) \leq c \cdot \sqrt{\frac{T}{2^\lambda}}$$

Consider an adversary $\mathcal{A}$ against G_0 with a positive Hellinger advantage and let $\emptyset_b$ denote $T_{\mathcal{A}}$-baseline adversary (Def. 8) against G_b for $b = 0, 1$. Then, we have the following inequalities. The first inequality is from the definition of baseline adversary, and the second is from the standard triangle inequality (Prop. 1(a)). For the third inequality, the first and third terms are bounded from the above note, and the second term is bounded by the fact that G_1 is λ-bit secure.

$$\begin{aligned} d_{\mathrm{H}}(P_{\mathcal{A}}^{G_0}, P_{\emptyset_0}^{G_0}) &\leq d_{\mathrm{H}}(P_{\mathcal{A}}^{G_0}, P_{\emptyset_1}^{G_0}) \\ &\leq d_{\mathrm{H}}(P_{\mathcal{A}}^{G_0}, P_{\mathcal{A}}^{G_1}) + d_{\mathrm{H}}(P_{\mathcal{A}}^{G_1}, P_{\emptyset_1}^{G_1}) + d_{\mathrm{H}}(P_{\emptyset_1}^{G_1}, P_{\emptyset_1}^{G_0}) \\ &\leq \quad c \cdot \sqrt{\frac{T_{\mathcal{A}}}{2^\lambda}} \quad + \quad \sqrt{\frac{T_{\mathcal{A}}}{2^\lambda}} \quad + \quad c \cdot \sqrt{\frac{T_{\mathcal{A}}}{2^\lambda}} \\ &= (2c + 1) \cdot \sqrt{\frac{T_{\mathcal{A}}}{2^\lambda}} \end{aligned}$$

Thus, we have $T_{\mathcal{A}} / d_{\mathrm{H}}(P_{\mathcal{A}}^{G_0}, P_{\emptyset_0}^{G_0})^2 \geq 2^\lambda / (2c + 1)^2$ and the theorem follows. □

Remark 11 (Comparison). Previous works [MW18, WY21] also prove similar theorems. However, their multi-page proof consists of a handful of computations, even after borrowing some results from [MW17].[12] In particular, they had to prove the theorem by dividing the case into decision and search primitives.

On the other hand, our proof is much simpler and only uses standard techniques. In particular, our proof fits on a single page. Moreover, our proof is *unified* in the sense that we do not handle search and decision primitives separately. We believe this indicates that our Hellinger-bit security is a more suitable definition, compared to the previous works.

4.2 Hybrid Argument

We prove the *hybrid argument* with respect to our Hellinger-bit security. Our proof is essentially no different from the proof for conventional bit security. The only difference is that we lose roughly $2 \log n$-bits of security, whereas we lose $\log n$-bits in the conventional setting. This gap seems natural as our Hellinger advantage is roughly the *square* of the conventional advantage for decision games (Remark 10). The gap appears also in the previous works [MW18, WY21].

[12] We note that their theorem covers a certain *class* of divergences, whereas we prove the theorem only for the Hellinger distance.

Theorem 3 (Hybrid Argument)**.** *Let n be a positive integer. For $0 \le i, j \le n$, let $G_{i,j}$ be a decision game (Def. 4), where the challenger acts as an algorithm $\mathcal{C}_i$ if $b = 0$ and $\mathcal{C}_j$ if $b = 1$. If $G_{i,i+1}$ are all λ-bit secure for $0 \le i < n$, then $G_{0,n}$ is $(\lambda - 2\log n - \alpha)$-bit secure with respect to Hellinger-bit security. Here, $\alpha = \log(8 - 4\sqrt{2})$.*

Proof. Let $\mathcal{A}$ be an adversary against $G_{0,n}$ with success probability $1/2+\varepsilon$ with $0 < \varepsilon \le 1/2$, i.e. the conventional advantage (Example 13) of $\mathcal{A}$ is ε. This implies the existence of an interactive algorithm $\bar{\mathcal{A}}$ with cost $T_{\mathcal{A}}$ such that

$$\left| \Pr\left[\bar{\mathcal{A}}(\mathcal{C}_0) = 0\right] - \Pr\left[\bar{\mathcal{A}}(\mathcal{C}_n) = 0\right] \right| = 2\varepsilon.$$

Then, by the triangle inequality, we have

$$\sum_{i=0}^{n-1} \left| \Pr\left[\bar{\mathcal{A}}(\mathcal{C}_i) = 0\right] - \Pr\left[\bar{\mathcal{A}}(\mathcal{C}_{i+1}) = 0\right] \right| \ge 2\varepsilon.$$

Thus, for some $0 \le i^* < n$, we have

$$\left| \Pr\left[\bar{\mathcal{A}}(\mathcal{C}_{i^*}) = 0\right] - \Pr\left[\bar{\mathcal{A}}(\mathcal{C}_{i^*+1}) = 0\right] \right| \ge 2\varepsilon/n,$$

which implies the existence of an adversary $\mathcal{A}^*$ against G_{i^*,i^*+1} with conventional advantage greater than ε/n and cost $T_{\mathcal{A}}$. Since G_{i^*,i^*+1} is λ-bit secure, $T_{\mathcal{A}}/d_{\mathrm{H}}(1/2+\varepsilon/n, 1/2)^2 \ge 2^{\lambda}$ must hold. Also, since $0 < \varepsilon \le 1/2$, we have

$$\frac{d_{\mathrm{H}}\left(\frac{1}{2} + \frac{\varepsilon}{n},\ \frac{1}{2}\right)^2}{d_{\mathrm{H}}\left(\frac{1}{2} + \varepsilon,\ \frac{1}{2}\right)^2} \ge \frac{d_{\mathrm{H}}\left(\frac{1}{2} + \frac{1}{2n},\ \frac{1}{2}\right)^2}{d_{\mathrm{H}}\left(\frac{1}{2} + \frac{1}{2},\ \frac{1}{2}\right)^2} \ge \frac{1}{8 - 4\sqrt{2}} \cdot \frac{1}{n^2}.$$

Thus,

$$\frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(P_{\mathcal{A}}^{G_{0,n}},\ \frac{1}{2}\right)^2} = \frac{d_{\mathrm{H}}\left(\frac{1}{2} + \frac{\varepsilon}{n},\ \frac{1}{2}\right)^2}{d_{\mathrm{H}}\left(\frac{1}{2} + \varepsilon,\ \frac{1}{2}\right)^2} \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(\frac{1}{2} + \frac{\varepsilon}{n},\ \frac{1}{2}\right)^2} \ge \frac{1}{8 - 4\sqrt{2}} \cdot \frac{1}{n^2} \cdot 2^{\lambda}.$$

□

Remark 12 (Comparison). Our proof outline is identical to the conventional proof. This is due to the structural similarity of our definition of bit security to the conventional definition (Remark 9). On the other hand, the proof of Micciancio-Walter [MW18] needs to take care of *aborts* which play a major role in their definition (Def. 17). The proof of Watanabe-Yasunaga [WY21] is intuitive and natural but requires a different approach as their definition of bit security (Def. 19) depends also on distributional information other than the success probability.

4.3 Decision-to-Search Reductions

In the literature, there are several pairs of decision and search games, which allow natural decision-to-search security reductions. These reductions are proved in the same outline: To decide whether a distribution is structured or not, the reduction algorithm calls the given adversary who solves the search problem related to the structure. If the adversary succeeds, the reduction algorithm answers that the distribution is structured, since if the distribution were random the adversary would have failed with a high probability. This approach gives a *tight* reduction, in a sense that the reduction preserves the conventional advantage up to a constant factor.

We show that this frame still works under our new definition of bit security. However, to achieve *tight* reductions, extra work must be done. Unlike the conventional proof where the search adversary is called only once, we have to call the adversary several times to amplify the success probability. Otherwise, at worst, we may lose *half* of the bit security along the reduction: The conventional reduction preserves the conventional advantage ε, but our Hellinger-bit security is roughly $\min \log(T/\varepsilon^2)$ for decision primitives and $\min \log(T/\varepsilon)$ for search primitives (Remark 10). We first prove that a PRG (Example 1) is a OWF (Example 4). We again leverage the Hellinger distance and Prop. 2.

Theorem 4 (Pseudorandomness to One-wayness)**.** *Let $f : \{0,1\}^\ell \to \{0,1\}^m$ be a PRG. If f is λ-bit secure, then f is also $(\lambda-\alpha)$-bit secure as a OWF, with respect to the Hellinger-bit security. Here, $\alpha = 8 + \log(1 + T_f)$, where T_f denotes the cost for evaluating f.*

Proof. Let $\mathcal{A}$ be an adversary against the one-wayness game G on f, with success probability $\varepsilon := P^G_{\mathcal{A}}$ such that $P^G_{\mathcal{A}} > P^G_{\emptyset}[T_{\mathcal{A}}]$. We construct an adversary $\mathcal{A}'$ against the pseudorandomness game G' as follows: Whenever a query is made, $\mathcal{A}'$ runs $\mathcal{A}$ on the sample y and checks whether the output x' satisfies $y = f(x')$. The adversary records `Yes` if the condition is satisfied and `No` if not. Note that $\Pr[\texttt{Yes}|b = 0] = \varepsilon$ and $\Pr[\texttt{Yes}|b = 1] \le \frac{2^\ell}{2^m}\varepsilon \le \frac{1}{2}\varepsilon$.

Now the goal is to distinguish the two cases of $b = 0$ and $b = 1$ using this `Yes`/`No` distribution. By Prop. 2, to distinguish two cases with a probability of at least $1 - \delta$, it is sufficient to count the number of `Yes`'s among $N = \ln(\frac{1}{2\delta})/d_{\mathrm{H}}(\varepsilon, \frac{1}{2}\varepsilon)^2$ samples. For such an adversary $\mathcal{A}'$, we have the following inequalities.

$$\begin{aligned}
\frac{T_{\mathcal{A}'}}{d_{\mathrm{H}}\big(P^{G'}_{\mathcal{A}'}, P^{G'}_{\emptyset}[T_{\mathcal{A}'}]\big)^2} &\le \frac{N \cdot (T_{\mathcal{A}} + T_f)}{d_{\mathrm{H}}\big(1-\delta, \frac{1}{2}\big)^2} \\
&< \frac{N}{d_{\mathrm{H}}\big(1-\delta, \frac{1}{2}\big)^2} \cdot (1 + T_f) \cdot T_{\mathcal{A}} \\
&= \frac{\ln(\frac{1}{2\delta})}{d_{\mathrm{H}}\big(1-\delta, \frac{1}{2}\big)^2} \cdot \frac{d_{\mathrm{H}}\big(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\big)^2}{d_{\mathrm{H}}\big(\varepsilon, \frac{1}{2}\varepsilon\big)^2} \cdot (1 + T_f) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\big(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\big)^2}
\end{aligned}$$

Meanwhile, we can bound the second term as follows.

$$\frac{d_{\mathrm{H}}\big(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\big)^2}{d_{\mathrm{H}}\big(\varepsilon, \frac{1}{2}\varepsilon\big)^2} \le \frac{d_{\mathrm{H}}(\varepsilon, 0)^2}{d_{\mathrm{H}}\big(\varepsilon, \frac{1}{2}\varepsilon\big)^2} \le 6 + 4\sqrt{2}$$

In total, if we choose $\delta \approx 0.1$, we have the following bound.

$$\frac{T_{\mathcal{A}'}}{d_{\mathrm{H}}\big(P^{G'}_{\mathcal{A}'}, P^{G'}_{\emptyset}[T_{\mathcal{A}'}]\big)^2} < 2^8 \cdot (1 + T_f) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\big(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\big)^2}$$

□

The proofs for DDH to CDH (Example 3 and 6) and IND-CPA to OW-CPA (Example 2 and 7) are similar to that of PRG and OWF.

Theorem 5 (DDH to CDH)**.** *If DDH on $(\mathbb{G}, g)$ is λ-bit hard, then CDH on $(\mathbb{G}, g)$ is also $(\lambda - \alpha)$-bit hard, with respect to the Hellinger-bit hardness. Here, $\alpha = 4 + \log(1 + T_{eq})$, where T_{eq} denotes the cost for checking two group elements are the same.*

Proof. Let $\mathcal{A}$ be an adversary against the CDH game G, with a positive advantage. We construct an adversary $\mathcal{A}'$ against the DDH game G' as follows: Whenever a query is made, $\mathcal{A}'$ runs $\mathcal{A}$ on (g^x, g^y) of the sample (g^x, g^y, g^z) and checks whether $g^z = \mathcal{A}(g^x, g^y)$. The adversary records `Yes` if the condition is satisfied and `No` if not. Note that $\Pr[\texttt{Yes}|b = 0] = P^G_{\mathcal{A}}$ and $\Pr[\texttt{Yes}|b = 1] = 1/|\mathbb{G}| = P^G_{\emptyset}[T_{\mathcal{A}}]$ (Example 11).

Now the goal is to distinguish the two cases of $b = 0$ and $b = 1$ using this `Yes`/`No` distribution. By Prop. 2, to distinguish two cases with a probability of at least $1 - \delta$, it is sufficient to count the number of `Yes`'s among $N = \ln(\frac{1}{2\delta})/d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2$ samples. For such an adversary $\mathcal{A}'$, we have the following bound.

$$\begin{aligned}\frac{T_{\mathcal{A}'}}{d_{\mathrm{H}}\left(P^{G'}_{\mathcal{A}'}, P^{G'}_{\emptyset}[T_{\mathcal{A}'}]\right)^2} &< \frac{\ln(\frac{1}{2\delta})}{d_{\mathrm{H}}\left(1-\delta, \frac{1}{2}\right)^2} \cdot (1 + T_{eq}) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2} \\ &< 2^4 \cdot (1 + T_{eq}) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}\end{aligned}$$

□

Theorem 6 (IND-CPA to OW-CPA). *Consider a public-key encryption scheme with a message space $\mathcal{M}$. If it is λ-bit secure against the IND-CPA game, then it is also $(\lambda - \alpha)$-bit secure against the OW-CPA game, with respect to the Hellinger-bit security. Here, $\alpha = 4 + \log(1 + T_{\mathcal{M}})$, where $T_{\mathcal{M}}$ denotes the cost for sampling a random message and checking if two messages are the same.*

Proof. Let $\mathcal{A}$ be an adversary against the OW-CPA game G, with success probability $\varepsilon := P^G_{\mathcal{A}}$ such that $P^G_{\mathcal{A}} > P^G_{\emptyset}[T_{\mathcal{A}}] = 1/|\mathcal{M}|$ (Example 12). We construct an adversary $\mathcal{A}'$ against the IND-CPA game G' as follows: The adversary $\mathcal{A}'$ always makes a query with randomly chosen messages m_0 and m_1. Whenever a query is made, $\mathcal{A}'$ runs $\mathcal{A}$ on the sample c and checks whether the output m' satisfies $m' = m_0$. The adversary records `Yes` if the condition is satisfied and `No` if not. Note that $\Pr[\texttt{Yes}|b = 0] = \varepsilon$ and $\Pr[\texttt{Yes}|b = 1] = \frac{1-\varepsilon}{|\mathcal{M}|-1}$.

Now the goal is to distinguish the two cases of $b = 0$ and $b = 1$ using this `Yes`/`No` distribution. By Prop. 2, to distinguish two cases with a probability of at least $1 - \delta$, it is sufficient to count the number of `Yes`'s among $N = \ln(\frac{1}{2\delta})/d_{\mathrm{H}}(\varepsilon, \frac{1-\varepsilon}{|\mathcal{M}|-1})^2$ samples. For such an adversary $\mathcal{A}'$, we have the following bound.

$$\frac{T_{\mathcal{A}'}}{d_{\mathrm{H}}\left(P^{G'}_{\mathcal{A}'}, P^{G'}_{\emptyset}[T_{\mathcal{A}'}]\right)^2} < \frac{\ln(\frac{1}{2\delta})}{d_{\mathrm{H}}\left(1-\delta, \frac{1}{2}\right)^2} \cdot \frac{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}{d_{\mathrm{H}}\left(\varepsilon, \frac{1-\varepsilon}{|\mathcal{M}|-1}\right)^2} \cdot (1 + T_{\mathcal{M}}) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}$$

Meanwhile, we can bound the second term, since $P^G_{\mathcal{A}} = \varepsilon > \frac{1}{|\mathcal{M}|} = P^G_{\emptyset}[T_{\mathcal{A}}]$.

$$\frac{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}{d_{\mathrm{H}}\left(\varepsilon, \frac{1-\varepsilon}{|\mathcal{M}|-1}\right)^2} < 1$$

In total, if we choose $\delta \approx 0.1$, we have the following bound.

$$\frac{T_{\mathcal{A}'}}{d_{\mathrm{H}}\left(P^{G'}_{\mathcal{A}'}, P^{G'}_{\emptyset}[T'_{\mathcal{A}}]\right)^2} < 2^4 \cdot (1 + T_{\mathcal{M}}) \cdot \frac{T_{\mathcal{A}}}{d_{\mathrm{H}}\left(P^G_{\mathcal{A}}, P^G_{\emptyset}[T_{\mathcal{A}}]\right)^2}$$

□

5 Comparisons

In this section, we compare our new definitions with the definitions in previous works of Micciancio-Walter [MW18] and Watanabe-Yasunaga [WY21]. We indicate several weaknesses of the previous works and strengths of our work to suggest that our Hellinger-bit security is a more *right* definition.

5.1 Definition of Security Game

We first review the definitions of security games in the previous works.

Definition 14 ([MW18, Def. 5])**.** An n-bit *security game* is played by an adversary $\mathcal{A}$ interacting with a challenger. At the beginning of the game, the challenger chooses a secret $x \in \{0,1\}^n$, represented by the random variable X, from some distribution $\mathcal{D}_X$. At the end of the game, $\mathcal{A}$ outputs some value, which is represented by the random variable A. The adversary $\mathcal{A}$ wins the game if it outputs a such that $(x, a) \in R$, where R is some relation. $\mathcal{A}$ may output a special symbol $\perp$ such that $R(x, \perp)$ and $\bar{R}(x, \perp)$ are both false.

Definition 15 ([WY21, Inner Game[13]])**.** An n-bit *security game* G consisting of an algorithm X, a relation R, and an oracle O, is played by an adversary $\mathcal{A}$ given oracle access to O. At the beginning of the game, a secret $u \in \{0,1\}^n$ is chosen uniformly at random, and the challenge x is computed as $X(u)$. Given x, the adversary $\mathcal{A}$ wins the game if it outputs a value a such that $(u, x, a) \in R$.

Previous definitions assert security games to begin with the challenger choosing a *secret* and characterize winning conditions as relations between this *secret* and the final answer of the adversary (plus a challenge in [WY21]). We regard this as an unnatural and restrictive formulation since we often consider security games where the winning condition is affected by the queries of an adversary during the game. This includes the EUF-CMA game for signature schemes (Example 5). Thus, previous definitions fail to capture even the very basic EUF-CMA game.

On the other hand, our definition (Def. 3) captures essentially all security games: we do not put any restrictions on the structure of the games and characterize winning conditions regarding the *view* of the challenger. In particular, our framework captures all security games that are covered by the definitions of falsifiable assumptions of Gentry-Wichs [GW11] and thus complexity assumptions of Goldwasser-Kalai [GK16].

5.2 Definition of Search Game

In the previous works of [MW18, WY21], decision and search games are defined by the length of the secret chosen by the challenger. That is, decision games are defined as 1-bit games, and search games as n-bit games with a large n. We claim that this characterization is problematic. In particular, we construct the following pathological examples.

Example 18 (Pathological Examples)**.** Let G be a decision game, i.e., 1-bit security game, with respect to the definition of [MW18] or [WY21]. We can naturally extend G into a redundant n-bit security game G' with an arbitrarily large n as follows: The first bit of n-bit secret is chosen following the secret distribution of G, and the remaining bits are chosen uniform randomly. The game G' proceeds exactly the same as G regarding the first bit of the secret as the secret bit of G. The winning condition of G' is also defined as the same as G by only reading the first bit of the secret.

In the example, G and G' are essentially the same game. However, G is a decision game while G' is a search game, according to the definition of [MW18, WY21]. Thus, we can say that the definition does not capture the core nature of search games in a suitable way. This is especially problematic when the definition of bit security depends on whether the game is decision or search, as in [WY21]. In fact, in Section 5.5, we show that G and G' generally have different bit security under the definitions of [MW18] and [WY21].

On the other hand, we do not precisely define search games (Remark 3). We rather take the extremely small baseline probability as a fuzzy characterization of search games. Recall

[13] A followup work by Watanabe-Yasunaga [WY23] considers game where adversaries are allowed to output $\perp$, as in [MW18]. However, this change does not affect discussions in this section.

that our definition of bit security only depends on baseline probability and is independent of any other structures of games. Thus, such a characterization suffices to apply our definitions to identify the quadratic gap of bit security between decision and search games (Remark 10). Elsewhere, we do not have to distinguish search games from decision games.

5.3 Bit Security of Micciancio-Walter

We review the definition of Micciancio-Walter [MW18] and compare it with our definition.

Definition 16 ([MW18, Def. 7])**.** For any security game G (regarding Def. 14) and adversary $\mathcal{A}$ against G, we denote and define the *MW-advantage* of $\mathcal{A}$ against G as the following.

$$\mathrm{adv}^{G}_{\mathrm{MW}}(\mathcal{A}) = \frac{I(X;Y)}{H(X)}$$

Here, $I(\cdot;\cdot)$ is the mutual information, $H(\cdot)$ is the Shannon entropy, and $Y(X,A)$ is the random variable with marginal distributions $Y_{x,a} = \{Y|X = x, A = a\}$ defined as

1. $Y_{x,\perp} = \perp$, for all x.
2. $Y_{x,a} = x$ for all $(x,a) \in R$.
3. $Y_{x,a} = \{x' \leftarrow \mathcal{D}_X | x' \neq x\}$, for all $(x,a) \in \bar{R}$.

Definition 17 ([MW18, Def. 8])**.** For any security game G (regarding Def. 14), we denote and define its *MW-bit security* as the following.

$$BS_{\mathrm{MW}}(G) = \min_{\mathcal{A}} \log \left(\frac{T_{\mathcal{A}}}{\mathrm{adv}^{G}_{\mathrm{MW}}(\mathcal{A})} \right)$$

Our definition and the definition of [MW18] are both *unified*, in the sense that the bit security is defined in a single framework regardless of game types (e.g., decision/search). The difference between the two works lies in naturalness and interpretability. Micciancio-Walter introduces a hypothetical random variable Y in Def. 16, which lacks intuitive meaning without a sufficient explanation. In particular, there remains a question of why Y should be defined in such a way in Case 3 of Def. 16. Moreover, there may be controversies on why the aborts must be allowed, why they must not be regarded as failures, and why they must affect the bit security in such a specific way. On the other hand, our definition is based on a simple and natural concept of *demonstrating advantage* (Def. 10), which has a firm operational meaning by nature.

5.4 Bit Security of Watanabe-Yasunaga

We review the definition of Watanabe-Yasunaga [WY21] and compare it with our definition.

Definition 18 ([WY21, Outer Game])**.** Let $\mathcal{A}$ be an adversary against a security game G, (regarding Def. 15).

- When G is a decision game (i.e. 1-bit game), the *outer game* of G with respect to $\mathcal{A}$ is played by an *outer* adversary $\mathcal{B}$, who wins the game if it outputs $u \in \{0,1\}$ given oracle access to $\mathcal{A}(X(u))$. See Figure 2a.
- When G is a search game (i.e. n-bit game with $n > 1$), the *outer game* of G with respect to $\mathcal{A}$ is played by an *outer* adversary $\mathcal{B}$, who invokes G several times and wins if there was any game $\mathcal{A}$ won. In other words, $\mathcal{B}$ has oracle access to $\mathcal{A}(X(u))$, where u is uniformly chosen from $\{0,1\}^n$ at the beginning of each query. See Figure 2b.

The outer game of G with respect to $\mathcal{A}$ is denoted as $\hat{G}^{\mathrm{WY}}_{\mathcal{A}}$.

Definition 19 ([WY21, Def. 1])**.** For any security game G (regarding Def. 15), we denote and define its *WY-bit security* (with respect to error probability $0 < \delta < 1/2$) as the following, where $N_{\mathcal{B}}$ denotes the query complexity of $\mathcal{B}$.

$$BS^{\delta}_{\mathrm{WY}}(G) = \min_{\mathcal{A},\mathcal{B}} \left\{ \log(T_{\mathcal{A}} \cdot N_{\mathcal{B}}) : \Pr[\mathcal{B} \text{ wins } \hat{G}^{\mathrm{WY}}_{\mathcal{A}}] \geq 1 - \delta \right\}$$

Our work and [WY21] both define bit security as the cost to win certain *meta*-games, yielding firm operational meanings to the definitions. The difference between the two works lies in generality. Although [WY21] defines bit security for both search and decision primitives as the cost for winning meta-games, the designated games for search and decision primitives differ *qualitatively*. On the other hand, our definition is *unified*, in the sense that the bit security is defined in terms of a single meta-game, the advantage observation game (Def. 10), regardless of game types (e.g., decision/search).

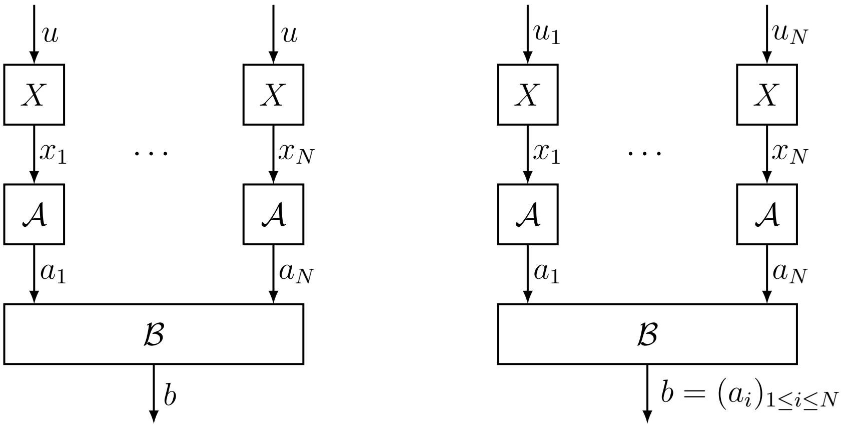

(a) Decision Game

(b) Search Game

Figure 2: Outer Games of [WY21]

5.5 Pathological Examples

It is easy to observe that the bit security is *finite* for n-bit security games with $n > 1$, under the definitions of [MW18, WY21]. This fact contradicts our expectation that unconditionally secure primitives (e.g., information-theoretic MAC) will have ∞-bit security. On the other hand, our definition reflects the expectation by its nature (Example 17).

Now consider a 1-bit game G, where an adversary wins the game if it correctly guesses the secret bit but is not allowed to make any queries. Indeed, the game G is unconditionally secure and achieves ∞-bit security under the definitions of [MW18, WY21]. However, the game G' (Example 18), which is essentially the same game but with a redundantly large secret, has *finite* (in fact, very small) bits of security under their definitions.

That is, some games that are essentially the same do not have the same number of bits of security under the previous definitions. We note that the game G is an extreme case considered for the ease of presentation, and this discrepancy happens in general for the games constructed in Example 18. On the other hand, our definition is not affected by such examples since our bit security depends only on success/baseline probabilities. These pathological examples suggest that the previous definitions of [MW18] and [WY21] are ill-defined.

Acknowledgments

We would like to thank Daniel J. Bernstein, Michael Walter, and the anonymous reviewers for their thoughtful comments and suggestions on earlier versions of this paper.

References

[ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 327–343. USENIX Association, August 2016.

[AGHP90] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k-wise independent random variables. In *31st FOCS*, pages 544–553. IEEE Computer Society Press, October 1990. doi:10.1109/FSCS.1990.89575.

[BDJR97] Mihir Bellare, Anand Desai, Eric Jokipii, and Phillip Rogaway. A concrete security treatment of symmetric encryption. In *38th FOCS*, pages 394–403. IEEE Computer Society Press, October 1997. doi:10.1109/SFCS.1997.646128.

[BH19] Daniel J. Bernstein and Andreas Hülsing. Decisional second-preimage resistance: When does SPR imply PRE? In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 33–62. Springer, Heidelberg, December 2019. doi:10.1007/978-3-030-34618-8_2.

[BL13] Daniel J. Bernstein and Tanja Lange. Non-uniform cracks in the concrete: The power of free precomputation. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 321–340. Springer, Heidelberg, December 2013. doi:10.1007/978-3-642-42045-0_17.

[BLL$^+$15] Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 3–24. Springer, Heidelberg, November / December 2015. doi:10.1007/978-3-662-48797-6_1.

[BR05] Mihir Bellare and Phillip Rogaway. Introduction to modern cryptography, 2005. URL: https://web.cs.ucdavis.edu/~rogaway/classes/227/spring05/book/main.pdf.

[BY02] Ziv Bar-Yossef. *The complexity of massive data set computations*. University of California, Berkeley, 2002.

[Can17] Clément Canonne. A short note on distinguishing discrete distributions, 2017. URL: https://github.com/ccanonne/probabilitydistributiontoolbox/blob/master/testing.pdf.

[GK16] Shafi Goldwasser and Yael Tauman Kalai. Cryptographic assumptions: A position paper. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 505–522. Springer, Heidelberg, January 2016. doi:10.1007/978-3-662-49096-9_21.

[GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32. ACM Press, May 1989. doi:10.1145/73007.73010.

[Gol04] Oded Goldreich. *The Foundations of Cryptography - Volume 2: Basic Applications*. Cambridge University Press, 2004. URL: http://www.wisdom.weizmann.ac.il/%7Eoded/foc-vol2.html, doi:10.1017/CBO9780511721656.

[GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011. doi:10.1145/1993636.1993651.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

[KL14] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014. URL: https://www.crcpress.com/Introduction-to-Modern-Cryptography-Second-Edition/Katz-Lindell/p/book/9781466570269.

[KM13] Neal Koblitz and Alfred Menezes. Another look at non-uniformity. *Groups Complex. Cryptol.*, 5(2):117–139, 2013. URL: https://doi.org/10.1515/gcc-2013-0008, doi:10.1515/GCC-2013-0008.

[MW17] Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 455–485. Springer, Heidelberg, August 2017. doi:10.1007/978-3-319-63715-0_16.

[MW18] Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 3–28. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78381-9_1.

[Nao03] Moni Naor. On cryptographic assumptions and challenges (invited talk). In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 96–109. Springer, Heidelberg, August 2003. doi:10.1007/978-3-540-45146-4_6.

[PDG14] Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In Lejla Batina and Matthew Robshaw, editors, *CHES 2014*, volume 8731 of *LNCS*, pages 353–370. Springer, Heidelberg, September 2014. doi:10.1007/978-3-662-44709-3_20.

[Pei16] Chris Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016. doi:10.1561/0400000074.

[PW] Yury Polyanskiy and Yihong Wu. Information theory. URL: https://people.lids.mit.edu/yp/homepage/data/itbook-export.pdf.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603.

[Ros21] Mike Rosulek. The joy of cryptography, 2021. URL: https://joyofcryptography.com.

[WY21] Shun Watanabe and Kenji Yasunaga. Bit security as computational cost for winning games with high probability. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 161–188. Springer, Heidelberg, December 2021. doi:10.1007/978-3-030-92078-4_6.

[WY23] Shun Watanabe and Kenji Yasunaga. Unified view for notions of bit security. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VI*, volume 14443 of *LNCS*, pages 361–389. Springer, Heidelberg, December 2023. doi:10.1007/978-981-99-8736-8_12.

A Proof of Prop. 2

In this section, we prove Prop. 2. It follows from Lemma 1 and 2. We first recall a standard fact in learning theory.

Proposition 3 (Simple Hypothesis Test [BY02, Prop. 2.58]). *For two discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ on the same domain $\mathcal{X}$, we say (possibly randomized) algorithm $\mathcal{A} : \mathcal{X} \to \{0,1\}$ is a simple hypothesis test for $\mathcal{D}_0, \mathcal{D}_1$ with error δ, if the following holds.*

$$\Pr\left[b \xleftarrow{\$} \{0,1\};\ x \leftarrow \mathcal{D}_b;\ \mathcal{A}(x) = b\right] = 1 - \delta$$

Let δ^* be the minimum error among all simple hypothesis tests for $\mathcal{D}_0, \mathcal{D}_1$. Then, we have the following equality.

$$d_{\mathrm{TV}}(\mathcal{D}_0, \mathcal{D}_1) = 1 - 2\delta^*$$

Lemma 1 (Sample Complexity Upper Bound). *For $0 < \delta < 1/2$, we have the following upper bound on the sample complexity of distinguishing discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ with probability at least $1 - \delta$.*

$$N_\delta(\mathcal{D}_0, \mathcal{D}_1) \ \leq \ \frac{\ln(\frac{1}{2\delta})}{d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2}$$

Proof. For convenience, let $\varepsilon := d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)$. Then, by Prop. 1(c), we have

$$d_{\mathrm{H}}(\mathcal{D}_0^N, \mathcal{D}_1^N)^2 = 1 - (1 - \varepsilon^2)^N \geq 1 - \exp(-N\varepsilon^2)$$

and therefore $d_{\mathrm{TV}}(\mathcal{D}_0^N, \mathcal{D}_1^N) \geq 1 - \exp(-N\varepsilon^2)$ holds by Prop. 1(d). If we choose $N \geq \ln(\frac{1}{2\delta})/\varepsilon^2$, we have $d_{\mathrm{TV}}(\mathcal{D}_0^N, \mathcal{D}_1^N) \geq 1 - 2\delta$. Then, Prop. 3 implies that we can distinguish $\mathcal{D}_0^N$ and $\mathcal{D}_1^N$ with probability $1 - \delta$ using a single sample. Equivalently, we can distinguish $\mathcal{D}_0$ and $\mathcal{D}_1$ with probability $1 - \delta$ using N samples. □

Lemma 2 (Sample Complexity Lower Bound). *For $0 < \delta < 1/2$, we have the following lower bound on the sample complexity of distinguishing discrete distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ with probability at least $1 - \delta$, when $d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2 \leq 1/2$.*

$$N_\delta(\mathcal{D}_0, \mathcal{D}_1) \ \geq \ \frac{1}{4\ln 2} \cdot \frac{\ln(\frac{1}{4\delta(1-\delta)})}{d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2}$$

Proof. Say we have an algorithm that distinguishes $\mathcal{D}_0$ and $\mathcal{D}_1$ with probability at least $1-\delta$ using N samples. Then, we can view it as an algorithm that distinguishes $\mathcal{D}_0^N$ and $\mathcal{D}_1^N$ with probability at least $1 - \delta$ using a single sample. By Prop. 3, $d_{\mathrm{TV}}(\mathcal{D}_0^N, \mathcal{D}_1^N) \geq 1 - 2\delta$ must hold. Applying this to Prop. 1(d), we have the following.

$$d_{\mathrm{H}}(\mathcal{D}_0^N, \mathcal{D}_1^N)^2 \geq 1 - \sqrt{1 - d_{\mathrm{TV}}(\mathcal{D}_0^N, \mathcal{D}_1^N)^2} \geq 1 - \sqrt{1 - (1 - 2\delta)^2} = 1 - \sqrt{4\delta(1 - \delta)}$$

Meanwhile, we can rewrite $d_{\mathrm{H}}(\mathcal{D}_0^N, \mathcal{D}_1^N)^2$ using Prop. 1(c). Then, denoting $\varepsilon := d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)$ when convenient, we have the followings. In the last inequality, we used the fact that $1 - x \geq 2^{-2x}$ for $0 \leq x \leq 1/2$, together with the assumption $\varepsilon^2 \leq 1/2$.

$$N \geq \frac{\ln(1/\sqrt{4\delta(1-\delta)})}{\ln(1/(1-\varepsilon^2))} = \frac{1}{2} \cdot \frac{\ln(\frac{1}{4\delta(1-\delta)})}{\ln(1/(1-\varepsilon^2))} \geq \frac{1}{4\ln 2} \cdot \frac{\ln(\frac{1}{4\delta(1-\delta)})}{d_{\mathrm{H}}(\mathcal{D}_0, \mathcal{D}_1)^2}$$

□