# Bit Security: optimal adversaries, equivalence results, and a toolbox for computational-statistical security analysis*

Daniele Micciancio Mark Schultz-Wu

University of California, San Diego
{daniele@cs.,mdshult@}ucsd.edu

**Abstract**

We investigate the notion of bit-security for decisional cryptographic properties, as originally proposed in (Micciancio & Walter, Eurocrypt 2018), and its main variants and extensions, with the goal clarifying the relation between different definitions, and facilitating their use. Specific contributions of this paper include: (1) identifying the optimal adversaries achieving the highest possible MW advantage, showing that they are deterministic and have a very simple threshold structure; (2) giving a simple proof that a competing definition proposed by (Watanabe & Yasunaga, Asiacrypt 2021) is actually equivalent to the original MW definition; and (3) developing tools for the use of the extended notion of computational-statistical bit-security introduced in (Li, Micciancio, Schultz & Sorrell, Crypto 2022), showing that it fully supports common cryptographic proof techniques like hybrid arguments and probability replacement theorems. On the technical side, our results are obtained by introducing a new notion of "fuzzy" distinguisher (which we prove equivalent to the "aborting" distinguishers of Micciancio and Walter), and a tight connection between the MW advantage and the *Le Cam* metric, a standard quantity used in statistics.

## 1 Introduction

The level of security provided by a cryptographic construction is customarily measured in "bits". The intuition is that breaking an application offering "$n$ bits of security" should have a cost[1] comparable to mounting a key recovery attack on an ideal cryptographic function with a key space of size $2^n$. Formalizing this intuition is not entirely trivial, because cryptographic attacks often exhibit a trade-off between the cost (e.g., the running time $T_A$) of the attack, and its success probability $\epsilon_A$. For (verifiable) search problems, like forging digital signatures, it is well established[2] that bit security can be defined as the quantity $\log_2(T_A/\epsilon_A)$, minimized over all possible adversaries $A$. However, the situation for decision problems (like indistinguishability of ciphertexts,



[1] Various measures of cost have been considered, and the reader is referred to [BL13, Appendix B] for a discussion. For simplicity, in this paper we identify the cost of an attack with its running time.

[2] This is justified by the fact that one can repeat the attack $O(1/\epsilon)$ times to make the success probability arbitrarily close to 1.

zero knowledge, pseudorandomness, etc.) is far less clear. We recall that in a decision game the goal of the adversary is to distinguish between two distributions $X_b$ for $b \in \{0,1\}$. So, a naive approach to measure security could be to mimic the definition for search problems, and replace the quantity $\log_2(T_A/\epsilon_A)$ with $\log_2(T_A/\delta_A)$, where $\delta_A = 2\epsilon_A - 1$ is the advantage (over a random choice) of guessing the bit $b$. But it is well known that this naive definition leads to paradoxical situations, where for example [DTT10] an algorithm $G$ is deemed more secure (i.e., it is attributed a higher level of bit security) as a pseudorandom generator than as a one-way function. This is at odds with cryptographic intuition because pseudorandomness is a stronger security requirement than one-wayness. (See [MW18] and references therein for a detailed discussion of this and other problematic examples.)

During the last few years, several papers have investigated the problem of giving meaningful definitions of bit security [MW18, WY21, WY23, LMSS22, Lee24], or using them to give a tight security analysis of cryptographic primitives (e.g., see [ALWW21, LMSS22]). A satisfactory definition of bit security for decision games was first proposed by Micciancio and Walter in [MW18]. A key element of their definition is to consider attackers that may output either a bit $b \in \{0,1\}$ (indicating a decision between $X_0$ and $X_1$) or a special "don't know" symbol $\perp$. Interestingly, [MW18] shows that this simple extension of traditional adversaries, together with an appropriate definition of advantage (already used by [Lev93] in a different context,) allows to resolve all the previously mentioned paradoxes, and argues (by means of examples) that this is the right definition of bit security.[3] Since then, a number of alternative definitions have appeared [WY21, WY23, LMSS22, Lee24], with various motivations. Watanabe and Yasunaga [WY21] proposed a competing framework to define bit security that directly admits what they call an "operational interpretation", and later argued [WY23] that it is actually equivalent to the original MW definition [MW18]. A seemingly attractive feature of their definition is that it only requires standard (non-aborting) adversaries with output in $\{0,1\}$. A variant of their definition that (similarly to [MW18]) interpolates between search and decision problems is given in [Lee24]. In a different and orthogonal direction, Li, Micciancio, Schultz and Sorrell [LMSS22] extend the MW definition to encompass both computational and statistical security. Informally, statistical security provides a strong measure of security even against computationally unbounded adversaries. When achievable, statistical security has the advantage of being easier to anaylze, and not requiring any computational assumptions. In practice, when setting parameters and optimizing efficiency, it is common to require lower levels of statistical bit security $s$, than computational bit security $c$. For example, $s = 80$ is usually considered more than acceptable, while computational security typically requires $c \geq 128$ or even higher values to anticipate possible improvements in the computational complexity of attacks. Li at al. [LMSS22] define $(c, s)$-security as satisfied by a protocol that provides *either* $c$ bits of computational security, *or* $s$ bits of statistical security against any possible attack. We remark that a protocol can admit both attacks with running time much less than $2^c$ (as long as their advantage is less than $2^{-s}$) and (different) attacks achieving advantage very close to 1 (as long as their running time is higher than $2^c$). In other words, a $(c, s)$-secure protocol can achieve neither $c$-bits of computational security nor $s$-bits of statistical security. Still, morally, it provides an acceptable level of security wherever

[3] This is at least for search (key recovery) and decision problems. The work [MW18] also proposes a more general definition based on information theory that interpolates between search and decision problems (e.g., encompassing password recovery problems with a polynomially large set of secrets,) but the corresponding notion of bit security for intermediate cases is largely unexplored. In this paper, we focus on the special case of decision problems which is the most relevant to cryptography. For search problems, the general bit-security definition of [MW18] reduces to $\log_2 T_A/\epsilon_A$, which is standard, and is adopted in this paper too.

$s$-bit statistical security and $c$-bit computational security are considered individually adequate. The advantage of $(c, s)$-security is that it allows to seamlessly combine statistical and computational cryptographic primitives (something very common in practice) and still be able to formally quantify the security level of an application. However, the notion of $(c, s)$-security has not been further explored, and, despite its potential usefulness, it has seen little adoption due to the lack of tools to simplify its usage.

**Our Contributions and Techniques** In this work, we examine the bit security definitions of [MW18, WY21, LMSS22], proving structural results about optimal (statistical) adversaries, clarifying the relation between the MW and WY bit security definitions, and then applying these results to the recent notion of $(c, s)$-bit security. Our main contributions, described in more details in the next subsections, can be summarized as follows:

- We characterize the MW adversaries achieving the optimal (statistical) bit-security advantage. Specifically, we show that these adversaries may be assumed to be deterministic (Corollary 1) and have a simple "threshold" structure (Theorems 3).
- We show (Theorem 4) that the WY notion of bit security is equivalent to the original MW bit security definition. In other words, the definition put forward in [WY21] is not a new security notion, but a different formulation of MW bit-security which, potentially, may be more convenient in some settings. We remark that a proof of this equivalence was already given in [WY23], but, as we are going to describe, that proof contained a gap. We clarify the relation between the two definitions by filling the gap and also giving a simpler proof of the equivalence.
- Despite the fact that the WY definition only uses traditional (non-aborting) adversaries, we show (Theorem 5) that the natural "maximum likelihood" distinguisher can offshoot the correct bit security level by a large margin. So, the advantages of using standard (non-aborting) adversaries in the characterization of bit security put forward in [WY21] are unclear.
- We show that common proof techniques widely used in the analysis of cryptographic protocols can be extended to work with the more general notion of computational-statistical security from [LMSS22]. Specifically, we show that $(c, s)$-security fully supports the use of hybrid arguments (Theorem 6) and probability substitution (Theorem 7).

On the technical side, many of our results rely on a new class of adversaries that further extends the MW (aborting) adversaries, and that may be of independent interest. Specifically, we make use of adversaries (for decision games) that may output not just $0, 1$ (representing a high confidence decision) or $\perp$ (representing no confidence), but an arbitrary value $\sigma \in [-1, 1]$, with the sign $\sigma/|\sigma| \in \{-1, 1\}$ representing the decision, and the magnitude $|\sigma| \in [0, 1]$ the confidence level that can vary continuously from 0 (no confidence) to 1 (perfect confidence). Interestingly, we show (Theorem 1) that these "fuzzy" adversaries still define precisely the same notion of bit security as the original MW "aborting" adversaries. The robustness of the notion of bit security with respect to such extensions further supports the use of [MW18] as the standard notion of bit security. Still, our equivalent definition using fuzzy adversaries with output in the continuous interval $[-1, 1]$ supports the use of analytical techniques, and it is useful to prove some of the results in this paper. We believe that the characterization of bit security in terms of these more general fuzzy adversaries may find other applications, and is of independent interest.

**Related Work** As mentioned, our work directly builds on the bit security frameworks of [MW18, WY21], so is directly related to these works. Our work is also tangentially related to the bit security framework of [Lee24], though this work mostly focuses on generalizing (a variant of) the framework of [WY21] to non-decision games, whereas we focus on decision games. Our work on the optimal adversary for the MW advantage is additionally related to the notion of (binary) hypothesis testing with an aborting option, see for example [GHVK11], though the measure optimized in that work does not appear to be related to the MW advantage. The similarity between our work and binary hypothesis testing with a rejection option is perhaps more obvious from [LJ20, Section 4], where (in a slightly different setting) optimality of threshold distinguishers was also highlighted. Still in relation to hypothesis testing, we discuss the implications of [PJL23] to the WY formulation of bit security.

**Paper organization** The rest of this paper is organized as follows. In the rest of this section we give a more detailed, still informal, description of our results and techniques. Section 2 defines the notation and preliminary results used in this paper. In Section 3 we formally define fuzzy adversaries, establish their equivalence (in both the computational and statistical setting) with the aborting adversaries of [MW18], and then use them to investigate the structure of the (statistical) MW adversaries achieving the optimal advantage. In Section 4, we explore the WY bit security definition and its equivalence with the MW bit security. Finally, in Section 5, we build our toolbox for the use of $(c, s)$-security in the analysis of cryptographic protocols, establishing the validity of hybrid arguments and probability replacement theorems. Section 6 concludes with some final remarks and open problems.

### 1.1 The Micciancio-Walter Advantage

Consider the problem of distinguishing between two distributions $\mathcal{X} = (X_0, X_1)$ over a set $\Omega$. (Everything applies more generally to the case of more complex decision games where an adversary interacts with one of two oracles.) Micciancio and Walter (following [Lev93]) define the advantage of an "aborting" adversary $A : \Omega \to \{0, 1, \perp\}$ as

$$\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) = \frac{(\beta_A - \bar{\beta}_A)^2}{\beta_A + \bar{\beta}_A}, \tag{1}$$

where $\beta_A = \Pr[A(x_b) = b]$ and $\bar{\beta}_A = \Pr[A(x_b) = 1 - b]$ are the probability that $A$ outputs the correct or incorrect bit, respectively, when $b \in \{0, 1\}$ is chosen at random and $x_b \leftarrow X_b$. For traditional (non-aborting) adversaries with output in $\{0, 1\}$, we have $\beta + \bar{\beta} = 1$, and it is well known (and quite intuitive) that the best advantage is achieved by an adversary $A(x)$ that on input a sample $x \in \Omega$, outputs the bit $b \in \{0, 1\}$ for which $\Pr[X_b = x]$ is highest. Moreover, the resulting optimal advantage equals precisely the square $\Delta_{\mathsf{SD}}(X_0, X_1)^2$ of the statistical distance between the two distributions. This allows to easily compute the bit security of $\mathcal{X}$ whenever the probability distributions are efficiently computable. This is a common scenario in cryptography, where, for example $X_0$ may be an ideal probability distribution used in the theoretical analysis of a cryptographic scheme (e.g., a discrete gaussian distribution in lattice-based cryptography) and $X_1$ is an approximate (more efficiently samplable) version of $X_0$ used when implementing the algorithm in practice (e.g. using floating point numbers). In fact, this was precisely the motivation in [MW17, MW18].

However, once the adversary is allowed to output $\perp$, it is no longer clear how to determine an optimal adversarial strategy, even when the probability distributions $X_0, X_1$ are efficiently computable. For example, while intuitively it is clear that the adversary $A(x)$ should output $\perp$ (and express low confidence) when the probabilities $p_0 = \Pr\{X_0 = x\}$ and $p_1 = \Pr\{X_1 = x\}$ are very close to each other, it is unclear how close is "very close" or even how to measure closeness, e.g., by $|p_0 - p_1|$, $p_0/p_1$, or some other function (of $x$) that depends on the global properties of $X_0$ and $X_1$. One of our main goals is to characterize the optimal aborting adversarial strategies, both to improve our understanding of the MW bit security definition, and offer a simple tool for the computation of the bit-security distance between specific distributions that may occur in practice.

To this end, we first show that one can equivalently phrase the study of aborting adversaries in terms of the class of *fuzzy adversaries* $\mathcal{A}_{\approx} := \{\tilde{A} \mid \tilde{A} : \Omega \to [-1, 1]\}$. These adversaries' output $y = A(x)$ represents not only a guess of which distribution they are given (via $y/|y| \in \{\pm 1\}$), but also a *confidence level* $|y| \in [0, 1]$. One then measures the advantage of fuzzy adversaries with a "continuous" analogue of (refeq:aborting-advantage) (Definition 6), which we write as $\mathsf{adv}^{\mathsf{MW},\approx}_{\mathcal{X}}(\tilde{A})$. We prove equivalence (Theorem 10) by giving efficient, advantage-preserving transformations between the two classes of adversaries. This shows that, when maximized over the set of all possible adversaries, $\mathsf{adv}^{\mathsf{MW}}(A)$ and $\mathsf{adv}^{\mathsf{MW},\approx}(A)$ are equivalent. Moreover, since the transformations between fuzzy and aborting adversaries used in our proofs also preserve the adversary's running time, they also establish the equivalence between the corresponding notions of *computational* (and, looking forward, *computational-statistical*) bit security.

We then prove a number of useful properties for aborting and fuzzy adversaries. For example, we show that the MW advantage is a convex function of randomized aborting adversary. As a simple corollary, we derive that the optimal advantage is always achieved by a deterministic aborting adversary (Corollary 1), obtained by fixing the randomness of the probabilistic adversary. This fact, while intuitively obvious[4] and often considered a *folk theorem*, is not generally true, and we give an explicit counterexample demonstrating how it can fail. (See Lemma 12.)

Next we dig deeper into the structure of the optimal fuzzy adversary when probabilities are efficiently computable (or adversaries are computationally unbounded.) We already established that optimal fuzzy adversaries may be assumed to always declare (for any given input) either *full* confidence or *no* confidence at all in their decision. Here we characterize when optimal fuzzy adversaries have full confidence, i.e., output $\pm 1$ instead of 0. Specifically, we show that the optimal adversary must have confidence 0 precisely when the quantity $|\log \Pr\{X_0 = x\}/\Pr\{X_1 = x\}|$ is below a certain threshold $\tau \in [0, \log 3]$, which is a simple function of the adversary's conditional success probability (Theorem 3).

### 1.2 Watanabe-Yasunaga Bit Security

We next investigate the optimal adversary for Watanabe-Yasunaga Bit Security. On the technical side, our work here is less novel, as information theorists had alrady identified [PJL23] a natural choice of adversaries that fit our purposes. Before discussing the precise results, we briefly provide some background. Watanabe-Yasunaga Bit security (as originally defined in [WY21]) is specified in terms of an "inner" adversary $A$ that on input a sample $x \leftarrow X_b$, outputs either 0 or 1. This adversary is run $n$ times $y_1 = A(x_1), \ldots, y_n = A(x_n)$ on independent samples $x_i \leftarrow X_b$ all chosen from the same unknown distribution. The number of samples $n$ is chosen large enough so that the

[4] We believe that this is the case because we tend to give convexity for granted.

value of the bit $b$ can be determined with very high probability $\mu \approx 1$ (say, $\mu \geq 0.99$) based on the output values $y_1, \ldots, y_n$. So, the total running time is given by $n \cdot T_A$, and [WY21] defines the bit security to be $\log_2(n \cdot T_A)$, minimized over all inner adversaries $A$ and number of repetitions $n$ such that $\mu \geq 0.99$. They also show that this quantity can be equivalently estimated as $\log(T_A/\mathcal{R}_{1/2}(A_0, A_1))$ where $\mathcal{R}_{1/2}$ is the Renyi divergence of order 1/2, and $A_b = A(X_b) \in \{0, 1\}$ is the Bernoulli random variable defined by the output of the adversary on input a sample from $X_b$.

At this point, it is natural to ask:

- What is the relation between the MW and WY bit security?
- What is the optimal adversary $A(x) \in \{0, 1\}$ for the WY definition?

Notice that since the WY adversaries always output either 0 or 1, they are potentially easier to use, as the attacker does not have to choose whether or not to abort.

Regarding the first question, [WY21] proves only the inequality[5] $MW \leq WY$, showing that WY is a more conservative notion of bit security, and leaving a more precise comparison as an open problem. In a follow-up paper [WY23] the same authors claimed the equivalence between MW and WY (up to an additive constant), but with a catch. Technically, they prove the equivalence between MW and WY bit security for the same class of *aborting* adversaries (with output in $\{0, 1, \perp\}$) introduced in [MW18]. Then, they claim equivalence with the original WY definition by informally stating that the definition in [WY21] does not explicitly depend on the size of the co-domain[6] of the adversary $A$. However, the justification is incorrect because the Renyi divergence $\mathcal{R}_{1/2}(A(X_0), A(X_1))$ implicitly depends on the size of the co-domain of $A$. Despite this gap in the proof, we show that the main claim of [WY23] (about the equivalence of MW and WY bit security) is correct, and in the process we give a simpler and tighter proof of this fact. (Theorem 4.)

So, at this point, the WY notion of bit security can be considered an alternative characterization of the MW bit security, rather than a new definition, and the question is whether this alternative characterization can help in evaluating the bit security of decision problems. One seemingly attractive feature of the WY is that it uses standard adversaries $A(x)$ which always output $0, 1$. This is because for these adversaries there is a particularly natural attack, that on input a sample $x$ outputs the bit $b$ for which the probability $\Pr[X_b = x]$ is highest. However, this does not seem to help in evaluating the bit security using the WY characterization: we show (Theorem 5) that there exist distinguishing games where this natural adversarial strategy yields bit security estimates that are far from optimal.[7]

### 1.3 Computational/Statistical Bit Security

Finally, we investigate the definition of $(c, s)$-bit security proposed in [LMSS22], to extend MW bit security to encompass both computational and statistical security. Recall that the MW (computational) bit security of a problem is the largest $c$ such that $T(A)/\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) \geq 2^c$ for all adversaries $A$. Similarly, statistical security can be defined as the largest $s$ such that $1/\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) \geq 2^s$ for

[5] Here and elsewhere we use $MW$ and $WY$ as a shorthand for the number of bits of security as computed according to the respective definitions.

[6] By *co-domain* we mean the set of possible outputs of the adversary, i.e., $|\{0, 1\}| = 2$ for traditional distinguishers and $|\{0, 1, \perp\}|$ for aborting adversaries

[7] Specifically, the estimates are twice as high as the optimal, correct value. Recall that bit security (roughly) measures the *exponent* of the running time of the adversary. So, a multiplicative factor in bit security estimation is quite large.

all adversaries $A$, where this time the running time of $A$ is ignored. Li et al. [LMSS22] define a protocol to be $(c, s)$-secure if for any adversary $A$

$$\text{either} \quad \frac{T(A)}{\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)} \geq 2^c \qquad \text{or} \quad \frac{1}{\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)} \geq 2^s.$$

As explained in the introduction, a protocol satisfying this definition seems to provide an adequate level of security whenever computational security and statistical security are considered individually acceptable. Here we point out that a protocol can offer neither $c$ bits of computational security nor $s$ bits of statistical security, and still be $(c, s)$-secure. Consider for example a protocol such that there exist a very efficient adversary $A_c$ with running time $T(A_c) = 1$ that achieves MW advantage $2^{-s}$, and some other adversary $A_s$ with very large running time $T(A_s) \geq 2^c$ that achieves MW advantage $\approx 1$. Then, the protocol is neither computationally nor statistically secure because $A_c$ breaks computational security (for $s < c$), and $A_s$ breaks statistical security. So, $(c, s)$-security is strictly weaker than both $c$-bits computational security, and $s$-bits of statistical security. In fact, one should expect this to be the case in any application that makes use of both computational and statistical security primitives, as an adversary can choose to attack the application by trying to break either one or the other type of primitives.

While $(c, s)$-bit security was introduced in [LMSS22] (and successfully used to analyze a practical protocol), this was done via direct manipulation of the definition. In this paper we establish a tight connection between the MW advantage $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)$ and a standard distance measure used in statistics: the (squared) *Le Cam distance* $\Delta^2_{\mathsf{LC}}(A(X_0), A(X_1))$ between the adversary's output probability distributions. Then, we use this connection to prove several useful properties of the $(c, s)$-bit security which support two of the most common proof techniques in cryptography:

- The "hybrid argument" (see Theorem 6 for the formal statement): consider a sequence of distributions $X_0, \ldots, X_k$. If the game defined by any pair of neighboring distributions $(X_{i-1}, X_i)$ is $(c, s)$-secure, then the game defined by the extremal distributions $(X_0, X_k)$ is also $(c', s')$-secure, for $c' \approx c - \log k$ and $s' \approx s - \log k$.

- The "distribution replacement" theorem (see Theorem 7 for the formal statement): Consider a decision game $(X_0^Y, X_1^Y)$ parameterized by a distribution $Y$. If distinguishing between $(Y, Y')$ is $(c, s)$-secure, and $(X_0^Y, X_1^Y)$ is $(c, s)$-secure, then $(X_0^{Y'}, X_1^{Y'})$ is also $(c', s')$-secure, for $c' \approx c$ and $s' \approx s$.

Our results improve or extend previous work. For example, [MW18] had already proved a hybrid theorem for computational bit security, and hybrid theorems for statistical bit security are essentially a form of (pythagorean) triangle inequality for the associated distance functions between distributions. The novelty here is to establish the validity of hybrid arguments in the more general computational-statistical setting, where each pair of neighboring distributions $(X_{i-1}, X_i)$ may be neither computationally nor statistically indistinguishable. Distribution replacement theorems for bit security were previously proved in [MW18, Yas21], but only for the setting where $(X_0^Y, X_1^Y)$ are computationally close and $(Y, Y')$ are *statistically* close (either in the max-log or Hellinger distance.) Our theorem allows both $(X_0^Y, X_1^Y)$ and $(Y, Y')$ to be close in the much weaker sense of computational-statistical bit security.

Both types of techniques are cornerstones for the modular analysis of complex cryptographic protocols that combine several cryptographic primitives. Our results support the uniform use of

computational-statistical bit-security to analyze both the final application and its building blocks, including neighboring hybrids $(X_{i-1}, X_i)$ and probability replacements $(Y, Y')$. Moreover, they support the seamless combination of computational and statistical security primitives, while at the same time offering tight security estimates, which, before our work, could only be done either informally or using ad-hoc methods. The connection with the Le Cam metric, which underlies our proofs, is also of independent interest, and may find other applications.

# 2 Preliminaries

We will make use of the following variant of the Cauchy-Schwarz inequality.

**Lemma 1** (Bergström's Inequality)**.** *For any real numbers $a_1, \ldots, a_n$, and positive reals $b_1, \ldots, b_n$, we have that*

$$\frac{(\sum_{i\in[n]} a_i)^2}{\sum_{i\in[n]} b_i} \leq \sum_{i\in[n]} \frac{a_i^2}{b_i}.$$

*Proof.* Rearrange the Cauchy-Schwarz inequality to $\frac{\langle c,d\rangle^2}{\|c\|_2^2} \leq \|d\|_2^2$ and let $c_i = \sqrt{b_i}$, $d_i = a_i/\sqrt{b_i}$. □

## 2.1 Distances between Distributions

We use several similarity measures between (discrete) probability distributions $X_0, X_1$. Below we write $X_b(x)$ as a shorthand for $\Pr\{X_b = x\}$.

- Statistical Distance: $\Delta_{\mathsf{SD}}(X_0, X_1) = \frac{1}{2}\sum_x |X_0[x] - X_1[x]|$,
- (Squared) Hellinger Distance: $\Delta^2_{\mathsf{H}}(X_0, X_1) = \frac{1}{2}\sum_x (\sqrt{X_0[x]} - \sqrt{X_1[x]})^2$
- (Squared) Le Cam Distance: $\Delta^2_{\mathsf{LC}}(X_0, X_1) = \frac{1}{2}\sum_x \frac{(X_0[x]-X_1[x])^2}{X_0[x]+X_1[x]}$
- Renyi Divergence of order 1/2: $\Delta_{1/2}(X_0, X_1) = -2\ln \sum_x \sqrt{X_0[x]X_1[x]}$

It is well known that $\Delta_{\mathsf{SD}}$, $\Delta_{\mathsf{H}}$ and $\Delta_{\mathsf{LC}}$ are distance functions, i.e., they satisfy the triangle inequality. They are also closely related as follows.

**Lemma 2** ([PW22, Section 7].)**.** *For any two distributions $X_0, X_1$ we have*

$$\Delta^2_{\mathsf{H}}(X_0, X_1) \leq \Delta_{\mathsf{SD}}(X_0, X_1) \leq \sqrt{2}\Delta_{\mathsf{H}}(X_0, X_1)$$
$$\Delta_{\mathsf{H}}(X_0, X_1) \leq \Delta_{\mathsf{LC}}(X_0, X_1) \leq \sqrt{2}\Delta_{\mathsf{H}}(X_0, X_1).$$

In other words, $\Delta_{\mathsf{H}}$ and $\Delta_{\mathsf{LC}}$ are equivalent (up to a constant factor), while $\Delta_{\mathsf{SD}}$ is polynomially related to them. As for the divergence $\Delta_{1/2}$, it easily follows from the definitions that it can be expressed as a monotonically increasing function of the Hellinger distance:

$$\Delta_{1/2}(X_0, X_1) = 2\ln \frac{1}{1 - \Delta^2_{\mathsf{H}}(X_0, X_1)}.$$

**Lemma 3.** *For any two distributions $X_0, X_1$ such that $\Delta_{1/2}(X_0, X_1) < \infty$, we have*

$$\Delta^2_{\mathsf{H}}(X_0, X_1) \leq \frac{1}{2}\Delta_{1/2}(X_0, X_1) \leq \frac{\Delta^2_{\mathsf{H}}(X_0, X_1)}{1 - \Delta^2_{\mathsf{H}}(X_0, X_1)}.$$

*In particular, if $\Delta^2_{\mathsf{H}}(X_0, X_1) \leq 1/2$, then $\Delta^2_{\mathsf{H}}(X_0, X_1) \leq \frac{1}{2}\Delta_{1/2}(X_0, X_1) \leq 2\Delta^2_{\mathsf{H}}(X_0, X_1)$.*

*Proof.* Easily follows from the bounds $1 - (1/t) \leq \ln t \leq t - 1$ and relation $\Delta_{1/2}(X_0, X_1) = -2\ln(1 - \Delta^2_{\mathsf{H}}(X_0, X_1))$. See [WY23] for details. □

## 2.2 Cryptographic Games

Cryptographic games are defined by one or more randomized, stateful programs $\mathfrak{Q}$ used by an adversary $A$ to carry out an attack $A^{\mathfrak{Q}}$. When running $A^{\mathfrak{Q}}$, the adversary only has black-box access to $\mathfrak{Q}$, which is used as an oracle. There are two main categories of cryptographic games. In a *search* game, the final output of $A^{\mathfrak{Q}}$ is determined by $\mathfrak{Q}$ and indicates if the attack was successful. A *decision* game $\mathfrak{Q} = (\mathfrak{Q}_0, \mathfrak{Q}_1)$ is a pair of oracles with identical interfaces, so that an adversary $A$ may interact with either of them $A^{\mathfrak{Q}_0}$, $A^{\mathfrak{Q}_1}$. This time it is $A$ that produces an output at the end of the interaction. We refer to the set of all possible outputs of $A$ as the *co-domain* of the adversary, and classify adversaries based on their co-domain. We consider three main classes of adversaries: traditional adversaries $A \in \mathcal{A}_{0,1}$ with co-domain $\{0, 1\}$, *aborting* adversaries $A \in \mathcal{A}_{\perp}$ with co-domain $\{0, 1, \perp\}$, and *fuzzy* adversaries $A \in \mathcal{A}_{\approx}$ with co-domain $[-1, 1]$. The goal of the adversary is to determine if it is interacting with either $\mathfrak{Q}_0$ or $\mathfrak{Q}_1$. $A$'s advantage in distinguishing between $\mathfrak{Q}_0$ and $\mathfrak{Q}_1$ is defined later on. We write $A(\mathfrak{Q}_b)$ for the random variable describing the final output of $A$ at the end of the interaction, and $A(\mathfrak{Q})$ as an abbreviation for the pair of output distributions $(A(\mathfrak{Q}_0), A(\mathfrak{Q}_1))$ over the co-domain of $A$. We remark that the output distribution $A(\mathfrak{Q}_b)$ is defined over the internal randomness of both $A$ and $\mathfrak{Q}_b$. We write $A(\mathfrak{Q}_b; r)$ when we want to make the randomness of $A$ explicit.

In the simplest, prototypical example $\mathfrak{Q} = (\mathfrak{Q}_0, \mathfrak{Q}_1)$ is a pair of probability distributions over a common set $\Omega$. The only interaction between $A$ and $\mathfrak{Q}_b$ is to receive a single sample $x \leftarrow \mathfrak{Q}_b$ from $\Omega$, upon which $A(x)$ produces its final output. In other words, the adversary $A$ is just a (possibly randomized) algorithm with input in $\Omega$. For simplicity, the reader may keep this simple case in mind throughout the paper, instead of arbitraty games. In any case, once a game $\mathfrak{Q}$ and adversary $A$ have been chosen, the output $A(\mathfrak{Q}) = (A(\mathfrak{Q}_0), A(\mathfrak{Q}_1))$ is always a pair of probability distributions.

Cryptographic protocols can be parameterized by other cryptographic primitives or distributions used as building blocks. So, for example, we may write $P^Y$ for a cryptographic program that uses a probability distribution $Y$, and $P^{Y'}$ for the same program run with a different distribution $Y'$. Similarly, security games $(\mathfrak{Q}_0^Y, \mathfrak{Q}_1^Y)$ can be parameterized by $Y$.

We remark that the running time of an adversary $A$ against a game $\mathfrak{Q}$ does not include the time required to run $\mathfrak{Q}$ in the interaction $A(\mathfrak{Q})$. In other words, we only account for the time taken by $A$ to write its oracle queries and read the answers. We consider adversaries running in strict (e.g. polynomial) time, i.e., we assume that the running time of $A$ in a run $A(\mathfrak{Q}_b)$ does not depend on how $\mathfrak{Q}_b$ answers the oracle queries. In particular, $A$ has the same running time in $A(\mathfrak{Q}_0)$ and $A(\mathfrak{Q}_1)$. The running time of an adversary $A$ is denoted by $T_A$ or $T(A)$.

In some settings it is useful to define also a notion of running time for the game $\mathfrak{Q}$. However, it should be clear that the (total) running time of $\mathfrak{Q}$ in an execution $A(\mathfrak{Q})$ typically depends on the

adversary $A$.[8] The time taken to run $\eth$ in an execution $A(\eth)$ is denoted $T_{\eth}^{A}$. Then, we can define the running time of a game $\eth$ relative to the running time of $A$ as follows.

**Definition 1.** *The (relative) running time of $\eth$ is defined as the maximum*

$$T_{\eth} = \sup_{A} \frac{T_{\eth}^{A}}{T_A}$$

*over all possible adversaries $A$.*

For decision games $(\eth_0, \eth_1)$ we always assume that $\eth_0$ and $\eth_1$ have the same running time. Using this definition, the total running time to run $A(\eth)$ (including both the time for $A$ and for $\eth$) can be bounded as

$$T_{A(\eth)} = T_A + T_{\eth}^{A} \leq T_A \cdot (1 + T_{\eth}) \,.$$

In the asymptotic setting both the game $\eth_\kappa = (\eth_{\kappa,0}, \eth_{\kappa,1})$ and adversary $A_\kappa$ are parametrized by a security parameter $\kappa$, and all quantities (e.g., $A_\kappa$'s advantage in winning a decision game, its running time $T_{A_\kappa}$, etc.) become functions of $\kappa$. Notice that $A_\kappa$ may depend arbitrarily on $\kappa$, i.e., we consider non-uniform adversaries. Technically, $A(\kappa, \tau_\kappa)$ is an algorithm that takes as input the security parameter $\kappa$ and an advice string $\tau_\kappa$ that depends on the security parameter. However, we will make only very limited use of non-uniformity: in most of our results $\tau_\kappa$ is a very short (typically constant size, independently of $\kappa$) string. So, the non-uniformity can be easily eliminated by running $A(\kappa, \tau)$ on all possible value of $\tau$, estimating $A$'s advantage, and then picking the best value of $\tau$ to carry out the attack.

### 2.3 Bit Security

Consider an adversary $A$ against a decision game $\eth = (\eth_0, \eth_1)$, where $A(\eth_b)$ may output $0, 1$ or some other values. Throughout the paper we will use the following definitions and notation:

$$\begin{array}{rrcl}
\text{(success probability)} & \beta_A & = & \Pr[A(\eth_b) = b] \\
\text{(failure probability)} & \bar{\beta}_A & = & \Pr[A(\eth_b) = 1 - b] \\
\text{(output probability)} & \alpha_A & = & \beta_A + \bar{\beta}_A \\
\text{(distinguishing gap)} & \delta_A & = & \beta_A - \bar{\beta}_A
\end{array}$$

where all probabilities are computed over the random choice of $b \leftarrow \{0,1\}$, and the randomness of $\eth_b$ and $A$. Notice that $\alpha_A$ equals the probability that the output of $A$ is in $\{0,1\}$. So, for standard adversaries $A \in \mathcal{A}_{0,1}$ that always output a bit $A(\eth_b) \in \{0,1\}$, we have $\alpha_A = 1$ and $\delta_A = 2\beta_A - 1 = \Pr\{A(\eth_1) = 1\} - \Pr\{A(\eth_0) = 1\}$. But we will use the definition of $\beta_A, \bar{\beta}_A, \alpha_A$ and $\delta_A$ also for unrestricted adversaries that may output values outside of $\{0,1\}$. It is well-known that, in the case of probability distributions $\mathcal{X} = (X_0, X_1)$, the highest possible distinguishing gap equals the statistical distance $\Delta_{\mathsf{SD}}(X_0, X_1) = \max_{A \in \mathcal{A}_{0,1}} \delta_A$ and it is achieved by a very simple adversary

$$A_{\mathsf{SD}}^{\mathcal{X}}(x) = \begin{cases} 0 & \text{if } \Pr[X_0 = x] > \Pr[X_1 = x] \\ 1 & \text{if } \Pr[X_0 = x] < \Pr[X_1 = x] \end{cases} \tag{2}$$

[8] This is most obvious when $\eth$ is a game where $A$ may issue an arbitrary number of calls to the game oracles.

(When $\Pr[X_0 = x] = \Pr[X_1 = x]$, the output of $A$ can be chosen arbitrarily without affecting the gap $\delta$.) This is easily generalized to arbitrary decision games $\eth = (\eth_0, \eth_1)$, where

$$\Delta_{\mathsf{SD}}(\eth_0, \eth_1) = \max_{A \in \mathcal{A}_{0,1}} \delta_A. \tag{3}$$

Since $\delta_A = \beta_A - \bar{\beta}_A$ is the difference between two probabilities, and the maximum over all $A$ is non-negative, we always have $\Delta_{\mathsf{SD}}(\eth_0, \eth_1) \in [0, 1]$.

### 2.3.1 The MW Bit Security Measure

Micciancio and Walter [MW18] suggested to use a more general class of adversaries $\mathcal{A}_\perp$, which output either 0, 1, or a special "don't know" symbol $\perp$, and demonstrated that these adversaries, together with an appropriate notion of advantage, allow to resolve several theoretical paradoxes related to the definition of a cryptographically meaningful notion of "bit security". (The reader is referred to [MW18] for intuition and justification of this definition.)

**Definition 2** (MW Advantage)**.** *For any (possibly randomized) MW distinguisher $A \in \mathcal{A}_\perp$ and decision game $\eth = (\eth_0, \eth_1)$, the advantage of $A$ is*[9]

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \frac{\delta_A^2}{\alpha_A} = \frac{(\beta_A - \bar{\beta}_A)^2}{\beta_A + \bar{\beta}_A}$$

*The (squared) MW distance between two distributions is*

$$\Delta^2_{\mathsf{MW}}(\eth_0, \eth_1) = \sup_{A \in \mathcal{A}_\perp} \mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \in [0, 1]. \tag{4}$$

If we restrict our attention to "non-aborting" adversaries $A \in \mathcal{A}_{0,1}$, we have $\alpha_A = 1$, and $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \delta_A^2$ is the square of the distinguishing gap. This immediately gives the following inequality.

**Lemma 4.** *For any decision game $\eth = (\eth_0, \eth_1)$, we have*

$$\Delta_{\mathsf{SD}}(\eth_0, \eth_1) \le \Delta_{\mathsf{MW}}(\eth_0, \eth_1).$$

*Proof.* Using (3) and the definition of $\Delta_{\mathsf{MW}}$, we get

$$\Delta_{\mathsf{SD}}(\eth_0, \eth_1) = \sup_{A \in \mathcal{A}_{0,1}} \delta_A = \sup_{A \in \mathcal{A}_{0,1}} \sqrt{\mathsf{adv}^{\mathsf{MW}}_{\eth}(A)} \le \Delta_{\mathsf{MW}}(\eth_0, \eth_1)$$

where the inequality follows from taking the supremum over a larger set $A \in \mathcal{A}_\perp$. □

It is also easy to see that the MW distance satisfies the data processing inequality.

**Lemma 5** (Data-Processing Inequality)**.** *For any decition game $\eth = (\eth_0, \eth_1)$ and game transformation*[10] *$\Gamma$, we have that*

$$\Delta_{\mathsf{MW}}(\Gamma(\eth_0), \Gamma(\eth_1)) \le \Delta_{\mathsf{MW}}(\eth_0, \eth_1).$$

[9] This is syntactically different, but perfectly equivalent to the definition given in [MW18], which defines the advantage as $\alpha_A \cdot (2\beta^*_A - 1)^2$, where $\beta^*_A = \beta_A / \alpha_A$.

[10] A transformation is simply a game $\Gamma(\eth_b)$ with oracle access to $\eth_b$. Applying $\Gamma$ to a game $\eth$ defines a new game $(\Gamma(\eth_0), \Gamma(\eth_1))$.

*Proof.* For any aborting adversary $A$, define $A^{\Gamma}(\eth_b) := A(\Gamma(\eth_b))$. It is straightforward to see that

$$\Delta^2_{\mathsf{MW}}(\Gamma(\eth_0), \Gamma(\eth_1)) = \sup_{A^{\Gamma}} \mathsf{adv}^{\mathsf{MW}}_{\eth}(A^{\Gamma}) \leq \sup_{A} \mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \Delta^2_{\mathsf{MW}}(\eth_0, \eth_1).$$

□

We will use the following construction from [MW18] to transform an aborting adversary $A \in \mathcal{A}_{\perp}$ to one with only two possible output values. For any adversary $A \in \mathcal{A}_{\perp}$, decision game $\eth = (\eth_0, \eth_1)$, and value $z \in \{0, 1, \perp\}$, let $\hat{b} = A^{A(\eth)}_{\mathsf{SD}}(z) \in \{0, 1\}$ be the bit such that $\Pr\{A(\eth_{\hat{b}}) = z\}$ is highest. (This bit can be given as a non-uniform advice, or estimated probabilistically by repeatedly running $A(\eth_{\hat{b}})$ for $\hat{b} \in \{0, 1\}$ with independent randomness.) Given $\hat{b}$, let

$$A^z(\eth_b) = \text{if } (A(\eth_b) = z) \text{ then } \hat{b} \text{ else } \perp \tag{5}$$

be the modified adversary that first runs $z' \leftarrow A(\eth_b)$, and then outputs $\hat{b}$ if $z' = z$ or $\perp$ otherwise. Notice that $A^z$ differs from $A$ just by a relabeling of its output. So, it has the same running time $T(A^z) = T(A)$.

**Lemma 6** ([MW18, Lemma 1])**.** *For any pair decision game $\eth = (\eth_0, \eth_1)$, aborting adversary $A \in \mathcal{A}_{\perp}$, and value $z \in \{0, 1, \perp\}$, the modified adversary $A^z$ in (5) has advantage*

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A^z) = \frac{(\Pr\{A(\eth_0) = z\} - \Pr\{A(\eth_1) = z\})^2}{2(\Pr\{A(\eth_0) = z\} + \Pr\{A(\eth_1) = z\})}.$$

### 2.3.2 The WY Bit Security Measure

In [WY21], an alternative bit security measure was introduced. The definition is parameterized by a "high enough" probability threshold $\mu \approx 1$, but it can be shown that the precise value of $\mu$ has only a marginal impact on the definition. An equivalent quantity (without the parameter $\mu$) is also defined in terms of the Renyi divergence of order 1/2.

**Definition 3.** *Let $\eth = (\eth_0, \eth_1)$ be a decision game, $\mu \in [0, 1]$, and $\epsilon_{A,B_k} := \Pr_b[B_k(A(\eth_b)^k) = b]$, where $A \in \mathcal{A}_{0,1}$, $k \in \mathbb{N}$, $B_k : \{0,1\}^k \to \{0,1\}$, and $X^k$ is the product distribution over $\{0,1\}^k$ of $k$ independent copies of $X = A(\eth_b)$. Define*

$$\mathsf{WY}^{\mu}_{\eth}(A) = \min_k \min_{B_k} \{\log_2(k \cdot T_A) \mid \epsilon_{A,B_k} \geq 1 - \mu\}, \qquad \mathsf{WY}^{\mu}_{\eth} := \min_{A \in \mathcal{A}_{0,1}} \mathsf{WY}^{\mu}_{\eth}(A). \tag{6}$$

$$\mathsf{WY}_{\eth}(A) := \log_2 T(A) + \log_2 \left\lceil \frac{1}{\Delta_{1/2}(A(\eth_0), A(\eth_1))} \right\rceil, \qquad \mathsf{WY}_{\eth} := \min_{A \in \mathcal{A}_{0,1}} \mathsf{WY}_{\eth}(A). \tag{7}$$

We say that two bit security measures are equivalent if they differ by an additive constant factor. While not highlighted as a formal statement, [WY21] shows that all these measures are essentially equivalent.

**Lemma 7** ([WY21, implicit])**.** *For any distinguishing game $\eth := (\eth_0, \eth_1)$, for any constants $\mu \leq \mu'$, one has that*

$$\left|\mathsf{WY}^{\mu}_{\eth} - \mathsf{WY}^{\mu'}_{\eth}\right| \leq O(1). \tag{8}$$

$$\left|\mathsf{WY}_{\eth} - \mathsf{WY}^{\mu}_{\eth}\right| \leq O(1), \tag{9}$$

*Proof.* The (stronger) bound

$$\forall A \in \mathcal{A}_{0,1} : \left|\mathsf{WY}^{\mu}_{\eth}(A) - \mathsf{WY}^{\mu'}_{\eth}(A)\right| \leq \ln(\ln(\frac{1}{4\mu^2})) \leq O(1) \tag{10}$$

follows from simple algebraic manipulations of [WY21, Lemmas 4 and 6], which bound the minimum $k$ in (6) via

$$\frac{\ln(\frac{1}{4\mu})}{\Delta_{1/2}(A(\eth_0), A(\eth_1))} \leq k \leq \left\lceil \frac{\ln(\frac{1}{4\mu^2})}{\Delta_{1/2}(A(\eth_0), A(\eth_1))} \right\rceil. \tag{11}$$

Multiplying by $T_A$ and taking logarithms yields nearly matching upper and lower bounds on $\mathsf{WY}^{\mu}_{\eth}$, which suffice to establish (10). One then gets the claimed result by minimizing (10) over $A$. □

Equivalence with the MW bit security is proved in [WY23], but technically only for aborting adversaries, which we denote $\mathsf{WY}^{\perp}_{\eth} = \min_{A \in \mathcal{A}_{\perp}} \mathsf{WY}_{\eth}(A)$.

**Lemma 8** ([WY23, Theorems 1 and 2])**.** *For any distinguishing game* $\eth := (\eth_0, \eth_1)$*,*

$$\left|\mathsf{WY}^{\perp}_{\eth} - \mathsf{MW}_{\eth}\right| \leq O(1).$$

Note that the measure $\mathsf{WY}^{\perp}_{\eth}$ is not *a priori* equal to $\mathsf{WY}_{\eth}$, as minimizing over a larger set $\mathcal{A}_{\perp}$ may produce smaller values. So, Lemma 8 does not imply that $\mathsf{WY}_{\eth}$ and $\mathsf{MW}_{\eth}$ are equivalent. Still, this is true, as we will show in Section 4.

### 2.3.3 Computational/Statistical Bit security

Sometimes, in cryptography, one can achieve a strong notion of security, where no adversary can break a cryptographic function with high probability, regardless of the computational cost incurred by the attack. In the MW bit-security framework, the number of bits of statistical security of a decision game $\eth$ can be defined as follows.

**Definition 4.** *A distinguishing game* $\eth = (\eth_0, \eth_1)$ *has* $s$ *bits of statistical security if for every adversary* $A$*,* $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq 2^{-s}$*.*

Contrast this with the definition of (computational) bit-security, where the requirement is that $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq T(A) \cdot 2^{-c}$. It immediately follows from the definition that any problem achieving $s$ bits of statistical security, also offers $s$ bits of computational security. So, statistical bit-security is a strengthening of computational bit-security. In particular, when combining computational and statistical primitives within a single protocols, one can treat all of them has achieving a given number $c = s$ of computational security bits. However, this is often undesirable in practice because one typically wants to use a higher value of $c$ than for $s$. In order to combine computational and statistical bit-security analysis in an efficient manner, [LMSS22] proposes the following notion of *computational-statistical* bit-security.

**Definition 5** ([LMSS22])**.** *A distinguishing game* $\eth$ *is said to have* $(c, s)$*-bits of security if for any adversary* $A \in \mathcal{A}_{\perp}$*,*

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq \max(T(A)2^{-c}, 2^{-s}),$$

*i.e., either* $c \leq \log_2 \frac{T(A)}{\mathsf{adv}^{\mathsf{MW}}_{\eth}(A)}$*, or* $s \leq \log_2 \frac{1}{\mathsf{adv}^{\mathsf{MW}}_{\eth}(A)}$*.*

The notions of computational and statistical security corresponds to the following special cases of $(c, s)$-security:

- A problem has $c$ bits of computational security iff it is $(c, \infty)$-bit secure
- A problem has $s$ bits of computational security iff if is $(\infty, s)$-bit secure.

Since any problem offering $s$ bits of statistical security also offers $s$ bits of computational security, $(c, s)$-bit security is equivalent to $(\max(c, s), s)$-bit security. In other words, one can always assume $c \geq s$. In particular, computational security can be equivalently formulated as $(c, c)$-bit security, rather than $(c, \infty)$.

$(c, s)$-security is easily defined for search problems as well: A search game $\eth$ has $(c, s)$-bits of security if any adversary $A$ has success probability[11] $\Pr\{A(\eth)\}$ at most $\max(T(A)2^{-c}, 2^{-s})$.

# 3 Structure and properties of optimal MW adversaries

In this section we characterize the MW adversaries achieving optimal advantage, and prove some useful properties about them. This is done by introducing an alternative, more general, class of adversaries (which we call "fuzzy" adversaries,) that still achieves the same optimal advantage (and bit security) of standard MW adversaries. We use the added flexibility provided by fuzzy adversaries to investigate optimal adversarial strategies.

MW adversaries are generalized as follows. Recall that the output of an MW distinguisher is either a bit $b \in \{0, 1\}$, representing a *high confidence* decision between the two distributions, or a special symbol $\perp$ expressing *no confidence*. We generalize this to distinguishers for which the output confidence level can vary continuously from 0 (no confidence) to 1 (high confidence). For this type of distinguishers, it is convenient to map the two values $b \in \{0, 1\}$ to a sign

$$\tilde{b} = (-1)^b = (1 - 2b) = \pm 1 \tag{12}$$

so that the output of $A$ can be described by a single number $\sigma \in [-1, 1]$, with $\mathsf{sign}(\sigma) = \sigma/|\sigma| = \tilde{b} \in \{\pm 1\}$ representing the decision bit and $|\sigma| \in [0, 1]$ the confidence level.[12] We also set $\tilde{\perp} = 0$, so that any MW distinguisher $A$ with output $A(\eth_b) = y \in \{0, 1, \perp\}$ can be represented by a fuzzy one $\tilde{A}$ with output $\tilde{A}(\eth_b) = \tilde{y} \in \{1, -1, 0\} \subset [-1, 1]$. Notice that this transformation preserves the cost of the adversary $T(\tilde{A}) = T(A)$ as the only difference between the two is the symbol used to encode the final output. We write $\tilde{\mathcal{A}}_\perp = \{\tilde{A} \mid A \in \mathcal{A}_\perp\}$ for the set of aborting adversaries with this alternative output representation.

**Definition 6** (Fuzzy Distinguisher)**.** *A fuzzy distinguisher for a decision game $\eth = (\eth_0, \eth_1)$ is a (possibly randomized) adversary $A$ with output in $[-1, 1]$. The* advantage *and* bit-security *of $A$ in the game $\eth$ are $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \frac{\tilde{\delta}_A^2}{\tilde{\alpha}_A}$ and $\mathsf{MW}_{\eth}(A) = \log_2(T(A)/\mathsf{adv}^{\mathsf{MW}}_{\eth}(A))$ where*

$$\tilde{\delta}_A = \mathbb{E}_b[\tilde{b} \cdot A(\eth_b)] \qquad \text{and} \qquad \tilde{\alpha}_A := \mathbb{E}_b\left[|A(\eth_b)|\right]$$

*are the correlation (between the correct result and the output of $A$) and expected confidence of $A$. The set of all possible fuzzy distinguishers is denoted $\mathcal{A}_{\approx}$.*

[11] We recall that for a search problem, the output of $A(\eth)$ is determined by the game $\eth$.

[12] When $\sigma = 0$, the confidence $|\sigma| = 0$ is zero, and the decision $\mathsf{sign}(\sigma)$ is irrelevant. For concreteness, we define $\mathsf{sign}(0) = 0$.

Note that $\tilde{\mathcal{A}}_\perp \subset \mathcal{A}_\approx$, so we can view aborting adversaries as a special case of fuzzy adversaries. The following lemma shows that the definition of advantage given in Definition 6 respects this identification. This justifies the use of the same notation $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A)$ and $\mathsf{MW}_{\eth}(A)$ for the advantage and bit security of both aborting $A \in \mathcal{A}_\perp$ and fuzzy adversaries $A \in \mathcal{A}_\approx$.

**Lemma 9.** *For any aborting adversary $A \in \mathcal{A}_\perp$ and corresponding fuzzy adversary $\tilde{A} \in \mathcal{A}_\approx$ we have $\tilde{\delta}_{\tilde{A}} = \delta_A$, $\tilde{\alpha}_{\tilde{A}} = \alpha_A$, $T(\tilde{A}) = T(A)$, $\mathsf{adv}^{\mathsf{MW}}_{\eth}(\tilde{A}) = \mathsf{adv}^{\mathsf{MW}}_{\eth}(A)$ and $\mathsf{MW}_{\eth}(\tilde{A}) = \mathsf{MW}_{\eth}(A)$.*

*Proof.* It is easy to check that $\tilde{\delta}_{\tilde{A}} = \delta_A$ and $\tilde{\alpha}_{\tilde{A}} = \alpha_A$ by evaluating the expectations over the set $0, 1, -1$ of all possible values. It follows that $\mathsf{adv}^{\mathsf{MW}}_{\eth}(\tilde{A}) = \tilde{\delta}^2_{\tilde{A}}/\tilde{\alpha}_{\tilde{A}} = \delta^2_A/\alpha_A = \mathsf{adv}^{\mathsf{MW}}_{\eth}(A)$. Finally, $A$ and $\tilde{A}$ have the same running time $T(A) = T(\tilde{A})$. So, we also have

$$\mathsf{MW}_{\eth}(\tilde{A}) = \log_2(T(\tilde{A})/\mathsf{adv}^{\mathsf{MW}}_{\eth}(\tilde{A})) = \log_2(T(A)/\mathsf{adv}^{\mathsf{MW}}_{\eth}(A)) = \mathsf{MW}_{\eth}(A).$$

□

### 3.1 Equivalence of Aborting and Fuzzy adversaries

Using fuzzy adversaries, we may define the maximum (statistical) advantage in attacking a decision game $\eth$ as

$$(\Delta^{\approx}_{\mathsf{MW}}(\eth))^2 = \sup\{\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \mid A \in \mathcal{A}_\approx\},$$

and similarly for bit security

$$\mathsf{MW}_{\approx}(\eth) = \inf\{\mathsf{MW}_{\eth}(A) \mid A \in \mathcal{A}_\approx\}.$$

Since we are optimizing over a larger class of adversaries $\mathcal{A}_\approx \supset \mathcal{A}_\perp$, it immediately follows from the definitions that $\Delta_{\mathsf{MW}}(\eth) \le \Delta^{\approx}_{\mathsf{MW}}(\eth)$ and $\mathsf{MW}(\eth) \ge \mathsf{MW}_{\approx}(\eth)$, and in principle these inequalities could be strict. But, as we will see, this is not the case, i.e., aborting and fuzzy adversaries define precisely the same notion of advantage and bit security for decision games. This is proved using the following transformation.

**Lemma 10.** *Let $\mathsf{N}\colon \mathcal{A}_\approx \to \mathcal{A}_\perp$ be the transformation*[13]

$$\mathsf{N}[A](\eth_b; r) = \begin{cases} \frac{1-\mathsf{sign}(A(\eth_b;r))}{2} & \text{with probability } |A(\eth_b; r)| \\ \perp & \text{otherwise.} \end{cases}$$

*Then, for any decision game $\eth = (\eth_0, \eth_1)$ and adversary $A \in \mathcal{A}_\approx$, we have*

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \mathsf{adv}^{\mathsf{MW}}_{\eth}(\mathsf{N}[A]).$$

[13] More precisely, $\mathsf{N}[A]$ is the aborting adversary that runs the fuzzy attack $a \leftarrow A(\eth_b) \in [-1, 1]$, and then outputs $(1 - \mathsf{sign}(a))/2$ with probability $|a|$ and $\perp$ with probability $1 - |a|$. Note that the output of $\mathsf{N}[A]$ is always in $\{0, 1, \perp\}$, i.e., $\mathsf{N}[A] \in \mathcal{A}_\perp$ is a valid aborting adversary.

*Proof.* We have that

$$\begin{aligned}
\delta_{\mathsf{N}[A]} &= \Pr_{b,r}[\mathsf{N}[A](\eth_b; r) = b] - \Pr_{b,r}[\mathsf{N}[A](\eth_b; r) = 1 - b] \\
&= \mathbb{E}_b \left[ |A(\eth_b)| \cdot \Pr\left[ \frac{1 - \mathsf{sign}(A(\eth_b))}{2} = b \right] \right. \\
&\qquad \left. - |A(\eth_b)| \cdot \Pr\left[ \frac{1 - \mathsf{sign}(A(\eth_b))}{2} = 1 - b \right] \right] \\
&= \mathbb{E}_b[|A(\eth_b)| \cdot (\Pr[\mathsf{sign}(A(\eth_b)) = 1 - 2b] - \Pr[\mathsf{sign}(A(\eth_b)) = -(1 - 2b)])] \\
&= \mathbb{E}_b[|A(\eth_b)| \cdot (\Pr[\mathsf{sign}(A(\eth_b)) = (-1)^b] - \Pr[\mathsf{sign}(A(\eth_b)) = -(-1)^b])] \\
&= \mathbb{E}_b[|A(\eth_b)| \cdot (\Pr[(-1)^b \cdot \mathsf{sign}(A(\eth_b)) = 1] - \Pr[(-1)^b \cdot \mathsf{sign}(A(\eth_b)) = -1])] \\
&= \mathbb{E}_b[|A(\eth_b)| \cdot \mathbb{E}[(-1)^b \cdot \mathsf{sign}(A(\eth_b))]] \\
&= \mathbb{E}_b \left[ (-1)^b \cdot A(\eth_b) \right] = \delta_A,
\end{aligned}$$

and

$$\alpha_{\mathsf{N}[A]} = \Pr_{b,r}[\mathsf{N}[A](\eth_b; r) \neq \perp] = \mathbb{E}_{b,r}[|A(\eth_b; r)|] = \alpha_A.$$

It then follows that $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \frac{{\delta_A}^2}{\alpha_A} = \frac{{\delta_{\mathsf{N}[A]}}^2}{\alpha_{\mathsf{N}[A]}} = \mathsf{adv}^{\mathsf{MW}}_{\eth}(\mathsf{N}[A])$, i.e. $\mathsf{N}$ preserves the advantage. □

Clearly, the transformation $\mathsf{N}$ also preserves the complexity of the adversary $T(\mathsf{N}[A]) \approx T[A]$, as the additional operations performed by $\mathsf{N}[A]$ have negligible cost. It immediately follows that aborting and fuzzy adversaries are equivalent, both for statistical and computational bit security.

**Theorem 1.** *Aborting and Fuzzy MW adversaries are equivalent, i.e., they define the same notions of advantage and bit security*

$$\begin{aligned}
\Delta^{\approx}_{\mathsf{MW}}(\eth) &= \Delta_{\mathsf{MW}}(\eth) \\
\mathsf{MW}_{\approx}(\eth) &= \mathsf{MW}(\eth).
\end{aligned}$$

*Proof.* We need to show that $\Delta_{\mathsf{MW}}(\eth) \geq \Delta^{\approx}_{\mathsf{MW}}(\eth)$ and $\mathsf{MW}(\eth) \leq \mathsf{MW}_{\approx}(\eth)$. For any $A \in \mathcal{A}_{\approx}$, the aborting adversary $\mathsf{N}[A] \in \mathcal{A}_{\perp}$ satisfies

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) = \mathsf{adv}^{\mathsf{MW}}_{\eth}(\mathsf{N}[A]) \leq \sup_{A'} \mathsf{adv}^{\mathsf{MW}}_{\eth}(A') = \Delta^2_{\mathsf{MW}}(\eth_0, \eth_1).$$

Therefore, $(\Delta^{\approx}_{\mathsf{MW}}(\eth))^2 = \sup_A \mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq \Delta^2_{\mathsf{MW}}(\eth)$. A similar argument works for bit security, using the fact that $T(A) \approx T(\mathsf{N}(A))$. □

### 3.2 Convexity and Determinism

In general, cryptographic adversaries can use randomness. Using *fuzzy* adversaries it is easy to turn any randomized adversary into a deterministic one. In the following lemma we give a simple transformation from (randomized) aborting adversaries to deterministic fuzzy ones. For simplicity, we present the lemma for the simple problem of distinguishing between two probability distributions $\mathcal{X} = (X_0, X_1)$. A more general statement for arbitrary games will be proved later in this section.

**Lemma 11.** *Let* $\mathsf{F}\colon \mathcal{A}_\perp \to \mathcal{A}_\approx$ *be the transformation*

$$\mathsf{F}[A](x) = \Pr_r[A(x;r)=0] - \Pr_r[A(x;r)=1]$$

*mapping a (randomized) aborting adversary* $A$ *to a deterministic fuzzy advversary* $\mathsf{F}[A] \in \mathcal{A}_\approx$. *Then, for any decision problem* $\mathcal{X} = (X_0, X_1)$ *and adversary* $A \in \mathcal{A}_\perp$ *we have*

$$\mathsf{adv}^{\mathsf{MW}}_{\ni}(A) \le \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(\mathsf{F}[A]).$$

*In particular, the optimal advantage* $\Delta^2_{\mathsf{MW}}(\ni)$ *is achieved by a* deterministic $A \in \mathcal{A}_\approx$.

*Proof.* We first show that $\delta_{\mathsf{F}[A]} = \delta_A$. We have that

$$\begin{aligned}
\delta_{\mathsf{F}[A]} &= \mathbb{E}_b[\tilde{b} \cdot \mathsf{F}[A](X_b)] \\
&= \mathbb{E}_b[(-1)^b \cdot (\Pr_r[A(X_b;r)=0] - \Pr_r[A(X_b;r)=1])] \\
&= \frac{1}{2}\left(\mathbb{E}[(-1)\cdot(\Pr_r[A(X_1;r)=0] - \Pr_r[A(X_1;r)=1])]\right) \\
&+ \frac{1}{2}\left(\mathbb{E}[(+1)\cdot(\Pr_r[A(X_0;r)=0] - \Pr_r[A(X_0;r)=1])]\right) \\
&= \frac{\mathbb{E}[\Pr_r[A(X_1;r)=1]] + \mathbb{E}[\Pr_r[A(X_0;r)=0]]}{2} \\
&- \frac{\mathbb{E}[\Pr_r[A(X_1;r)=0]] + \mathbb{E}[\Pr_r[A(X_0;r)=1]]}{2} \\
&= \mathbb{E}_b[\Pr_r[A(X_b;r)=b]] - \mathbb{E}_b[\Pr_r[A(X_b;r)=1-b]] \\
&= \beta_A - \bar{\beta}_A = \delta_A.
\end{aligned}$$

We next show that $\alpha_{\mathsf{F}[A]} \le \alpha_A$. We have that

$$\begin{aligned}
\alpha_{\mathsf{F}[A]} &= \mathbb{E}_b[|\mathsf{F}[A](X_b)|] \\
&= \mathbb{E}_b\left[\left|\Pr_r[A(X_b;r)=0] - \Pr_r[A(X_b;r)=1]\right|\right] \\
&\le \mathbb{E}_b\left[\Pr_r[A(X_b;r)=0] + \Pr_r[A(X_b;r)=1]\right] = \alpha_A.
\end{aligned}$$

It follows that $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) = \frac{{\delta_A}^2}{\alpha_A} \le \frac{{\delta_{\mathsf{F}[A]}}^2}{\alpha_{\mathsf{F}[A]}} = \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(\mathsf{F}[A])$. □

Notice that the result of the transformation $\mathsf{F}[A]$ is not in general an efficient algorithm, because it requires the computation of the probabilities[14] $\Pr_r[A(x;r)=b]$ for $b=0,1$. So, Lemma 11 says little about the (computational) bit security under deterministic attacks. Moreover, it says nothing about the existence of deterministic *aborting* adversaries $A \in \mathcal{A}_\perp$ because $\mathsf{F}[A]$ is fuzzy.[15]

We would like to prove a similar result (for arbitrary decision games) that produces deterministic aborting adversaries, and address the efficiency issue (at least in the non-uniform setting). We will show that for any randomized aborting adversary $A$ there is a value of the randomness $r$ such that the deterministic adversary $A_r(\cdot) = A(\cdot;r)$ is at least as good as $A$. But before doing so, we observe that (perhaps contrary to intuition) this is not generally true for arbitrary notions of advantage.

[14] Naturally, one could approximate these probabilities in a relatively efficient manner by repeatedly running $A(x;r_i)$ on a given input $x$ and many independent random $r_i$. However, this would result in a randomized algorithm.

[15] Note that turning $\mathsf{F}[A]$ into an aborting adversary $\mathsf{N}[\mathsf{F}[A]]$ using Lemma 10 does not work, because the result of $\mathsf{N}$ is generally a randomized algorithm.

**Lemma 12.** *There is a pair of efficiently samplable distributions $\mathcal{X} = (X_0, X_1)$, randomized distinguisher $A(x; r)$ and advantage function $\mathsf{adv}^\star_{\mathcal{X}}$ such that $\mathsf{adv}^\star_{\mathcal{X}}(A)$ is strictly bigger than $\mathsf{adv}^\star_{\mathcal{X}}(A(\cdot; r))$ for all $r$.*

*Proof.* Let $X_0$ and $X_1$ be the uniform distributions over $\{0\}$ and $\{0, 1, 2, 3\}$ rexpectively, and consider a randomized distinguisher $A(x; r)$ using a single bit of randomness $r \in \{0, 1\}$ that works as follows: if $x \leq 2r$ then $A(x; r) = 0$, and $A(x; r) = \perp$ otherwise.

Consider the output of $A(x; r)$ when $x \leftarrow X_b$ for $b \leftarrow \{0, 1\}$. It is easy to see that $A(\cdot; r)$ is correct precisely when $b = 0$. So, $A$ has success probability $\beta = 1/2$ regardless of the value of the randomness $r$. On the other hand, the failure probability is $\bar{\beta}_0 = 1/8$ when $r = 0$ (for $b = 1$ and $x = 0$), $\bar{\beta}_1 = 3/8$ when $r = 1$ (for $b = 1$ and $x \in \{0, 1, 2\}$) and $\bar{\beta} = (\bar{\beta}_0 + \bar{\beta}_1)/2 = 1/4$ when $r$ is chosen at random. Now define the advantage function[16]

$$\mathsf{adv}^\star_{\mathcal{X}}(A) = \left|\beta_A - \sin(2\pi\bar{\beta}_A)\right| \cdot (\beta_A - \bar{\beta}_A).$$

Using this function we can compute $\mathsf{adv}^\star_{\mathcal{X}}(A) \approx 0.125$, $\mathsf{adv}^\star_{\mathcal{X}}(A(\cdot; 0)) \approx 0.077$ and $\mathsf{adv}^\star_{\mathcal{X}}(A(\cdot; 1)) \approx 0.025$. So, the advantage of the randomized adversary $A$ is strictly bigger than both $A(\cdot; 0)$ and $A(\cdot; 1)$. □

We prove the existence of deterministic optimal aborting adversaries using a convexity argument. For any adversaries $A, B \in \mathcal{A}_\perp$ and $\theta \in [0, 1]$, define the convex combination $C = \theta \cdot A + (1 - \theta) \cdot B$ as the (randomized) adversary that runs $A$ with probability $\theta$ and $B$ with probability $1 - \theta$. Notice that the convex combination is taken over the randomness, not the output of the adversaries, so that the result is still an aborting adversary in $\mathcal{A}_\perp$.

**Theorem 2.** *For any decision game $\ni$, the advantage $\mathsf{adv}^{\mathsf{MW}}_{\ni}(A)$ is a convex function of $A \in \mathcal{A}_\perp$, i.e., for any two adversaries $A, B \in \mathcal{A}_\perp$ and $\theta \in (0, 1)$, the convex combination $C = \theta \cdot A + (1-\theta) \cdot B \in \mathcal{A}_\perp$ satisfies*

$$\mathsf{adv}^{\mathsf{MW}}_{\ni}(C) \leq \theta \cdot \mathsf{adv}^{\mathsf{MW}}_{\ni}(A) + (1 - \theta) \cdot \mathsf{adv}^{\mathsf{MW}}_{\ni}(B).$$

*Proof.* Using the definition of $C$, we see that

$$\begin{aligned}\beta_C = \Pr[C(\ni_b) = b] &= \theta \cdot \Pr[A(\ni_b) = b] + (1 - \theta) \cdot \Pr[B(\ni_b) = b] \\ &= \theta \cdot \beta_A + (1 - \theta) \cdot \beta_B\end{aligned}$$

and similarly for $\bar{\beta}_C$, $\alpha_C$ and $\delta_C$. Therefore, by Lemma 1,

$$\begin{aligned}\mathsf{adv}^{\mathsf{MW}}_{\ni}(C) = \frac{\delta_C^2}{\alpha_C} &= \frac{(\theta \cdot \delta_A + (1 - \theta) \cdot \delta_B)^2}{\theta \cdot \alpha_A + (1 - \theta) \cdot \alpha_B} \\ &\leq \theta \cdot \frac{\delta_A^2}{\alpha_A} + (1 - \theta) \cdot \frac{\delta_B^2}{\alpha_B} \\ &= \theta \cdot \mathsf{adv}^{\mathsf{MW}}_{\ni}(A) + (1 - \theta) \cdot \mathsf{adv}^{\mathsf{MW}}_{\ni}(B).\end{aligned}$$

□

[16] Similarly to the MW advantage $(\beta_A - \bar{\beta}_A)^2/(\beta_A + \bar{\beta}_A)$ and statistical distance $(\beta_A - \bar{\beta}_A)$, we define this function as a simple combination of $\beta_A$ and $\bar{\beta}_A$. We included the $(\beta_A - \bar{\beta}_A)$ factor so that the advantage measure retains the appealing feature that "trivial adversaries" (with $\beta_A = \bar{\beta}_A$) have advantage 0. Our definition is otherwise rather arbitrary.

An immediate consequence of convexity is that optimal aborting adversaries $A \in \mathcal{A}_\perp$ can be easily derandomized by fixing the value of the randomness that achieves the highest advantage.

**Corollary 1.** *For any decision game $\eth$ and (randomized) adversary $A(\cdot; r)$, there is a value of $r$ such that the deterministic adversary $A_r(\cdot) = A(\cdot; r)$ has advantage at least $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A_r) \geq \mathsf{adv}^{\mathsf{MW}}_{\eth}(A)$.*

*Proof.* Any randomized adversary $A \in \mathcal{A}_\perp$ can be written as a convex combination $A = \sum_r \Pr[r] \cdot A_r$ of deterministic adversaries $A_r(\cdot) = A(\cdot; r)$ indexed by the randomness $r$. It follows by Theorem 2 that

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq \sum_r \Pr[r] \cdot \mathsf{adv}^{\mathsf{MW}}_{\eth}(A_r) \leq \max_r \mathsf{adv}^{\mathsf{MW}}_{\eth}(A_r). \tag{13}$$

Choosing the value of $r$ that achieves the maximum gives a deterministic adversary $A_r$ with an advantage which is at least as good as $A$. □

Note that the deterministic adversary of Corollary 1 has the same running time $T(A_r) = T(A)$ as the original randomized adversary because we are just fixing the randomness. So, Corollary 1 says that the optimal bit-security is achieved by a deterministic adversary. However, this is only true in a non-uniform setting, where the optimal randomness $r$ can be hardwired in the code of $A$. In a uniform setting, when considering probability ensembles over larger and larger sets indexed by a security parameter $\kappa$, determinining the optimal value of $r$ can be computationally difficult. In particular, trying all possible values of $r$ and estimating which one is best is not computationally feasible because there are exponentially (in $\kappa$) possible values of $r$, and, in any case, it would result again in a randomized adversary. This is the only result in this paper that makes essential use of the non-uniform model.

### 3.3 Threshold Adversaries are Optimal

In this subsection we focus on the simple problem of distinguishing between two probability distributions $\mathcal{X} = (X_0, X_1)$ over a set $\Omega$ (as opposed to arbitrary distinguishing games), in the statistical security setting, i.e., when the computational cost of the distinguisher is not taken into account. This problem reduces to determining the highest possible advantage $\Delta^2_{\mathsf{MW}}(\mathcal{X}) = \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)$ achieved by a (computationally unbounded) adversary $A$. All our adversaries can be implemented very efficiently given oracle access to the probabilities $\Pr\{X_b = x\}$. So, the results apply to the computational security setting as well when the probability distributions $X_0, X_1$ are efficiently computable.

In the case of traditional (non-aborting) adversaries, it is well known that this problem admits a very simple, closed form optimal distinguisher

$$A^{\mathcal{X}}_{\mathsf{SD}}(x) = \mathsf{sign}(\Pr\{X_0 = x\} - \Pr\{X_1 = x\})$$

which, on input a sample $x$, outputs the bit $b \in \{0, 1\}$ such that the probability $\Pr\{X_b = x\}$ is highest. Note that the distinguisher $A^{\mathcal{X}}_{\mathsf{SD}}$ is efficient only when the probabilities $\Pr\{X_0 = x\}, \Pr\{X_1 = x\}$ are efficiently computable. In this subsection we explore if a similar, closed form optimal distinguisher can also be described for the more general aborting $\mathcal{A}_\perp$ and fuzzy $\mathcal{A}_\approx$ adversaries.

The next lemma shows that even for fuzzy distinguishers, the "sign" of the output should be set to $\mathsf{sign}(A(x)) = \mathsf{sign}(\Pr\{X_0 = x\} - \Pr\{X_1 = x\})$ and the only extra freedom afforded by fuzzy adversaries is the choice of the confidence $|A(x)|$.

**Lemma 13.** *For any $A \in \mathcal{A}_{\approx}$ and $\mathcal{X} = (X_0, X_1)$, define the modified adversary*

$$\hat{A}(x) = |A(x)| \cdot \mathsf{sign}(\Pr\{X_0 = x\} - \Pr\{X_1 = x\})$$

*that on input $x$ outputs the same confidence $|A(x)|$ as $A$, and fixes the sign of the output to match $A^{\mathcal{X}}_{\mathsf{SD}}(x)$. Then, this modified adversary satisfies $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(\hat{A}) \geq \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)$.*

*Proof.* It is straightforward to verify that $\left|\hat{A}(x;r)\right| \leq |A(x;r)|$. So, the expected confidence of the modified adversary satisfies $\alpha_{\hat{A}} \leq \alpha_A$. We also have

$$\begin{aligned}
|\delta_A| &= \left|\mathbb{E}_{b;r}[(-1)^b \cdot A(X_b;r)]\right| \\
&= \left|\sum_x \mathbb{E}_r[A(x;r) \cdot \frac{\Pr\{X_0 = x\} - \Pr\{X_1 = x\}}{2}]\right| \\
&\leq \sum_x \mathbb{E}_r[|A(x;r)| \cdot \frac{|\Pr\{X_0 = x\} - \Pr\{X_1 = x\}|}{2}] \\
&= \sum_x \mathbb{E}_r[\hat{A}(x;r) \cdot \frac{\Pr\{X_0 = x\} - \Pr\{X_1 = x\}}{2}] \\
&= \left|\mathbb{E}_{b;r}[(-1)^b \cdot \hat{A}(X_b;r)]\right| = \delta_{\hat{A}}.
\end{aligned}$$

It follows that $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) = \delta_A^2/\alpha_A \leq \delta_{\hat{A}}^2/\alpha_{\hat{A}} = \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(\hat{A})$. □

Notice that if $A \in \tilde{\mathcal{A}}_{\perp}$, then $\hat{A} \in \tilde{\mathcal{A}}_{\perp}$. So, when applied to aborting adversaries, Lemma 13 shows that the adversary achieving the optimal advantage $\Delta^2_{\mathsf{MW}}(\mathcal{X})$ must agree with $A^{\mathcal{X}}_{\mathsf{SD}}$, except possibly for replacing the output with $\perp$ when confidence is low.

At this point we are left with the problem of determining how a fuzzy adversary should set the output confidence, and, as a special case, when an aborting adversary should output $\perp$. To this end, define the function

$$\ell_{\mathcal{X}}(x) = \log \frac{\Pr\{X_0 = x\}}{\Pr\{X_1 = x\}} = \log \Pr\{X_0 = x\} - \log \Pr\{X_1 = x\} \tag{14}$$

and consider the class of adversaries that output $\perp$ when $|\ell_{\mathcal{X}}(x)|$ is below a given threshold. Note that $\left|\ell_{(X_0,X_1)}(x)\right| = \left|\ell_{(X_1,X_0)}(x)\right|$ is symmetric in the ordering of the two distributions, and $\mathsf{sign}(\ell_{\mathcal{X}}(x)) = A^{\mathcal{X}}_{\mathsf{SD}}(x)$ because log is a monotonically increasing function.

**Definition 7.** *We say that $A \in \mathcal{A}_{\approx}$ is a threshold distinguisher between two distributions $\mathcal{X} = (X_0, X_1)$ over a set $\Omega$ if there is a threshold $\tau \geq 0$ such that*[17]

$$A(x) = \begin{cases} 0 & \text{if } |\ell_{\mathcal{X}}(x)| \leq \tau \\ \mathsf{sign}(\ell_{\mathcal{X}}(x)) & \text{if } |\ell_{\mathcal{X}}(x)| > \tau \end{cases}$$

[17]The choice that $A(x) = 0$ when $|\ell_{\mathcal{X}}(x)| = \tau$ is somehow arbitrary. We will use a threshold $\tau$ such that $|\ell_{\mathcal{X}}(x)| = \tau$ with probability 0.

**Theorem 3.** *Let $\mathcal{X} = (X_0, X_1)$ be a pair of probability distributions on a set $\Omega$. Then, the optimal advantage $\Delta^2_{\mathsf{MW}}(\mathcal{X})$ is achieved by a threshold distinguisher $A$. Moreover, the threshold*

$$\tau^* = \log\left(\frac{4}{3 - 2\beta^*_A} - 1\right)$$

*is a simple function of the conditional success probability $\beta^*_A = \beta_A/\alpha_A$. In particular, as $\beta^*_A \in [1/2, 1]$, the threshold satisfies $\exp(\tau) \in [1, 3]$.*

*Proof.* Let $A \in \mathcal{A}_{\approx}$ be an optimal fuzzy adversary, and assume without loss of generality that $\alpha_A, \delta_A > 0$, i.e., $A$ is non-trivial. Now, fix any point $x^*$ in the support of $X_0, X_1$, and define

$$\begin{aligned} \alpha^* &= \frac{\Pr\{X_0 = x^*\} + \Pr\{X_1 = x^*\}}{2} > 0 \\ \delta^* &= \frac{|\Pr\{X_0 = x^*\} - \Pr\{X_1 = x^*\}|}{2}. \end{aligned}$$

We will prove that

- $\frac{\delta^*}{\alpha^*} \le \frac{\delta_A}{2\alpha_A}$ if and only if $|\ell_{\mathcal{X}}(x^*)| \le \tau^*$ (and similarly for $\ge$),
- if $\frac{\delta^*}{\alpha^*} \le \frac{\delta_A}{2\alpha_A}$ then $|A(x^*)| = 1$, and
- if $\frac{\delta^*}{\alpha^*} \ge \frac{\delta_A}{2\alpha_A}$ then $|A(x^*)| = 0$.

In particular, since $|A(x^*)| = 0$ and $|A(x^*)| = 1$ are mutually exclusive, it must be $|\ell_{\mathcal{X}}(x^*)| \ne \tau^*$.

Note that the values $\alpha^*, \delta^*$ and $\tau^*$ satisfy

$$\begin{aligned} \frac{\delta_A}{2\alpha_A} &= \beta^*_A - \frac{1}{2} = 1 - \frac{2}{\exp(\tau^*) + 1} \\ \frac{\delta^*}{\alpha^*} &= \frac{\exp(|\ell(x^*)|) - 1}{\exp(|\ell(x^*)|) + 1} = 1 - \frac{2}{\exp(|\ell(x^*)|) + 1}. \end{aligned}$$

Since $\tau \mapsto 1 - 2/(\exp(\tau) + 1)$ is is a monotonically increasing function, this proves that $|\ell_{\mathcal{X}}(x^*)| \le \tau^*$ if and only if $(\delta^*/\alpha^*) \le \delta_A/(2\alpha_A)$.

Now assume $(\delta^*/\alpha^*) \le \delta_A/(2\alpha_A)$ and (for contradiction) $|A(x^*)| < 1$. Consider a modified adversary $A^*$ which is identical to $A$, except that $|A^*(x^*)| = |A(x^*)| + \epsilon$, for some $\epsilon < 1 - |A(x^*)|$. Using the definition of $\delta_A$ and $\alpha_A$, we get $\delta_{A^*} = \delta_A + \epsilon \cdot \delta^*$, and $\alpha_{A^*} = \alpha_A + \epsilon \cdot \alpha^*$. So, this modification increases the advantage of $A$ by

$$\begin{aligned} \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^*) - \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) &= \frac{(\delta_A + \epsilon \cdot \delta^*)^2}{\alpha_A + \epsilon \cdot \alpha^*} - \frac{\delta_A^2}{\alpha_A} \\ &= \frac{\epsilon^2(\delta^*)^2 + 2\epsilon\delta_A\alpha^*_A\left(\frac{\delta_A}{2\alpha_A} - \frac{\delta^*}{\alpha^*}\right)}{\alpha_A + \epsilon\alpha^*} \\ &\ge \frac{\epsilon^2(\delta^*)^2}{\alpha_A + \epsilon\alpha^*} > 0. \end{aligned}$$

This is a contradiction to the optimality of $A$.

Similarly, if $(\delta^*/\alpha^*) \geq \delta_A/(2\alpha_A)$ and (for contradiction) $|A(x^*)| > 0$, we may define a modified adversary $A^*$ that reduces the confidence $|A^*(x^*)| = |A(x^*)| - \epsilon$ of $A$ on $x^*$ by some $\epsilon < |A(x^*)|$. This increases the advantage of $A$ by

$$\begin{aligned}
\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^*) - \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) &= \frac{(\delta_A - \epsilon \cdot \delta^*)^2}{\alpha_A - \epsilon \cdot \alpha^*} - \frac{\delta_A^2}{\alpha_A} \\
&= \frac{\epsilon^2(\delta^*)^2 - 2\epsilon\delta_A\alpha_A^*\left(\frac{\delta_A}{2\alpha_A} - \frac{\delta^*}{\alpha^*}\right)}{\alpha_A - \epsilon\alpha^*} \\
&\geq \frac{\epsilon^2(\delta^*)^2}{\alpha_A - \epsilon\alpha^*} > 0,
\end{aligned}$$

again contradicting the optimality of $A$. □

# 4 Equivalence of MW and WY bit security

In [WY23], it is claimed that for any decision game $\eth$, the quantities $\mathsf{WY}(\eth)$ and $\mathsf{MW}(\eth)$ are equal up to an additive constant, i.e., the MW and WY notions of bit-security are equivalent. However, [WY23] only proves the statement for a variant of the WY security definition that uses aborting adversaries (i.e., the MW adversaries with output in $\{0, 1, \perp\}$ introduced in [MW18]), rather than the traditional (non-aborting, inner) adversaries used in [WY21] to define WY security. To close this gap, [WY23] informally states that changing the class of adversaries does not affect the definition of $\mathsf{WY}(\eth)$, and justifies the assertion saying that the definition does not *explicitly* depend on the size of the co-domain[18] of the adversary $A$. However, this reasoning is incorrect because the quantity $\Delta_{1/2}(A(\eth))$ used in the definition *implicitly* depends on the size of the co-domain of $A$. Still, the equivalence claimed in [WY23] holds true, as shown in the following theorem which gives a direct proof that $\mathsf{WY}_{\eth}$ and $\mathsf{MW}_{\eth}$ are equivalent.

The theorem makes use of the following technical lemma to modify an aborting adversary in such a way that it uses only two of the output symbols in $\{0, 1, \perp\}$.

**Lemma 14.** *For any decision game $\eth = (\eth_0, \eth_1)$, and aborting adversary $A \in \mathcal{A}_\perp$, there exists a modified adversary $A' \in \mathcal{A}_\perp$ with output in $\{\hat{b}, \perp\}$ (for some fixed $\hat{b} \in \{0, 1\}$) and similar running time $T(A) = T(A')$, such that*

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A') \geq \tfrac{1}{2} \cdot \mathsf{adv}^{\mathsf{MW}}_{\eth}(A).$$

*Proof.* Let $A' = A^z$ be the modified adversary from Lemma 6 with $z$ the value in $\{0, 1\}$ that maximizes the advantage $\mathsf{adv}^{\mathsf{MW}}_{\eth}(A^z)$. For $i \in \{0, 1\}, j \in \{0, 1, \perp\}$, let $p_{i,j} = \Pr\{A(\eth_i) = j\}$, so that $\beta_A = (p_{0,0} + p_{1,1})/2$, $\bar{\beta}_A = (p_{0,1} + p_{1,0})/2$ and, by Lemma 6,

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A^j) = \frac{(p_{0,j} - p_{1,j})^2}{2(p_{0,j} + p_{1,j})}.$$

[18] Recall that the co-domain of $A$ is the set of all possible outputs of $A$, e.g., $\{0, 1\}$ or $\{0, 1, \perp\}$.

We can then bound

$$\begin{aligned}\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) &= \frac{(\beta_A - \bar{\beta}_A)^2}{\beta_A + \bar{\beta}_A} \\ &= \frac{1}{2}\frac{((p_{0,0} - p_{1,0}) - (p_{0,1} - p_{1,1}))^2}{(p_{0,0} + p_{1,0}) + (p_{0,1} + p_{1,1})} \\ &\leq \frac{(p_{0,0} - p_{1,0})^2}{2(p_{0,0} + p_{1,0})} + \frac{(p_{0,1} - p_{1,1})^2}{2(p_{0,1} + p_{1,1})} \\ &= \mathsf{adv}^{\mathsf{MW}}_{\eth}(A^0) + \mathsf{adv}^{\mathsf{MW}}_{\eth}(A^1) \\ &\leq 2\mathsf{adv}^{\mathsf{MW}}_{\eth}(A^z)\end{aligned}$$

where the first inequality is Lemma 1, and the second one follows by our choice of $z$. □

We also need a variant of Lemma 14 which gives a tight connection between the MW advantage and the (squared) Le Cam distance of the adversary output probability distributions $A(\eth)$. A similar statement was previously proved in [WY23] under the condition that $\Delta_{1/2}(A(\eth)) \leq 1$, and with worse multiplicative constants.

**Lemma 15.** *For any decision game $\eth = (\eth_0, \eth_1)$ and aborting adversary $A \in \mathcal{A}_\perp$, there is a modified adversary $A' \in \mathcal{A}_\perp$ with similar running time $T(A) \approx T(A')$, such that*

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq \Delta^2_{\mathsf{LC}}(A(\eth)) \leq 3\mathsf{adv}^{\mathsf{MW}}_{\eth}(A').$$

*Proof.* The proof proceeds as in Lemma 14, using the same notation, except that this time $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^z)$ is maximized over $z \in \{0, 1, \perp\}$. As in the proof of Lemma 14, we still have

$$\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A) \leq \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^0) + \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^1).$$

To prove the new lemma we notice that

$$\Delta^2_{\mathsf{LC}}(A(X_0), A(X_1)) = \sum_{j \in \{0,1,\perp\}} \frac{(p_{0,j} - p_{1,j})^2}{2(p_{0,j} + p_{1,j})} = \sum_{j \in \{0,1,\perp\}} \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^j)$$

which is at least $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^0) + \mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^1)$ and at most $3\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A^z)$. □

**Theorem 4.** *For any decision game $\eth$, $\mathsf{WY}(\eth) = \mathsf{MW}(\eth) + \Theta(1)$.*

*Proof.* The inequality $\mathsf{MW}_\eth \leq \mathsf{WY}_\eth$ was already proved in [WY21]. Here we prove $\mathsf{WY}_\eth \leq \mathsf{MW}_\eth + O(1)$. Note that by Lemma 15, Lemma 2 and Lemma 3, for any adversary $A$, we have

$$\mathsf{adv}^{\mathsf{MW}}_{\eth}(A) \leq \Delta^2_{\mathsf{LC}}(A(\eth)) \leq 2 \cdot \Delta^2_{\mathsf{H}}(A(\eth)) \leq \Delta_{1/2}(A(\eth)).$$

So, by Lemma 14, for any adversary $A$ there is an adversary $A'$ such that

$$\begin{aligned}\mathsf{MW}_\eth(A) &= \log_2 \frac{T(A)}{\mathsf{adv}^{\mathsf{MW}}_{\eth}(A)} \\ &\geq \log_2 \frac{T(A')}{2\mathsf{adv}^{\mathsf{MW}}_{\eth}(A')} \\ &\geq \log_2 \frac{T(A')}{\Delta_{1/2}(A^z(\eth))} - 1.\end{aligned}$$

Note that $A'$ has co-domain $\{b, \perp\}$ (rather than $\{0,1\}$). But since $\Delta_{1/2}$ does not give any special meaning to the symbols output by the adversary, we can view $A'$ as a valid adversary for $\mathsf{WY}_{\eth}$. So, we get that $\mathsf{MW}_{\eth}(A) \geq \mathsf{WY}_{\eth}(A') - 1$. Since $A$ was arbitrary, this proves the theorem. □

The previous theorem shows that one can use $\mathsf{WY}(\eth)$ as an alternative characterization of $\mathsf{MW}(\eth)$. This is potentially interesting, as $\mathsf{WY}(\eth)$ only makes use of traditional (non-aborting) adversaries, which are perhaps more intuitive and easier to use. (This was indeed one of the motivations of [WY21].) In particular, it is tempting to assume that, since the inner adversary of [WY21] always outputs either 0 or 1 (i.e., it never aborts), the optimal WY advantage in distinguishing between two distributions $\mathcal{X} = (X_0, X_1)$ is achieved by the maximum likelihood distinguisher $A^{\mathcal{X}}_{\mathsf{SD}}$. Perhaps counterintuitively, the following theorem shows that this is not the case, and even if [WY21] does not make use of aborts, the obvious (inner) distinguishing strategy $A^{\mathcal{X}}_{\mathsf{SD}}$ is not optimal, and can in fact substantially overestimate the number of bits of security by a factor[19] close to 2.

**Theorem 5.** *There exist (efficiently samplable, efficiently computable) distributions $\mathcal{X} = (X_0, X_1)$ such that*

$$\mathsf{WY}_{\mathcal{X}}(A^{\mathcal{X}}_{\mathsf{SD}}) \geq 2 \cdot \mathsf{MW}(\mathcal{X}) - O(1).$$

*Proof.* The choice of $\mathcal{X}$ below is from [Sur21, Lemma 2], where it was used to show the suboptimality of distinguishing a product distribution $\mathcal{X}^{\otimes n} = (X_0^{\otimes n}, X_1^{\otimes n})$ by first computing $A^{\mathcal{X}}_{\mathsf{SD}}$ "coordinate-wise" (sometimes called *Scheffé's test*). Consider the distributions $\mathcal{X} = (X_0, X_1)$ shown in following table, where $\epsilon \leq 1/8$:

| | 0 | 1 | 2 |
|---|---|---|---|
| $X_0$ | $0.5$ | $0.5 - \epsilon$ | $\epsilon$ |
| $X_1$ | $0.5 - \epsilon$ | $0.5 + \epsilon$ | $0$ |
| $A^{\mathcal{X}}_{\mathsf{SD}}$ | $0$ | $1$ | $0$ |
| $A^{\mathcal{X}}_{\mathsf{MW}}$ | $\perp$ | $\perp$ | $0$ |
| $A^{\mathcal{X}}_{\mathsf{SD}}(X_0)$ | $0.5 + \epsilon$ | $0.5 - \epsilon$ | |
| $A^{\mathcal{X}}_{\mathsf{SD}}(X_1)$ | $0.5 - \epsilon$ | $0.5 + \epsilon$ | |

The table also shows the optimal $A^{\mathcal{X}}_{\mathsf{SD}}$ distinguisher, its output distribution on input $X_0$ and $X_1$, and a candidate[20] MW distinguisher which we will use in our proof. The intuition is clear: if the sample is 2, then it certainly comes from distribution $X_0$, but for the other samples the distinguisher does not have enough confidence to make the call. This distinguisher succeeds with probability $\beta = \epsilon/2$, but it never fails. So, it achieves advantage $(\beta - \bar{\beta})^2/(\beta + \bar{\beta}) = \beta = \epsilon/2$. Since $A_{\mathsf{MW}}$ runs in constant time, the decisional problem $\mathcal{X}$ has at most $\log_2(2/\epsilon) = 1 + \log_2(1/\epsilon)$ bits of security.

Let's now estimate the advantage achieved by $A_{\mathsf{SD}}$ as an inner distinguisher. We first evaluate the Hellinger distance

$$\Delta^2_{\mathsf{H}}(A^{\mathcal{X}}_{\mathsf{SD}}(X_0), A^{\mathcal{X}}_{\mathsf{SD}}(X_1)) = 1 - \sqrt{1 - 4\epsilon^2} \leq 4\epsilon^2$$

where we have used the inequality $1 - \sqrt{1-x} \leq x$, which is valid for all $x \in [0,1]$. Finally, using Lemma 3, we bound

$$\Delta_{1/2}(A^{\mathcal{X}}_{\mathsf{SD}}(\mathcal{X})) \leq 4\Delta^2_{\mathsf{H}}(A^{\mathcal{X}}_{\mathsf{SD}}(\mathcal{X})) \leq 16\epsilon^2.$$

[19] This is a doubling of the number of security bits $k$, so it corresponds to overestimating the cost of the attack by an exponential factor $2^k$.

[20] This is indeed the optimal MW distinguisher when $\epsilon \leq 1/8$. When $\epsilon \geq 1/8$, then $A^{\mathcal{X}}_{\mathsf{SD}}$ becomes optimal.

Since $A_{\mathsf{SD}}$ also runs in constant time, the upper bound on bit security it gives is $\log_2(1/(16\epsilon^2)) = 2\log_2(1/\epsilon) - 4$. In summary, if $\epsilon = 2^{-k}$ (for any $k \geq 3$), the bit security is at most $k+1$, but the WY framework with non-aborting distinguisher $A_{\mathsf{SD}}$ only provides a very weak bound of $2k-4$ $\square$

# 5 A Toolbox for Analysis of $(c, s)$-Bit Security

In this section we use the close relation between the MW and the Le Cam distance (Lemma 15) to establish two fundamental tools for the use of computational-statistical bit security in the analysis of complex cryptograhic protocols: the hybrid proof technique, and the probability replacement theorem.

**Theorem 6.** *Let $X_0, \dots, X_k$ be a sequence of cryptographic games. If for all $i = 1, \dots, k$, $\mathcal{X}_i = (X_{i-1}, X_i)$ is $(c_i, s_i)$-bit secure, then $\mathcal{X} = (X_0, X_k)$ is $(c, s)$-bit secure for*

$$\begin{aligned} c &= \min_i(c_i) - 2\log_2(\sqrt{3}k) \\ s &= \min_i(s_i) - 2\log_2(\sqrt{3}k) \end{aligned}$$

*Proof.* Using Lemma 15 we get the upper bound

$$\begin{aligned} \sqrt{\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)} &\leq \Delta_{\mathsf{LC}}(A(X_0), A(X_k)) \\ &\leq \sum_i \Delta_{\mathsf{LC}}(A(X_i), A(X_{i+1})) \\ &\leq \sqrt{3}\sum_i \max_{z_i} \sqrt{\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}_i}(A^{z_i})} \\ &\leq \sqrt{3}k\sqrt{\max_i(T(A^{z_i})2^{-c_i}, 2^{-s_i})}. \end{aligned}$$

So, since $T(A) \approx T(A^{z_i})$ for all $i$, the advantage $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)$ is at most

$$3k^2 \max(T(A)2^{-\min_i c_i}, 2^{-\min_i s_i}) = \max(T(A)2^{-c}, 2^{-s}).$$

This proves that $\mathcal{X}$ is at least $(c, s)$-secure. $\square$

This may be seen as an extension of [MW18, Theorem 7], which is an analogous result for $(c, c)$-bit security, though with slightly smaller[21] loss of $\log_2(2k^2) = 2\log_2(\sqrt{2}k)$ bits.

We next establish a distribution replacement theorem for $(c, s)$-bit security for games $\mathcal{D}^Y$ parameterized by a distribution $Y$. This was done in [MW18] under the assumption that $(Y_0, Y_1)$ is statistically ($(\infty, s)$-bit) secure, and in [WY23] under the assumption that $(Y_0, Y_1)$ is computationally ($(c, c)$-bit) secure. We extend this to a $(c, s)$-bit security assumption below.

[21] One can recover the exact same loss ($\log_2(2k^2) = 2\log_2(\sqrt{2}k)$) by giving a variant of Lemma 15 with constant factor 2 rather than 3. This can be done by comparing $\mathsf{adv}^{\mathsf{MW}}_{\mathcal{X}}(A)$ to $\Delta^2_{\mathsf{LC}}(X'_0, X'_1)$, where $X'_b \in [0,1]^2$ is the first two coordinates of $A(X_b) \in [0,1]^3$. This is to say that one can exactly generalize [MW18, Theorem 7] by working with $(X'_0, X'_1)$ that are positive measures of total mass $\leq 1$ rather than *probability* measures of total mass $= 1$. We avoid doing this as the quantitative improvement is small, at the cost of a large amount of conceptual overhead.

**Theorem 7.** *Let $\eth, \mathcal{Y}$ be decision games. If $\eth^{Y_0}$ is $(c,s)$-bit secure, and $\mathcal{Y}$ is $(c',s')$-bit secure, then $\eth^{Y_1}$ is $(c'',s'')$-bit secure, where $c'' = \min(c-2, c'-3-\log_2(1+T_\eth))$, and $s'' = \min(s-2, s'-3)$. In particular, if $\mathcal{Y}$ and $\eth^{Y_0}$ are $(c,s)$-bit secure and[22] $T_\eth = O(1)$, then $\eth^{Y_1}$ is almost $(c,s)$-bit secure, up to a small additive constant term in bit security.*

*Proof.* Let $A$ be any adversary. By Lemma 15 and the triangle inequality (for $\Delta_{\mathsf{LC}}$), we have We compute

$$\begin{aligned}\sqrt{\mathsf{adv}^{\mathsf{MW}}_{\eth^{Y_1}}(A)} &\le \Delta_{\mathsf{LC}}(A(\eth_0^{Y_1}), A(\eth_1^{Y_1})) \\ &\le \Delta_{\mathsf{LC}}(A(\eth_0^{Y_1}), A(\eth_0^{Y_0})) + \Delta_{\mathsf{LC}}(A(\eth_0^{Y_0}), A(\eth_1^{Y_0})) + \Delta_{\mathsf{LC}}(A(\eth_1^{Y_0}), A(\eth_1^{Y_1})).\end{aligned}$$

We bound each term in the last sum separately. For the middle term, using the upper bound in Lemma 15 and $T(A) = T(A^z)$, we get

$$\Delta_{\mathsf{LC}}(A(\eth_0^{Y_1}), A(\eth_0^{Y_0})) \le \sqrt{3 \max_z \mathsf{adv}^{\mathsf{MW}}_{\eth^{Y_0}}(A^z)} \le \sqrt{3\max(T_A 2^{-c}, 2^{-s})}$$

The other terms are bound constructing distinguishers $A_0, A_1$ against the game $\mathcal{Y}$ as follows. $A_0^Y$ simulates the execution of $A$ in the game $\eth_0^Y$ and flips the answer, i.e., it outputs $1-a$ when $A$ outputs $a \in \{0,1\}$, and $\perp$ otherwise. $A_1^Y$ simulates the execution of $A$ in the game $\eth_1^Y$, and outputs the same result as $A$. Then, we have

$$\begin{aligned}\Delta_{\mathsf{LC}}(A(\eth_0^{Y_1}), A(\eth_0^{Y_0})) &= \Delta_{\mathsf{LC}}(A_0(Y_0), A_0(Y_1)) \\ &\le \sqrt{3 \max_z \mathsf{adv}^{\mathsf{MW}}_{\mathcal{Y}}(A_0^z)} \\ &\le \sqrt{3\max(T(A)(1+T_\eth)2^{-c'}, 2^{-s'})}\end{aligned}$$

and similarly for the last term $\Delta_{\mathsf{LC}}(A(\eth_1^{Y_0}), A(\eth_1^{Y_0}))$ using adversary $A_1$. Combining the three terms gives the bound in the theorem. □

# 6 Conclusion and Open Problems

We developed a number of useful tools to evaluate the bit security of decisional cryptographic properties, in the statistical and computational setting, or even combinations of the two. These include a characterization of the structure of the optimal statistical "aborting" adversaries to facilitate the use of approximate probability distributions (like uniform or discrete gaussians), and general hybrid arguments and probability replacement theorems to combine subprotocols together and support modular security analysis. More tools may be added to the toolbox in the future, but we believe that the results presented in this paper already demonstrate that computational-statistical bit-security can be quite usable and useful.

For all results in this paper we focused on decision problems, which are the hardest case, but combining decisional primitives with search ones should be fairly straightforward, as the definition of bit security for search problems is standard. An interesting direction for future work is to

[22]Recall from Definition 1 that $T_\eth$ is the relative running time of $\eth$. So, $T_\eth = O(1)$ is quite common, e.g., when oracle calls can be answered in linear time.

explore the space between search and decision problems. These include, for example, problems with small (polynomially sized) search space, like password authenticated key exchange. Two works [MW18, Lee24] offer general definitions that interpolate between search and decision problems, but the significance of those definitions for intermediate problems is unclear. Similarly to what was done in [MW18] for search and decision problems, it would be interesting to analyze a representative set of protocols falling in-between search and decision primitives, possibly in conjunction with standard search and decision primitives, to see if the bit-security estimates provided by those definitions match the cryptographic intuition behind the informal notion of bit-security.

Another interesting direction for further work is to make good use of the definition of computational-statistical bit-security (proposed in [LMSS22] and studied in this work) to formally analyze concrete protocols of practical interest, and make provable (still tight) claims about their security.

## References


[ALWW21] Parhat Abla, Feng-Hao Liu, Han Wang, and Zhedong Wang. Ring-based identity based encryption - asymptotically shorter MPK and tighter security. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part III*, volume 13044 of *Lecture Notes in Computer Science*, pages 157–187. Springer, 2021.

[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, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part II*, volume 8270 of *Lecture Notes in Computer Science*, pages 321–340. Springer, 2013.

[DTT10] Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 649–665. Springer, 2010.

[GHVK11] Naira Grigoryan, Ashot Harutyunyan, Svyatoslav Voloshynovskiy, and Oleksiy Koval. On multiple hypothesis testing with rejection option. In *2011 IEEE Information Theory Workshop*, pages 75–79. IEEE, 2011.

[Lee24] Keewoo Lee. Bit security as cost to observe advantage: Towards the definition from the book. *Communications in Cryptology*, 2024.

[Lev93] Leonid A. Levin. Randomness and non-determinism. *Journal of Symbolic Logic*, 58:1102–1103, 1993.

[LJ20] Anusha Lalitha and Tara Javidi. On error exponents of almost-fixed-length channel codes and hypothesis tests. *arXiv preprint arXiv:2012.00077*, 2020.

[LMSS22] Baiyu Li, Daniele Micciancio, Mark Schultz, and Jessica Sorrell. Securing approximate homomorphic encryption using differential privacy. In Yevgeniy Dodis and Thomas

Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 560–589, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany.

[MW17] Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In *CRYPTO (2)*, volume 10402 of *Lecture Notes in Computer Science*, pages 455–485. Springer, 2017.

[MW18] Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 3–28, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.

[PJL23] Ankit Pensia, Varun Jog, and Po-Ling Loh. Communication-constrained hypothesis testing: Optimality, robustness, and reverse data processing inequalities. *IEEE Transactions on Information Theory*, 2023.

[PW22] Yury Polyanskiy and Yihong Wu. Information theory: From coding to learning. *Book draft*, 2022.

[Sur21] Ananda Theertha Suresh. Robust hypothesis testing and distribution estimation in Hellinger distance. In *International Conference on Artificial Intelligence and Statistics*, pages 2962–2970. PMLR, 2021.

[WY21] Shun Watanabe and Kenji Yasunaga. Bit security as computational cost for winning games with high probability. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part III*, volume 13092 of *Lecture Notes in Computer Science*, pages 161–188, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany.

[WY23] Shun Watanabe and Kenji Yasunaga. Unified view for notions of bit security. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4-8, 2023, Proceedings, Part VI*, volume 14443 of *Lecture Notes in Computer Science*, pages 361–389. Springer, 2023.

[Yas21] Kenji Yasunaga. Replacing probability distributions in security games via Hellinger distance. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography, ITC 2021, July 23-26, 2021, Virtual Conference*, volume 199 of *LIPIcs*, pages 17:1–17:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.