Affine Advantage in Guess-type Attacks on Goldreich's PRGs: Generalization, Characterization and Countermeasures

Ximing Fu[1] Mo Li[2] Shihan Lyu[2]

[1]Harbin Institute of Technology, Shenzhen
[2]The Chinese University of Hong Kong, Shenzhen

Abstract

The recent guess-type attack (Eurocrypt'26) on Goldreich's pseudorandom generators exploits the strong bias toward constants that can arise when some inputs of majority-based predicates are fixed. This attack has demonstrated substantial effectiveness against several concrete parameter instances (Eurocrypt'24&'25, Asiacrypt'24). In this paper, we generalize this attack paradigm to arbitrary affine approximations, characterize its capabilities, and develop a new predicate as a countermeasure.

We introduce a guess-and-affine-approximation (GAA) attack framework that generalizes existing guess-type attacks from constant approximations to arbitrary affine approximations of residual nonlinear predicates. To characterize the capacity of GAA attack, we introduce the bit-fixing affine advantage $\mathrm{adv}_{b'}(M_b)$, defined as the maximum affine advantage of a b-ary nonlinear predicate M_b over all residuals with b' unfixed inputs. For polynomial stretch $m = n^s$ and logarithmic nonlinear arity $b = \gamma \log_2 n$, we obtain the optimal attack complexities $C^\star$ linking bit-fixing affine advantages: (1) $C^\star > 2^{cn}$ for every $c < 2^{-(s-1)/\gamma}$ iff $\mathrm{adv}_{\theta b}(M_b) \to 0$ for every constant $\theta \in (0, 1]$; (2) $C^\star = 2^{o(n)}$ iff fixing only $o(b)$ predicate inputs can yield a residual with affine advantage $\frac{1}{2} - o(1)$, and the sufficient direction also holds for $b = \Omega(\log n)$. In particular, if $\mathrm{adv}_{b-r}(M_b) = \frac{1}{2} - O\left(\frac{\log n}{n}\right)$ for $r \le s - 1$, then $C^\star$ can further decrease to $n^{O(1)}$. These results complement prior analyses for arities $O(1)$, $\omega(1) \cap O(\log n)$, and $\Omega(n^{1/s})$.

Guided by the results above, we introduce the XOR_a–$\mathsf{MOD12}_b$ family as a predicate candidate. We prove that every b'-input residual of its nonlinear component retains algebraic degree and algebraic immunity $\Theta(b')$, matching the asymptotic orders of MAJ (STOC'16). Moreover, its bit-fixing affine advantage is at most $2^{-\Omega(b')}$, whereas $\mathrm{adv}_{\theta b}(\mathsf{MAJ}_b) \approx \frac{1}{2}$ for every fixed $\theta \in (0, 1)$. Analysis shows that, for $n = 1024$, $m = n^2$, and $b = 54$, XOR–MOD12 can achieve $\log C^\star \approx 737$, while XOR–MAJ and XOR–MOD3 have cost exponents of only 182 and 599, respectively, under the specified attack. Furthermore, our experiments show that MOD12 offers stronger resistance to belief-propagation-based attacks.

1 Introduction

Pseudorandom generators (PRGs) are a fundamental primitive in modern cryptography, expanding a short, uniformly random seed into a longer string that is computationally indistinguishable from uniform. In addition to security and stretch, computational efficiency is a key consideration in the design of PRGs. Goldreich's PRGs are particularly attractive in this regard, as they exhibit *locality*: each output bit depends on only a small number of seed bits [Gol00]. This property enables highly parallel evaluation by low-depth circuits, making Goldreich's PRGs well suited to cryptographic constructions that seek to minimize computational overhead. Uses include small-overhead secure computation [IKOS08, ADI+17, AK23], candidate constructions of indistinguishability obfuscation [LV17], and cryptographic capsules for preprocessing-based multiparty computation [BCG+17, CDM+18]. Other related applications include MPC-friendly symmetric-key primitives [GRR+16], oblivious transfer [BCM+24, CDD+24], garbling circuit [ZLY+26, LWYY25], and VOLE-based zero-knowledge proofs [ABBS25].

In Goldreich's PRG, each output bit is computed by applying the same predicate to publicly selected seed bits, making the choice of predicate central to security. The requirements on this predicate become more stringent when moving from linear stretch $m = \Theta(n)$ to polynomial stretch $m = n^s$ for a constant $s > 1$. We first review the main predicate properties identified in prior work for resisting known attacks.

1.1 Asymptotic Security and Predicate Design Criteria

A predicate P has *resilience* k if it is uncorrelated with the parity of every subset of at most k of its inputs. Suppose P has nonzero correlation with a parity involving ℓ inputs. When $m = \omega(n^{\ell/2})$, a birthday-paradox argument shows that with high probability there are two output bits whose XOR has constant bias. Taking the XOR of these bits therefore gives an $\mathbb{F}_2$-linear distinguisher. Consequently, resistance to such distinguishers at polynomial stretch $m = n^s$ requires $k \geq 2s - 1$ [OW14, AL16].

A predicate P with low *algebraic degree* $\deg(P)$ over $\mathbb{F}_2$ is intuitively vulnerable to linearization attacks, and resisting such attacks at polynomial stretch $m = n^s$ requires $\deg(P) > s$. However, high algebraic degree alone is insufficient, since fixing some inputs may reduce the degree and expose additional linear dependencies. To capture this vulnerability, Applebaum and Lovett [AL16] introduced the *r-bit-fixing degree* $\deg_r(P)$, defined as the minimum algebraic degree over all restrictions obtained by fixing r input positions to arbitrary Boolean values. They showed that if $r + \deg_r(P) < s$, a linear attack with constant advantage exists for almost all choices of the input positions. Thus, resistance to these attacks requires $r + \deg_r(P) \geq s$ for every r.

Applebaum and Lovett also showed that if the *algebraic immunity* $\mathrm{AI}(P) < s$, the generator can be efficiently distinguished from uniform by an algebraic refutation algorithm. They further proved that $\mathrm{AI}(P) > 8s + 1$ suffices for security against the class of algebraic attacks [AL16].

These criteria have guided the search for predicates that resist both linear and algebraic attacks. A prominent example is the XOR–MAJ predicate, which combines parity on a inputs with majority on b inputs [AL16], and a more recent alternative is the XOR_a–$\mathsf{MOD3}_b$ predicate [DMÜ26]. Although these predicates exhibit favorable asymptotic properties, their security at concrete parameter sizes remains a separate concern.

1.2 Concrete Security

Recent studies have investigated the concrete security of Goldreich's PRGs through various seed-recovery attacks, which we briefly review below.

Guess-and-Determine (GAD) attacks on XOR_3–AND_2 PRGs [CDM+18] exploit the fact that an AND_2 gate becomes linear once either input is known. By guessing some seed bits, the attacker can linearize output equations and recover the remaining seed bits via Gaussian Elimination (GE), provided that the resulting linear system has sufficient rank.

Unlike GAD, which exploits the specific property of AND_2, Guess-and-Decode (GADec) adopts a different approach by extracting probabilistic information from nonlinear constraints that remain after guessing [YGJL22]. Specifically, GADec employs belief propagation (BP) to iteratively refine the log-likelihood ratios (LLRs) of unknown seed bits, without requiring the remaining nonlinear constraints to be linearized. This framework can also be extended to XOR–MAJ through appropriate BP update rules, although computing these updates becomes increasingly expensive for larger nonlinear components.

The Group-and-Solve (GAS) attack also targets XOR–MAJ predicates [FLLL26]. It exploits the observation that fixing some inputs of the majority component can make it strongly biased toward a constant, allowing the original equations to be approximated by affine equations with high probability and subsequently solved using learning parity with noise (LPN) algorithms [MMT11, EKM17]. More recently, Guess-Filter-Solve (GFS) [LLMF26] and Bias-Amplified Pooled Recovery (BAPR) [WWLM26] were proposed to relax equation selection and retain more sufficiently biased equations.

In a complementary direction, [LLMF26] introduce Filter-Guess-Propagate (FGP), demonstrating that the number of equations involved in BP can be substantially reduced with little loss in success rate, while dynamic programming can significantly accelerate BP updates for nonlinear components. These improvements broaden the applicability of BP-based attacks, enabling FGP to attack several recent instances, including some instances of the XOR–MOD3 challenge [DMÜ26].

Together, these effective concrete attacks have affected many useful applications [CDD+24, BCM+24, ABBS25, BCPR25], and motivate the design of predicates that combine favorable asymptotic security properties with strong concrete security at short seed lengths, thereby enabling efficient instantiations of Goldreich's PRGs.

1.3 Our Contributions

To this end, we generalize existing guess-type attacks, characterize their cost within an idealized model, derive a new predicate design criterion, and construct a predicate correspondingly. Our contributions are threefold.

1. **A Generalized Attack Framework.** We introduce the Guess-and-Affine-Approximation (GAA) framework, which unifies the mechanisms underlying existing guess-based attacks and extends exact linearization and near-constant approximation to arbitrary affine approximations of residual predicates. Focusing on the GE-based strategies, including [CDM+18, FLLL26, WWLM26, LLMF26], we formulate an idealized cost model and derive a finite-parameter lower bound that jointly accounts for guessing probability, approximation reliability, and equation availability.

2. **Characterization of the Attack Capacity.** We introduce the bit-fixing affine advantage $\mathrm{adv}_{b'}(M_b)$, which measures the maximum affine advantage over all residuals of a nonlinear predicate M_b with b' unfixed inputs. For Goldreich's PRGs with polynomial stretch ($m = n^s$) and nonlinear arity $b = \gamma \log n$, we establish two asymptotic characterizations linking bit-fixing affine advantage to the optimal cost $C^\star$: (1) $C^\star > 2^{cn}$ for every $c < 2^{-(s-1)/\gamma}$ iff $\mathrm{adv}_{\theta b}(M_b) \to 0$ for every constant $\theta \in (0, 1]$, the optimal cost exponent can be made arbitrarily close to that of exhaustive seed search; (2) $C^\star = 2^{o(n)}$ iff fixing only $o(b)$ predicate

inputs can yield a residual with affine advantage $\frac{1}{2} - o(1)$, and the sufficient direction also holds for $b = \Omega(\log n)$. In particular, if $\mathrm{adv}_{b-r}(M_b) = \frac{1}{2} - O\left(\frac{\log n}{n}\right)$ for $r \leq s-1$, then C can further decrease to $n^{O(1)}$, which is polynomial. Together, they establish bit-fixing affine advantage as a theoretically motivated criterion for predicate design. Specifically, our framework characterizes which predicates admit subexponential cost with arity $b = \Omega(\log n)$, complementing the generic subexponential attacks for constant locality [App16], the $\omega(1) \cap O(\log n)$ arity [WWLM26] and the polynomial arity $b = \Omega(n^{1/s})$ [FLLL26].

3. **The XOR–MOD12 Predicate.** Guided by the proposed criterion, we introduce the XOR_a–$\mathsf{MOD12}_b$ predicate family. We prove that every b'-input residual of $\mathsf{MOD12}_b$ retains algebraic degree and algebraic immunity $\Theta(b')$, matching the asymptotic result of MAJ. Moreover, $\mathsf{MOD12}_b$ achieves $\mathrm{adv}_{b'}(\mathsf{MOD12}_b) \leq 2^{-\Omega(b')}$ uniformly over all partial input fixings, whereas $\mathrm{adv}_{\theta b}(\mathsf{MAJ}_b) \approx \frac{1}{2}$ for every constant $\theta \in (0,1)$. We further establish that 12 is the smallest fixed modulus for which a modular predicate can simultaneously achieve vanishing bit-fixing affine advantage and desired algebraic properties. We complement these theoretical results with concrete comparisons of the attack costs for MOD12, MOD3, and MAJ. For $n = 1024$, $m = n^2$, and $b = 54$, the attack cost for MOD12 is approximately 2^{737}, while the costs for MOD3 and MAJ under the specified attack are approximately 2^{599} and 2^{182}, respectively. Similarly, MOD12 offers stronger resistance to the BP decoding methods, according to our experimental evaluation.

The remainder of the paper is organized as follows. Section 2 introduces the preliminaries, including relevant properties of predicates and the definition of bit-fixing affine advantage. Section 3 presents the GAA attack framework and evaluates its optimal cost and establishes its asymptotic characterizations in terms of bit-fixing affine advantage. Section 4 introduces the XOR_a–$\mathsf{MOD12}_b$ construction, establishes its properties, and presents concrete evaluations. Finally, Section 5 concludes this paper and discusses the future work. All proofs are deferred to Appendices.

2 Preliminaries

We begin by introducing some basic notation used throughout the paper. For an integer $n \geq 0$, let $[n] \triangleq \{0, \ldots, n-1\}$. For a Boolean vector $\mathbf{x}$, let $\|\mathbf{x}\|$ denote its Hamming weight. For an index sequence $\sigma = (i_0, \ldots, i_{\ell-1})$ with $i_j \in [n]$, define $\mathbf{x}[\sigma] \triangleq (x_{i_0}, \ldots, x_{i_{\ell-1}})$. When a set is used as an index sequence, its elements are taken in increasing order. For every integer $d \geq 0$, let U_d denote a uniformly distributed random vector over $\{0,1\}^d$. For simplicity, we omit floor and ceiling operations in asymptotic statements whenever they do not affect the results, with integer-valued parameters implicitly rounded up when necessary.

A predicate $P_d : \{0,1\}^d \to \{0,1\}$ is symmetric if its value depends only on the Hamming weight of its input. For $d \geq 1$, define $\mathsf{XOR}_d(\mathbf{x}) \triangleq \bigoplus_{j \in [d]} x_j$ and $\mathsf{AND}_d(\mathbf{x}) \triangleq \prod_{j \in [d]} x_j$. For $1 \leq t \leq d$, the threshold predicate $\mathsf{THR}_d^t(\mathbf{x})$ equals 1 if $\|\mathbf{x}\| \geq t$ and 0 otherwise. The majority predicate is $\mathsf{MAJ}_d \triangleq \mathsf{THR}_d^{d/2}$, with ties resolved in favor of 1. For $t \geq 2$ and $S_0 \subseteq [t]$, define

$$\mathsf{MOD}_d^{t,S_0}(\mathbf{x}) \triangleq \begin{cases} 0, & \text{if } \|\mathbf{x}\| \bmod t \in S_0, \\ 1, & \text{otherwise.} \end{cases}$$

For positive integers a, b and a predicate $M_b : \{0,1\}^b \to \{0,1\}$, let $d = a + b$ and define $P_d \triangleq \mathsf{XOR}_a$–$M_b$ by

$$P_d(\mathbf{x}) \triangleq \mathsf{XOR}_a(x_0, \ldots, x_{a-1}) \oplus M_b(x_a, \ldots, x_{d-1}),$$

where $\mathbf{x} \in \{0,1\}^d$, and input-length subscripts are omitted whenever clear from context.

2.1 Goldreich's PRG and Seed-Recovery Attacks

Let n, d, m be positive integers with $d \le n$, and let $P_d : \{0,1\}^d \to \{0,1\}$ be a public predicate. Goldreich's PRG is specified by P_d and a public sequence $\Sigma_m = (\sigma_0, \ldots, \sigma_{m-1})$, where each σ_i is sampled independently and uniformly from all ordered d-tuples of distinct indices in $[n]$. For a seed $\mathbf{x} \in \{0,1\}^n$, the generator produces an output $\mathbf{y} \in \{0,1\}^m$ defined by

$$\mathbf{y} = G_{P_d,\Sigma_m}(\mathbf{x}) \triangleq (P_d(\mathbf{x}[\sigma_0]), \ldots, P_d(\mathbf{x}[\sigma_{m-1}])) .$$

For $P_d = \mathsf{XOR}_a$–M_b with $d = a + b$, write $\sigma_i = (\sigma_{i,1}, \sigma_{i,2})$, where $\sigma_{i,1}$ and $\sigma_{i,2}$ are the ordered subtuples consisting of the first a and last b indices of σ_i, respectively. The i-th output bit is then given by

$$y_i = \mathsf{XOR}_a(\mathbf{x}[\sigma_{i,1}]) \oplus M_b(\mathbf{x}[\sigma_{i,2}]).$$

Pseudorandomness requires the distribution ensembles

$$\{(\Sigma_m, G_{P_d,\Sigma_m}(U_n))\} \quad \text{and} \quad \{(\Sigma_m, U_m)\}$$

to be computationally indistinguishable, where U_n and U_m are sampled independently of Σ_m. In this work, we consider polynomial stretch $m = n^s$ for a constant $s > 1$.

In the seed-recovery problem, a hidden seed $\mathbf{x}$ is sampled uniformly from $\{0,1\}^n$, and the adversary is given $(n, m, P_d, \Sigma_m, \mathbf{y})$, where $\mathbf{y} = G_{P_d,\Sigma_m}(\mathbf{x})$. Unless stated otherwise, a seed-recovery attack succeeds if it outputs any $\widehat{\mathbf{x}} \in \{0,1\}^n$ satisfying

$$G_{P_d,\Sigma_m}(\widehat{\mathbf{x}}) = \mathbf{y}.$$

In particular, the recovered seed $\widehat{\mathbf{x}}$ need not equal the original seed $\mathbf{x}$. The optimal cost model in Section 3 accounts for the probability that the planted seed $\mathbf{x}$ belongs to the candidate set, whereas verification accepts any valid preimage.

2.2 Basic Properties of Predicates

We next define some useful properties of a predicate.

Definition 1 (Constraint and residual)**.** Let $d \ge 0$ be an integer. A *constraint* on d Boolean inputs is a pair $(T, \mathbf{v})$, where $T \subseteq [d]$ and $\mathbf{v} \in \{0,1\}^{|T|}$. An input $\mathbf{x} \in \{0,1\}^d$ satisfies the constraint $(T, \mathbf{v})$ if $\mathbf{x}[T] = \mathbf{v}$.

For a predicate $P_d : \{0,1\}^d \to \{0,1\}$, its *residual* under the constraint $(T, \mathbf{v})$ is the predicate $P_{d|T,\mathbf{v}} : \{0,1\}^{d-|T|} \to \{0,1\}$ defined by

$$P_{d|T,\mathbf{v}}(\mathbf{y}) \triangleq P_d(\mathbf{z}), \qquad \mathbf{z}[T] = \mathbf{v}, \quad \mathbf{z}[[d] \setminus T] = \mathbf{y}.$$

If P_d is symmetric, its residual depends only on $|T|$ and $\|\mathbf{v}\|$.

For $0 \le d' \le d$, let $\mathcal{R}_{d'}(P_d)$ denote the set of all residuals obtained by fixing any $d - d'$ inputs of P_d to arbitrary values.

Definition 2 (Balance, correlation immunity, and resiliency)**.** Let $P : \{0,1\}^d \to \{0,1\}$ be a predicate. It is balanced if $\Pr[P(U_d) = 1] = \frac{1}{2}$.

For an integer $0 \le k \le d$, P is k-correlation immune if, for every $0 \le r \le k$ and every $P' \in \mathcal{R}_{d-r}(P)$,

$$\Pr[P'(U_{d-r}) = 1] = \Pr[P(U_d) = 1].$$

P is k-resilient if it is both balanced and k-correlation immune.

For any predicate M_b and integer $a \geq 1$, the predicate XOR_a–M_b is $(a-1)$-resilient, since fixing at most $a-1$ inputs leaves at least one XOR input unfixed. Since this resiliency guarantee is independent of M_b, we do not analyze resiliency further in this work.

Definition 3 (Algebraic degree and algebraic immunity)**.** Let $P : \{0,1\}^d \to \{0,1\}$ be a predicate. Its algebraic normal form (ANF) is its unique multilinear representation over $\mathbb{F}_2$:

$$P(\mathbf{x}) = \bigoplus_{A \subseteq [d]} \lambda_A \prod_{j \in A} x_j, \qquad \lambda_A \in \mathbb{F}_2,$$

where the empty product is defined as 1. The *algebraic degree* of P is defined as $\deg(P) \triangleq \max\big(\{0\} \cup \{|A| : \lambda_A = 1\}\big)$.

A Boolean function $f_d \not\equiv 0$ is an annihilator of P if $f_d P \equiv 0$, where multiplication of Boolean functions is evaluated pointwise. The *algebraic immunity* of P is defined as

$$\mathrm{AI}(P) \triangleq \min\left\{\deg(f_d) : f_d \not\equiv 0, \quad f_d P \equiv 0 \text{ or } f_d(P \oplus 1) \equiv 0\right\},$$

where $P \oplus 1$ is the Boolean complement of P.

For any positive integers a, b and any predicate M_b, the predicate XOR_a–M_b satisfies $\deg(\mathsf{XOR}_a\text{–}M_b) = \max\{1, \deg(M_b)\}$ and $\mathrm{AI}(\mathsf{XOR}_a\text{–}M_b) = \mathrm{AI}(\mathsf{XOR}_1\text{–}M_b) \geq \mathrm{AI}(M_b)$. Thus, both properties are independent of the number a of XOR inputs.

2.3 Bit-fixing Affine Advantage

We introduce a new property called bit-fixing affine advantage to quantify how well a predicate or a residual can be approximated by an affine Boolean function.

Definition 4 (Affine approximation)**.** Let $P : \{0,1\}^d \to \{0,1\}$ be a predicate and $\mathcal{A}_d$ denote the set of affine Boolean functions on d inputs. For $f \in \mathcal{A}_d$, define the *affine agreement probability* between P and f by

$$q(P, f) \triangleq \Pr[P(U_d) = f(U_d)].$$

Definition 5 (Bit-fixing affine advantage)**.** Let $P : \{0,1\}^d \to \{0,1\}$ be a predicate. For $0 \leq d' \leq d$, let $\mathcal{R}_{d'}(P)$ denote the set of all residuals obtained by fixing any $d - d'$ inputs of P to arbitrary values. The *bit-fixing affine advantage* of P at d' remaining inputs is defined as

$$\mathrm{adv}_{d'}(P) \triangleq \max_{\substack{P' \in \mathcal{R}_{d'}(P) \\ f \in \mathcal{A}_{d'}}} \left(q(P', f) - \frac{1}{2}\right).$$

For any predicate P, we have $0 \leq \mathrm{adv}(P) \leq 1/2$, and $\mathrm{adv}(\mathsf{XOR}_a\text{–}M_b) = \mathrm{adv}(M_b)$ for any $a \geq 1$. We therefore focus on the affine advantage of the nonlinear component M_b.

3 Guess-and-Affine-Approximation Attack and Bit-Fixing Affine Advantage

This section introduces the Guess-and-Affine-Approximation (GAA) framework, which generalizes existing guess-based attacks to arbitrary affine approximations while allowing different solving strategies. We then develop an idealized cost model for the GE-based GAA, and characterize the optimal cost in terms of bit-fixing affine advantage, providing a theoretical basis for its use as a predicate design criterion.

3.1 GAA Attack

Consider Goldreich's PRG with predicate $P_d = \mathsf{XOR}_a$–M_b and seed $\mathbf{x} \in \{0,1\}^n$. A GAA attack proceeds in three stages: a global seed guess, affine approximation of the residual predicates induced by overlaps with the guessed positions, and solving the resulting equations.

A global guess consists of a set $R \subseteq [n]$ of seed positions and an assignment $\mathbf{u} \in \{0,1\}^R$. For each output equation i, let $\eta_i : [b] \to [n]$ map the nonlinear predicate inputs to their corresponding seed positions. The global guess induces a local restriction $(T_i, \mathbf{v}_i)$, where

$$T_i \triangleq \eta_i^{-1}(R), \qquad \mathbf{v}_i \triangleq (u_{\eta_i(j)})_{j \in T_i}.$$

Conditioned on a correct guess, the nonlinear component of equation i becomes the residual predicate $M_{b|T_i,\mathbf{v}_i}$ on $b - |T_i|$ unknown inputs.

For each residual $M_{b|T_i,\mathbf{v}_i}$, GAA attacks may choose an affine approximation $f_i \in \mathcal{A}_{b-|T_i|}$. Its agreement probability under uniform inputs is

$$q_i \triangleq q(M_{b|T_i,\mathbf{v}_i}, f_i).$$

Exact linearization and constant approximation are included as special cases. For example, when applied to the MAJ predicate, GAA includes the GFS attack [LLMF26] as a constant-approximation special case.

Each affine approximation yields a potentially noisy linear equation in the remaining seed bits. The attacker selects suitable equations based on their availability and approximation quality, and employs an appropriate solving algorithm, such as GE or noisy linear-equation decoding based on BP, to recover candidate seeds. The candidates are then verified against the original PRG outputs.

3.2 The Idealized Cost Model

We now focus on the GE-based GAA, covering GAD-, GAS-, GFS-, and BAPR-type strategies [CDM+18, FLLL26, LLMF26, WWLM26] while excluding BP-based decoding [YGJL22]. We introduce an idealized cost model for GE-based attacks, adopting attacker-favorable assumptions to estimate the optimal cost. In particular, the model assigns each overlap class the maximum affine approximation reliability among all corresponding residuals, deliberately relaxing the constraints imposed by the specific values of the global guess.

To evaluate a fixed global guess $(R, \mathbf{u})$, we condition on the seed weight $\|\mathbf{x}\| = w$ for mathematical analysis, without assuming that the attacker knows w.

For $0 \le w - \|\mathbf{u}\| \le n - |R|$, the conditional guessing probability is

$$p_{\text{guess}}(R, \mathbf{u}) \triangleq \Pr[\mathbf{x}[R] = \mathbf{u} \mid \|\mathbf{x}\| = w] = \frac{\binom{n-|R|}{w-\|\mathbf{u}\|}}{\binom{n}{w}}. \tag{1}$$

Let r denote the number of nonlinear predicate inputs fixed by the global guess, referred to as the overlap class. For a fixed R independent of the random output indices, the expected number of equations in class r is

$$\mu_r(R) \triangleq m \frac{\binom{|R|}{r}\binom{n-|R|}{b-r}}{\binom{n}{b}}. \tag{2}$$

We use this expectation as the capacity of class r in the cost model.

Conditioned on a correct guess, the remaining seed is uniformly distributed over all $(n-|R|)$-bit vectors of weight $w - \|\mathbf{u}\|$. For $b' \le n - |R|$, let $Z_{b'}^{(R,\mathbf{u})}$ denote the values at a uniformly sampled ordered b'-tuple of distinct positions in the remaining seed.

Then, for each feasible overlap class r, define

$$q_r(R, \mathbf{u}) \triangleq \max_{\substack{M' \in \mathcal{R}_{b-r}(M_b) \\ f \in \mathcal{A}_{b-r}}} \Pr\left[M'(Z_{b-r}^{(R,\mathbf{u})}) = f(Z_{b-r}^{(R,\mathbf{u})})\right], \tag{3}$$

as the maximum fixed-weight agreement probability among all residuals obtained by fixing r inputs of M_b and their affine approximations. The attacker is allowed to assign this maximum agreement probability to every selected equation in class r, regardless of whether the corresponding residuals can be realized under the global guess.

Let k_r denote the number of selected equations from overlap class r, and write $\mathbf{k} \triangleq (k_r)_r$ and $k \triangleq \sum_r k_r$. The guess procedure imposes the capacity constraints

$$0 \le k_r \le \mu_r(R), \qquad \sum_r k_r \le n - |R|. \tag{4}$$

Following the discussion in prior works [FLLL26, WWLM26, LLMF26], we assume that the approximation errors of the selected equations are independent conditioned on a correct guess. Under this assumption, the modeled probability that all selected approximations are correct is

$$p_{\mathrm{sub}}(\mathbf{k}) \triangleq \prod_r q_r(R, \mathbf{u})^{k_r}.$$

We further assume that the selected equations are linearly independent, leaving $n - |R| - k$ free variables and $2^{n-|R|-k}$ candidate assignments.

Under these assumptions, we define the following cost, measured in units of GE with polynomial factors suppressed:

$$\begin{aligned} C(R, \mathbf{u}, \mathbf{k}) &\triangleq \frac{2^{n-|R|-k}}{p_{\mathrm{guess}}(R, \mathbf{u}) p_{\mathrm{sub}}(\mathbf{k})} \\ &= \frac{\binom{n}{w}}{\binom{n-|R|}{w-\|\mathbf{u}\|}} 2^{n-|R|-k} \prod_r q_r(R, \mathbf{u})^{-k_r}. \end{aligned} \tag{5}$$

The corresponding optimal cost is therefore

$$C^\star \triangleq \min_{R, \mathbf{u}, \mathbf{k}} C(R, \mathbf{u}, \mathbf{k}), \tag{6}$$

where the minimum is over feasible global guesses and nonnegative real-valued equation counts satisfying (4). This cost measures candidate enumeration work divided by the modeled probability of a correct guess and error-free affine approximations. The underlying GE-based strategy can be implemented using the actual residual predicates. Under our parameter setting, affine approximation, equation selection, and Gaussian elimination incur only polynomial overhead per trial. The model omits this overhead and evaluates success probabilities under the stated modeling assumptions, yielding an idealized work factor. Under the same modeling assumptions, $C^\star$ lower-bounds the modeled costs of covered instantiations with Gauss/SubGauss solvers [CDM+18, FLLL26, WWLM26, LLMF26].

Notably, the minimization is purely analytical and does not assume that the attacker knows w or can identify the optimal configuration.

3.3 Asymptotic Analysis

We now study how the bit-fixing affine advantages of M_b affect the cost as n grows. Fix a family $(M_b)_{b\geq 1}$ and a constant $s > 1$. For any constant $\gamma > 0$, let

$$m = n^s, \qquad b = \gamma \log_2 n.$$

We consider seed weights close to $n/2$, defined by

$$\mathcal{W}_n \triangleq \left\{ w \in [n+1] : \left| w - \frac{n}{2} \right| \leq \sqrt{n \ln n} \right\}.$$

Here, we consider a relatively balanced seed, strengthening our results. The imbalance can naturally lead to bias and has been exploited in [DMÜ26]. The Hamming weight of a uniformly random seed belongs to $\mathcal{W}_n$ with probability at least $1 - 2/n^2$. We write $C^\star(n, w; \gamma)$ to make its dependence on these parameters explicit. Throughout this subsection, $\gamma > 0$ is a constant, and all asymptotic bounds are uniform over the stated range of w.

Theorem 1 (Exponential cost)**.** *For any predicate family $(M_b)_{b\geq 1}$ and every $c \in (0, 2^{-(s-1)/\gamma})$, the optimal cost satisfies*

$$C^\star(n, w; \gamma) > 2^{cn}$$

for all sufficiently large n, if and only if

$$\lim_{b\to\infty} \mathrm{adv}_{\theta b}(M_b) = 0$$

for every fixed $\theta \in (0, 1]$.

The equivalence above is between the stated advantage condition and the cost bound holding for all admissible pairs (γ, c). In particular, under this condition, increasing γ makes the optimal cost exponent arbitrarily close to that of exhaustive seed search.

Corollary 1. *For any predicate family $(M_b)_{b\geq 1}$, the optimal cost satisfies*

$$C^\star(n, w; \gamma) \leq 2^{(2^{-(s-1)/\gamma} + o(1))n}.$$

Together with Theorem 1, this establishes the tightness of its exponential optimal cost bound.

Theorem 2 (Subexponential cost)**.** *For any predicate family $(M_b)_{b\geq 1}$, the optimal cost satisfies*

$$C^\star(n, w; \gamma) = 2^{o(n)}$$

if and only if

$$\lim_{b\to\infty} \mathrm{adv}_{\theta b}(M_b) = \frac{1}{2}$$

for every fixed $\theta \in (0, 1)$.

Theorem 2 characterizes the subexponential regime of the cost model. Its condition is equivalent to the existence of restrictions fixing only $o(b)$ predicate inputs, under which the residual predicates have affine agreement approaching one. Together with Theorem 1, this shows how the bit-fixing affine advantage profile captures distinct asymptotic cost regimes. For intermediate advantage profiles, the optimal cost may still be exponential, but its exponent is not guaranteed to approach one as γ increases.

Table 1: Comparison of asymptotic complexity.

Method	Condition	Complexity
Generic inversion [App16]	$b = O(1)$	$2^{o(n)}$
BAPR* [WWLM26]	$b = \omega(1) \cap O(\log n)$	$2^{o(n)}$
GAS* [FLLL26]	$b = \Omega(n^{1/s})$	$2^{n(1-\Theta(s \cdot b^{-1/2}))}$
Our work	$b = \Omega(\log n) \wedge \mathrm{adv}_{\theta b}(M_b) \to \frac{1}{2}$	$2^{o(n)}$
	$b = \Omega(\log n) \wedge \mathrm{adv}_{b-r}(M_b) = \frac{1}{2} - O\left(\frac{\log n}{n}\right)$	$\mathrm{poly}(n)$

*The BAPR and GAS results reported here are only for XOR–MAJ predicates. The BAPR additionally requires that the XOR arity a is constant and the MAJ arity b is odd.

For comparison, the generic inversion method described in [App16] yields subexponential-time attacks against constant-locality random local functions with polynomial stretch. BAPR [WWLM26] gives a $2^{o(n)}$ seed-recovery attack against XOR–MAJ for $b \in \omega(1) \cap O(\log n)$ when the majority arity b is odd and the XOR arity $a \geq 3$ is constant. The GAS attack can achieve $2^{n(1-\Theta(s \cdot b^{-1/2}))}$ complexity for $b = \Omega(n^{1/s})$. By comparison, our method extends the result to $b = \Theta(\log n)$ and $b = \Omega(\log n)$. Table 1 summarizes the complexity comparison.

Moreover, when the bit-fixing affine advantage approaches $1/2$ sufficiently rapidly after fixing only a constant number of inputs, the optimal complexity can degrade to polynomial, as shown in Corollary 2.

Corollary 2. *For any predicate family $(M_b)_{b \geq 1}$, the optimal cost satisfies*

$$C^\star(n, w; \gamma) = n^{O(1)}$$

if

$$\mathrm{adv}_{b-r}(M_b) = \frac{1}{2} - O\left(\frac{\log n}{n}\right)$$

for some nonnegative integer $r \leq s - 1$.

Corollary 3. *For any predicate family $(M_b)_{b \geq 1}$ with $b = \Omega(\log n)$, the optimal cost satisfies*

$$C^\star(n, w) = 2^{o(n)}$$

if

$$\lim_{b \to \infty} \mathrm{adv}_{\theta b}(M_b) = \frac{1}{2}$$

for every fixed $\theta \in (0, 1)$.

4 Predicate Design

Guided by Theorem 1, we seek nonlinear predicates M_b with vanishing bit-fixing affine advantage $\mathrm{adv}_{b'}(M_b) \to 0$ as $b' \to \infty$, while every residual $M' \in \mathcal{R}_{b'}(M_b)$ retains algebraic degree and algebraic immunity $\Theta(b')$. The latter two requirements match the asymptotic order of MAJ, while the first provides stronger resistance to affine approximation attacks.

4.1 Properties of Modulus Functions

Intuitively, the periodic dependence of modular predicates on input Hamming weight suggests that their bit-fixing affine advantage may increase slowly as more inputs are fixed. Prior work [DMÜ26] established that $\mathsf{MOD3}$ satisfies the desired algebraic requirements, but its bit-fixing affine advantage does not vanish. We therefore investigate the broader family of modular predicates.

Theorem 3 (Bit-fixing affine advantage criterion for modular functions)**.** *Fix an integer $t \geq 2$ and a residue set $S_0 \subseteq [t]$. Then*

$$\lim_{b' \to \infty} \sup_{b \geq b'} \mathrm{adv}_{b'}(\mathsf{MOD}_b^{t,S_0}) = 0$$

if and only if $4 \mid t$ and

$$|S_0 \cap 2\mathbb{Z}| = |S_0 \cap (2\mathbb{Z}+1)| = \frac{t}{4}.$$

Theorem 3 shows that vanishing bit-fixing affine advantage is possible only when the modulus is divisible by four. For $t \in \{4, 8\}$, the simplified value vector of any modular predicate MOD_b^{t,S_0} is periodic with period t. By the relationship between periodicity and algebraic degree for symmetric Boolean functions [CV05, Theorem 1], its algebraic degree is at most $t-1$, independently of b. Consequently, 12 is the smallest modulus not excluded by the joint requirements of vanishing bit-fixing affine advantage and unbounded algebraic degree.

4.2 Properties of $\mathsf{MOD12}$

Notably, satisfying the criterion in Theorem 3 does not necessarily guarantee the desired algebraic properties. A simple search led us to the residue set

$$S_0 = \{0, 1, 3, 4, 6, 9\}.$$

For positive integers a and b, we define

$$\mathsf{XOR}_a\text{–}\mathsf{MOD12}_b \triangleq \mathsf{XOR}_a\text{–}\mathrm{MOD}_b^{12,S_0}.$$

We establish its remaining properties below.

Theorem 4 (Bit-fixing algebraic degree)**.** *Let $b \geq 6$, $6 \leq b' \leq b$. For every residual $M' \in \mathcal{R}_{b'}(\mathsf{MOD12}_b)$, we have $b' - 5 \leq \deg(M') \leq b'$, and hence $\deg(M') = \Theta(b')$.*

Theorem 5 (Bit-fixing algebraic immunity)**.** *Let $b \geq 6$, $6 \leq b' \leq b$. For every residual $M' \in \mathcal{R}_{b'}(\mathsf{MOD12}_b)$, we have*

$$\mathrm{AI}(M') \geq \left\lfloor \frac{b'+1}{3} \right\rfloor.$$

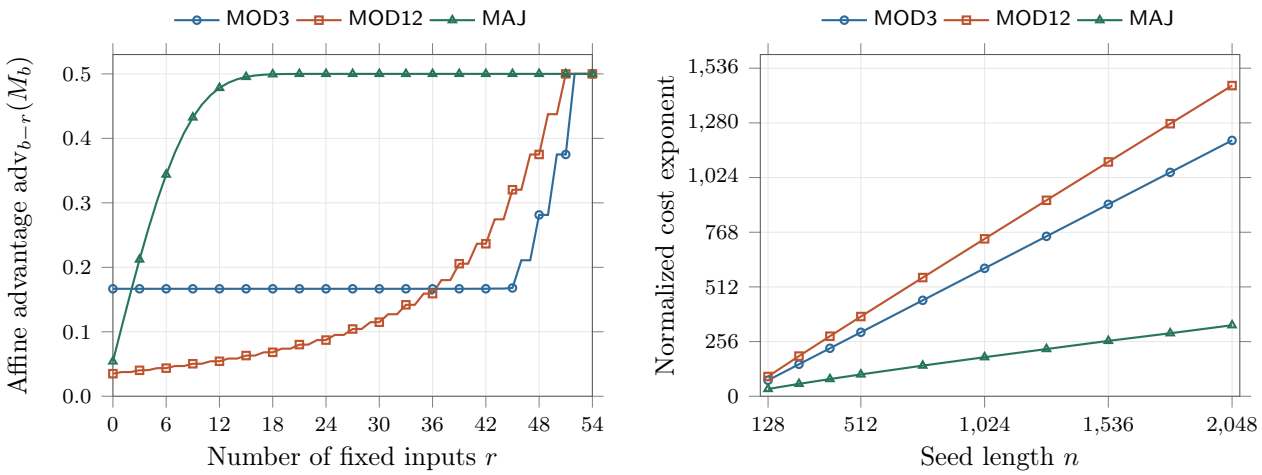
Consequently, $\mathrm{AI}(M') = \Theta(b')$.

Together, these results show that $\mathsf{MOD12}_b$ provides a concrete predicate construction that simultaneously meets our algebraic and affine-approximation design criteria.

Moreover, we give an explicit upper bound on the bit-fixing affine advantage of MOD12.

Theorem 6 (Bit-fixing affine advantage of $\mathsf{MOD12}$)**.** *For integers $b \geq 1$ and $1 \leq b' \leq b$, define*

$$B_{12}(b') = \frac{\sqrt{2}}{6}\left[\left(\cos\frac{\pi}{12}\right)^{b'} + \left(\sin\frac{\pi}{12}\right)^{b'} + 2^{-b'/2}\right] + \frac{1}{\sqrt{3}}\left(\frac{\sqrt{3}}{2}\right)^{b'}.$$



(a) Affine advantage after fixing r inputs.

(b) Solving costs as the seed length increases.

Figure 1: Concrete comparison of MOD3, MOD12, and MAJ with $(s, a, b) = (2, 5, 54)$. Left: bit-fixing affine advantage under uniformly random unfixed inputs. Right: base-two logarithms of normalized solving costs as n increases, with $m = n^2$ and $w = n/2$, using optimal cost for MOD12, GFS with SubGauss for MAJ, and constant-prediction SubGauss adaptation for MOD3.

Then

$$\mathrm{adv}_{b'}(\mathsf{MOD12}_b) \leq \min\left\{\frac{1}{2}, B_{12}(b')\right\}. \tag{7}$$

4.3 Concrete Comparison

4.3.1 Bit-fixing Affine Advantage and Cost Evaluation

We compare $\mathsf{MOD3}_{54}$, $\mathsf{MOD12}_{54}$, and MAJ_{54} at the same nonlinear arity $b = 54$. Figure 1(a) plots their bit-fixing affine advantages under uniformly random unfixed inputs. The advantage of MAJ rapidly increases and reaches $1/2$ after fixing 27 inputs. The advantage of MOD3 remains close to $1/6$ over a broad range of restrictions, whereas MOD12 has a smaller advantage for $0 \leq r \leq 36$.

For the cost comparison, we fix $s = 2$ and $a = 5$, set $m = n^2$, and evaluate ten seed lengths from 128 to 2048 under the balanced seed-weight condition $w = n/2$. We compute the optimal cost for MOD12, optimize GFS with SubGauss for MAJ [LLMF26], and apply the constant-prediction adaptation of SubGauss to MOD3. For the latter, optimization selects the no-guess strategy at every tested seed length. All cost estimates use the conditional fixed-weight marginal and retain the corresponding independence and rank assumptions. We report normalized solving costs, excluding filtering, verification, polynomial GE costs, and constant repetition factors.

Figure 1(b) shows that MOD12 has the largest cost exponent throughout the evaluated range. At $n = 1024$, the exponents are approximately 737, 599, and 182 for MOD12, MOD3, and MAJ, respectively. Thus, even under the optimistic model, the evaluated cost for MOD12 exceeds the estimated solving work of the specified attacks on the other two predicates.

4.3.2 Empirical Evaluation of FGP

The preceding analysis evaluates seed-recovery methods based on GE. FGP [LLMF26], however, uses belief propagation (BP) to extract information from the remaining nonlinear equations. We therefore

Table 2: Estimated FGP attack work on XOR–MOD3 and XOR–MOD12 with $m = n^s$.

Parameters				XOR–MOD3		XOR–MOD12	
n	s	a	b	t	$\log_2 \widehat{T}$	t	$\log_2 \widehat{T}$
512	2	5	54	33	61	89	123
512	2	5	114	33	63	282	319
512	2	11	54	182	209	222	253
512	2	11	114	183	212	320	355
1024	2	5	54	51	81	145	180
1024	2	5	114	51	82	450	489
1024	2	11	54	332	361	415	448
1024	2	11	114	334	364	590	626
512	3	11	54	88	116	119	152
1024	3	11	54	147	177	203	239

Here, t denotes the number of guessed seed bits. The estimated work $\widehat{T}$ is computed following the complexity model of FGP [LLMF26].

evaluate FGP separately to compare XOR–MOD3 and the proposed XOR–MOD12 predicate under BP-based seed recovery.

We compare the two predicates using matching parameter settings (n, s, a, b), with $m = n^s$. For both predicates, we use 40 BP iterations and 2000 trials for each tested value of t. Table 2 reports t and the corresponding estimated work exponent $\log_2 \widehat{T}(t)$ for each predicate, using the complexity estimation method adopted in the FGP attack.

The results show that XOR–MOD12 is substantially more resistant to the evaluated BP-based FGP attack than XOR–MOD3. For small a, FGP is particularly effective against XOR–MOD3. The substantial increase in estimated work obtained by replacing MOD3 with MOD12 therefore highlights the benefit of the proposed predicate in a parameter regime favorable to FGP. For larger a, FGP itself is less well suited to seed recovery, resulting in high estimated attack costs for both predicates. Nevertheless, XOR–MOD12 continues to impose a substantial additional recovery cost.

The advantage of MOD12 becomes particularly pronounced as b increases. For example, with $n = 1024$ and $a = 5$, increasing b from 54 to 114 raises the estimated work exponent for XOR–MOD12 from 180 to 489, while the corresponding exponent for XOR–MOD3 changes only from 81 to 82.

5 Conclusion

In this work, we generalize existing guess-type attacks to the GAA framework. We establish asymptotic connections between bit-fixing affine advantage and the optimal cost, providing a theoretical basis for predicate design. Guided by this criterion, we propose the XOR_a–$\mathsf{MOD12}_b$ predicate.

As the result for $\omega(1) \cap O(\log n)$ is specifically targeted to XOR–MAJ, it is an interesting direction to discuss the results for general predicates for the future work.

A Proofs of the Asymptotic Results

Throughout this appendix, for a feasible global guess $(R, \mathbf{u})$, we write

$$t \triangleq |R|, \qquad n' \triangleq n - t, \qquad w' \triangleq w - \|\mathbf{u}\|.$$

The residual weight w' is used only for analysis and is not assumed to be known to the attacker.

A.1 Distributional Lemmas

Lemma 1 (Fixed-weight and uniform agreement)**.** *Let $(R, \mathbf{u})$ be a feasible guess with $n' \geq 1$, and let $0 \leq b' \leq \min\{b, n'\}$. For every $M' \in \mathcal{R}_{b'}(M_b)$ and every affine function $f \in \mathcal{A}_{b'}$,*

$$\begin{aligned} &\left|\Pr\left[M'(Z_{b'}^{(R,\mathbf{u})}) = f(Z_{b'}^{(R,\mathbf{u})})\right] - q(M', f)\right| \\ &\quad \leq \frac{b'(b'-1)}{2n'} + b'\left|\frac{w'}{n'} - \frac{1}{2}\right|. \end{aligned} \tag{8}$$

Proof. Let $(z_0, \ldots, z_{b'-1})$ denote the random vector $Z_{b'}^{(R,\mathbf{u})}$. Suppose the first i bits contain h ones, where $0 \leq i < b'$, with an empty conditioning sequence when $i = 0$. Since the positions are sampled without replacement, the conditional probability that the next bit is 1 is

$$p_i \triangleq \Pr[z_i = 1 \mid z_0, \ldots, z_{i-1}] = \frac{w' - h}{n' - i}.$$

Since $w' = h + (n' - i)p_i$, we have

$$\left|p_i - \frac{w'}{n'}\right| = \frac{|ip_i - h|}{n'} \leq \frac{i}{n'}.$$

Consequently,

$$\left|p_i - \frac{1}{2}\right| \leq \frac{i}{n'} + \left|\frac{w'}{n'} - \frac{1}{2}\right|.$$

For independent uniform bits, the conditional probability that each bit equals 1, given the preceding bits, is $1/2$.

We compare the two input distributions by replacing one bit at a time. For $0 \leq i \leq b'$, let a_i denote the agreement probability between M' and f when the first i inputs follow the fixed-weight marginal distribution and the remaining $b' - i$ inputs are independent and uniform. Thus,

$$a_0 = q(M', f), \qquad a_{b'} = \Pr\left[M'(Z_{b'}^{(R,\mathbf{u})}) = f(Z_{b'}^{(R,\mathbf{u})})\right].$$

For each $0 \leq i < b'$, the distributions defining a_i and a_{i+1} differ only in the conditional distribution of the next bit. Given the first i bits, its probability of being 1 changes from $1/2$ to p_i. Hence,

$$|a_{i+1} - a_i| \leq \frac{i}{n'} + \left|\frac{w'}{n'} - \frac{1}{2}\right|.$$

By the triangle inequality,

$$\begin{aligned} |a_{b'} - a_0| &\leq \sum_{i=0}^{b'-1} |a_{i+1} - a_i| \\ &\leq \sum_{i=0}^{b'-1} \left(\frac{i}{n'} + \left|\frac{w'}{n'} - \frac{1}{2}\right|\right) \\ &= \frac{b'(b'-1)}{2n'} + b'\left|\frac{w'}{n'} - \frac{1}{2}\right|. \end{aligned}$$

This proves (8). □

Lemma 2 (Uniform-to-fixed-weight transfer)**.** *Let $(R, \mathbf{u})$ be a feasible guess and $0 \le b' \le \min\{b, n'\}$, and set $\ell \triangleq \min\{w', n' - w'\}$. For $0 \le j \le \min\{b', \ell\}$, define*

$$p_j(n', w', b') \triangleq \frac{2^j \binom{\ell}{j}}{\binom{n'}{b'}} \sum_{h=0}^{\lfloor (b'-j)/2 \rfloor} \binom{\ell - j}{h} \cdot \binom{n' - 2\ell}{b' - j - 2h},$$

and set $p_j(n', w', b') = 0$ for $\ell < j \le b'$. Define

$$\Phi(n', w', b') \triangleq \sum_{j=0}^{b'} p_j(n', w', b') \operatorname{adv}_j(M_b). \tag{9}$$

Then, for every $M' \in \mathcal{R}_{b'}(M_b)$ and every affine function $f \in \mathcal{A}_{b'}$,

$$\Pr\left[M'(Z_{b'}^{(R,\mathbf{u})}) = f(Z_{b'}^{(R,\mathbf{u})})\right] \le \frac{1}{2} + \Phi(n', w', b'). \tag{10}$$

Consequently, for every feasible overlap class r,

$$q_r(R, \mathbf{u}) \le \frac{1}{2} + \Phi(n', w', b - r). \tag{11}$$

Proof. We first construct a uniformly random n'-bit vector of Hamming weight w'. Choose ℓ disjoint pairs uniformly at random among the n' positions. Independently assign each pair either $(0, 1)$ or $(1, 0)$ with equal probability. Set all unpaired positions to 0 if $w' \le n'/2$, and to 1 otherwise. The resulting vector has Hamming weight w' and is uniformly distributed over all such vectors by permutation symmetry.

Let $J \subseteq [n']$ denote a fixed set of b' sampled input positions. Let κ denote the number of pairs containing exactly one position in J. Condition on the chosen pairs and the assignments to all pairs entirely contained in J. Then only the κ positions paired with positions outside J remain random. These κ bits are independent and uniform, while all other sampled inputs are fixed.

Under this conditioning, M' becomes a residual belonging to $\mathcal{R}_\kappa(M_b)$, and f remains affine on the κ remaining inputs. By the definition of bit-fixing affine advantage, their conditional agreement probability is at most $1/2 + \operatorname{adv}_\kappa(M_b)$. Averaging over the random pairing and assignments gives

$$\Pr\left[M'(Z_{b'}^{(R,\mathbf{u})}) = f(Z_{b'}^{(R,\mathbf{u})})\right] \le \frac{1}{2} + \mathbb{E}[\operatorname{adv}_\kappa(M_b)].$$

To compute the distribution of κ, we may fix the ℓ pairs and instead choose J uniformly among all b'-element subsets of $[n']$. For $\kappa = j$, exactly one position must be selected from each of j pairs, giving $\binom{\ell}{j} 2^j$ choices. Let h denote the number of additional pairs from which both positions are selected. There are $\binom{\ell - j}{h}$ ways to choose these pairs. The remaining $b' - j - 2h$ positions are selected from the $n' - 2\ell$ unpaired positions, giving $\binom{n' - 2\ell}{b' - j - 2h}$ choices.

Summing over h and dividing by $\binom{n'}{b'}$ gives

$$\Pr[\kappa = j] = p_j(n', w', b').$$

Therefore,

$$\mathbb{E}[\mathrm{adv}_\kappa(M_b)] = \Phi(n', w', b'),$$

which proves (10).

The bound holds for every residual in $\mathcal{R}_{b-r}(M_b)$ and every affine function in $\mathcal{A}_{b-r}$. Taking their maximum as in (3) proves (11). □

A.2 Finite-Parameter Cost Bounds

For each feasible pair (n', w') induced by a global guess, write

$$\mu_r = m\frac{\binom{n-n'}{r}\binom{n'}{b-r}}{\binom{n}{b}}.$$

Define the guessing exponent as

$$g(n', w') \triangleq \log_2 \frac{\binom{n}{w}}{\binom{n'}{w'}}.$$

For each feasible overlap class r, define

$$\eta_r(n', w') \triangleq \log_2\left(1 + 2\Phi(n', w', b-r)\right).$$

Define an upper bound on the total benefit as

$$v(n', w') \triangleq \max_{\substack{0\le k_r\le \mu_r \\ \sum_r k_r \le n'}} \sum_r k_r \eta_r(n', w'). \tag{12}$$

All sums and optimizations over overlap classes include only feasible classes. Since $0 \le \Phi(n', w', b-r) \le 1/2$, we have $0 \le \eta_r(n', w') \le 1$ and $0 \le v(n', w') \le n'$.

Lemma 3. *The optimal cost satisfies*

$$\begin{aligned}\log_2 C^\star \ge \min_{(n',w') \text{ feasible}} \Bigg\{ & \log_2 \frac{\binom{n}{w}}{\binom{n'}{w'}} + n' \\ & - \max_{\substack{0\le k_r\le \mu_r \\ \sum_r k_r \le n'}} \sum_r k_r \log_2\left(1 + 2\Phi(n', w', b-r)\right) \Bigg\}.\end{aligned} \tag{13}$$

Equivalently,

$$\log_2 C^\star \ge \min_{(n',w') \text{ feasible}} \left\{ g(n', w') + n' - v(n', w') \right\}.$$

Proof. Fix feasible guessing parameters (n', w') and equation counts $\mathbf{k}$. By the definition of the cost,

$$\begin{aligned}\log_2 C(R, \mathbf{u}, \mathbf{k}) &= g(n', w') + n' - k - \sum_r k_r \log_2 q_r(R, \mathbf{u}) \\ &= g(n', w') + n' - \sum_r k_r \log_2(2q_r(R, \mathbf{u})).\end{aligned}$$

By Lemma 2,

$$q_r(R, \mathbf{u}) \le \frac{1}{2} + \Phi(n', w', b-r).$$

Therefore,

$$\begin{aligned}\log_2 C(R,\mathbf{u},\mathbf{k}) &\geq g(n',w') + n' - \sum_r k_r \eta_r(n',w') \\ &\geq g(n',w') + n' - v(n',w').\end{aligned}$$

Minimizing over all feasible global guesses and real-valued equation counts proves the claim. □

Lemma 4. *Fix any constants $\gamma > 0$ and $\theta \in (0,1]$, and let $b' = \theta b$. Then*

$$\log_2 C^\star(n,w;\gamma) \leq \left[1 - \theta \log_2\left(1 + 2\,\mathrm{adv}_{b'}(M_b)\right) + o(1)\right] n,$$

uniformly over $w \in \mathcal{W}_n$.

Proof. Let $r_0 = b - b'$. To obtain sufficiently many equations with exactly r_0 fixed nonlinear inputs, we choose t such that t/n is close to r_0/b. Specifically, let t be the even integer nearest to $(r_0/b)n$, and choose a global guess $(R,\mathbf{u})$ satisfying $|R| = t$ and $\|\mathbf{u}\| = t/2$. Then

$$n' = n - t, \qquad w' = w - t/2.$$

Consequently,

$$\frac{t}{n} = \frac{r_0}{b} + O(1/n), \qquad \frac{n'}{n} = \frac{b'}{b} + O(1/n) \longrightarrow \theta, \qquad w' - \frac{n'}{2} = w - \frac{n}{2}.$$

Since $n'/n \to \theta > 0$ and $|w' - n'/2| \leq \sqrt{n \ln n} = o(n')$, the guess is feasible for all $w \in \mathcal{W}_n$ when n is sufficiently large.

We next show that there are sufficiently many equations with exactly r_0 fixed nonlinear inputs. Recall that

$$\mu_r = m \frac{\binom{t}{r}\binom{n'}{b-r}}{\binom{n}{b}}.$$

For consecutive classes with positive expected counts,

$$\frac{\mu_{r+1}}{\mu_r} = \frac{(t-r)(b-r)}{(r+1)(n'-b+r+1)}.$$

This ratio is at least 1 if and only if

$$r + 1 \leq \frac{(b+1)(t+1)}{n+2}.$$

By our choice of t,

$$\frac{(b+1)(t+1)}{n+2} = r_0 + \frac{r_0}{b} + O(b/n).$$

For $1 \leq r_0 < b$, this quantity lies strictly between r_0 and $r_0 + 1$ for sufficiently large n. Hence, μ_{r_0} is the largest expected class count. When $r_0 = 0$, we have $\mu_0 = m$, so the same conclusion holds. Since $\sum_r \mu_r = m$ and there are at most $b+1$ classes,

$$\mu_{r_0} \geq \frac{m}{b+1}. \tag{14}$$

Since $m = n^s$ with $s > 1$ and $b = O(\log n)$, we have $\mu_{r_0}/n \to \infty$.

By Lemma 1, for every $M' \in \mathcal{R}_{b'}(M_b)$ and every affine function $f \in \mathcal{A}_{b'}$,

$$\left|\Pr\left[M'(Z_{b'}^{(R,\mathbf{u})}) = f(Z_{b'}^{(R,\mathbf{u})})\right] - q(M', f)\right| \le \xi_n,$$

where

$$\xi_n \triangleq \frac{b'(b'-1)}{2n'} + b'\left|\frac{w'}{n'} - \frac{1}{2}\right|.$$

Since $b' = O(\log n)$, $n' = \Theta(n)$, and $|w' - n'/2| \le \sqrt{n \ln n}$, we have $\xi_n = o(1)$ uniformly over $w \in \mathcal{W}_n$.

By the definition of bit-fixing affine advantage, there exist $M' \in \mathcal{R}_{b'}(M_b)$ and $f \in \mathcal{A}_{b'}$ attaining agreement $1/2+\mathrm{adv}_{b'}(M_b)$ under uniform inputs. Since the maximization defining $q_{r_0}(R, \mathbf{u})$ includes all such residuals and affine functions,

$$q_{r_0}(R, \mathbf{u}) \ge \frac{1}{2} + \mathrm{adv}_{b'}(M_b) - \xi_n.$$

Stirling's formula gives $g(n', w') = t + O(\log n)$ uniformly over $w \in \mathcal{W}_n$. By (14), we may select $k_{r_0} = n'$ and $k_r = 0$ for $r \ne r_0$ in the cost model. Therefore,

$$\begin{aligned}\log_2 C^\star(n, w; \gamma) &\le g(n', w') + n' \log_2 q_{r_0}(R, \mathbf{u})^{-1} \\ &= g(n', w') + n' - n' \log_2(2q_{r_0}(R, \mathbf{u})) \\ &\le n - n' \log_2 (1 + 2\,\mathrm{adv}_{b'}(M_b)) + o(n).\end{aligned}$$

Since $n'/n = \theta + o(1)$ and $0 \le \log_2(1 + 2\,\mathrm{adv}_{b'}(M_b)) \le 1$, we conclude that

$$\log_2 C^\star(n, w; \gamma) \le [1 - \theta \log_2 (1 + 2\,\mathrm{adv}_{b'}(M_b)) + o(1)]\, n.$$

□

A.3 Proofs of Theorem 1

Proof: Sufficiency. Assume that $\mathrm{adv}_{\theta b}(M_b) \to 0$ for every fixed $\theta \in (0, 1]$. Fix any $\gamma > 0$ and any $c \in (0, 2^{-(s-1)/\gamma})$. Equivalently,

$$s + \gamma \log_2 c < 1.$$

All estimates below are uniform over $w \in \mathcal{W}_n$.

We first consider feasible guessing parameters (n', w') satisfying $g(n', w') \le cn$. For $w \in \mathcal{W}_n$, Stirling's formula gives

$$\log_2 \binom{n}{w} = n - O(\log n).$$

Since $\binom{n'}{w'} \le 2^{n'}$, we have $g(n', w') \ge (n - n') - O(\log n)$ and $g(n', w') + n' \ge n - O(\log n)$. Consequently,

$$\begin{aligned} n - n' &\le cn + O(\log n), \\ n' &\ge (1-c)n - O(\log n) = \Omega(n), \\ g(n', w') + n' &\ge n - O(\log n). \end{aligned} \tag{15}$$

We also need to bound the remaining seed weight w'. Let $h_2(x) \triangleq -x \log_2 x - (1-x)\log_2(1-x)$ denote the binary entropy function, with $h_2(0) = h_2(1) = 0$. Using $g(n', w') \le cn$, we obtain

$$(1-c)n - O(\log n) \le \log_2 \binom{n'}{w'} \le n' h_2(w'/n') \le n h_2(w'/n').$$

Choose a constant $\eta \in (0, 1/2)$ such that $h_2(\eta) < 1 - c$. For sufficiently large n, this implies

$$\eta \le w'/n' \le 1 - \eta.$$

We now show that $v(n', w') = o(n)$. To this end, we separate equations according to the number $b' = b - r$ of remaining nonlinear inputs.

Since $s + \gamma \log_2 c < 1$, we can choose a sufficiently small constant $\lambda \in (0, 1/2)$ such that

$$s + \gamma \left[h_2(\lambda) + (1 - \lambda) \log_2 c\right] < 1.$$

Set $d_0 \triangleq \lambda b$.

The total expected number of equations with at most d_0 remaining nonlinear inputs satisfies

$$\begin{aligned}
\mu_{\le d_0} &\triangleq \sum_{b-r \le d_0} \mu_r \\
&\le m \binom{b}{d_0} \left(\frac{n - n'}{n}\right)^{b - d_0} \\
&\le n^{s + \gamma[h_2(\lambda) + (1-\lambda)\log_2 c] + o(1)} \\
&= o(n).
\end{aligned}$$

Next, consider equations with $b' > d_0$. We show that their fixed-weight agreement probabilities are at most $1/2 + o(1)$.

Recall from Lemma 2 that

$$\Phi(n', w', b') = \mathbb{E}[\mathrm{adv}_\kappa(M_b)],$$

where κ counts the pairs containing exactly one of the b' selected positions. Let

$$\ell \triangleq \min\{w', n' - w'\}, \qquad \vartheta \triangleq \frac{2\ell}{n'} \ge 2\eta.$$

To estimate κ, fix the pairs and consider selecting b' positions independently with replacement. Let h count the selections that fall in paired positions. Then

$$h \sim \mathrm{Bin}(b', \vartheta).$$

If no two selections belong to the same pair or repeat an unpaired position, the number of pairs selected exactly once equals h. For any two draws, the probability that they belong to the same pair or repeat an unpaired position is at most $2/n'$. Coupling the with-replacement and without-replacement samples and applying a union bound gives

$$\left|\Phi(n', w', b') - \mathbb{E}[\mathrm{adv}_h(M_b)]\right| \le \frac{b'(b' - 1)}{n'}. \tag{16}$$

Since $\mathbb{E}h = \vartheta b' \ge 2\eta b'$, the Chernoff bound gives

$$\Pr[h < \eta b'] \le e^{-\eta b'/4}.$$

Let $u_b \triangleq \eta\lambda b$. Since $b' > d_0$, we have $b' > \lambda b$, so $h \ge \eta b'$ implies $h \ge u_b$. By the monotonicity of bit-fixing affine advantage,

$$\mathbb{E}[\mathrm{adv}_h(M_b)] \le \mathrm{adv}_{u_b}(M_b) + \frac{1}{2} e^{-\eta b'/4}.$$

Here monotonicity follows because further fixing inputs preserves affine approximations, and at least one assignment achieves agreement probability no smaller than their average.

Combining this with (16) gives

$$\sup_{d_0<b'\le b} \Phi(n', w', b') \le \operatorname{adv}_{\eta\lambda b}(M_b) + \frac{1}{2}e^{-\eta\lambda b/4} + \frac{b(b-1)}{n'} = o(1).$$

The first term tends to zero by the assumed advantage condition, while the remaining terms vanish since $b = \Theta(\log n)$ and $n' = \Omega(n)$.

We can now bound the total benefit from the selected equations. For classes with $b' \le d_0$, we use $\eta_r(n', w') \le 1$. For classes with $b' > d_0$, we use

$$\eta_r(n', w') \le \frac{2}{\ln 2}\Phi(n', w', b - r).$$

Since at most n' equations can be selected,

$$\begin{aligned} v(n', w') &\le \mu_{\le d_0} + \frac{2n'}{\ln 2} \sup_{d_0<b'\le b} \Phi(n', w', b') \\ &= o(n). \end{aligned}$$

Combining this with (15), we obtain

$$g(n', w') + n' - v(n', w') \ge n - o(n) > cn$$

for sufficiently large n.

For guessing parameters with $g(n', w') > cn$, we instead use $v(n', w') \le n'$ to obtain

$$g(n', w') + n' - v(n', w') \ge g(n', w') > cn.$$

Thus, taking the minimum over all feasible guessing parameters and applying Lemma 3 gives

$$\log_2 C^\star(n, w; \gamma) \ge \min_{(n', w') \text{ feasible}} \left\{ g(n', w') + n' - v(n', w') \right\} > cn.$$

Therefore $C^\star(n, w; \gamma) > 2^{cn}$, proving sufficiency.

Necessity. Suppose the advantage condition fails. Then there exist constants $\theta \in (0, 1]$ and $\delta > 0$, and a sequence $b_i \to \infty$ such that

$$\operatorname{adv}_{\theta b_i}(M_{b_i}) \ge \delta.$$

Choose a constant c_0 satisfying

$$1 - \theta \log_2(1 + 2\delta) < c_0 < 1.$$

Choose $\gamma > 0$ sufficiently large that

$$c_0 < 2^{-(s-1)/\gamma}.$$

Choose n_i to be the even integer nearest to $2^{(b_i - 1/2)/\gamma}$. For sufficiently large i,

$$b_i = \gamma \log_2 n_i.$$

Taking $w_i = n_i/2$, Lemma 4 gives

$$\log_2 C^\star(n_i, w_i; \gamma) \le [1 - \theta \log_2(1 + 2\delta) + o(1)]\, n_i.$$

Consequently, infinitely many n_i satisfy

$$C^\star(n_i, w_i; \gamma) < 2^{c_0 n_i},$$

where $c_0 < 2^{-(s-1)/\gamma}$ and $w_i \in \mathcal{W}_{n_i}$. This contradicts the asserted cost property for the chosen γ and c_0, completing the proof. □

A.4 Proof of Corollary 1

Proof. Fix $\gamma > 0$ and an arbitrary predicate family $(M_b)_{b \ge 1}$. Let

$$\rho_n \triangleq \left(\frac{n}{m}\right)^{1/b}.$$

Since $m = n^s$ and $b = \gamma \log_2 n$, we have

$$\rho_n = 2^{-(s-1)/\gamma} + o(1).$$

Choose a global guess $(R, \mathbf{u})$ with $|R| = \rho_n n + O(1)$ and $\|\mathbf{u}\| = |R|/2 + O(1)$. For every $w \in \mathcal{W}_n$, Stirling's approximation gives

$$\begin{aligned} p_{\text{guess}}(R, \mathbf{u}) &= \frac{\binom{n-|R|}{w-\|\mathbf{u}\|}}{\binom{n}{w}} \\ &= 2^{-|R| + O(\log n)}, \end{aligned}$$

uniformly over $w \in \mathcal{W}_n$.

Consider overlap class $r = b$, in which all nonlinear predicate inputs are fixed. Its expected equation capacity satisfies

$$\begin{aligned} \mu_b(R) &= m \frac{\binom{|R|}{b}}{\binom{n}{b}} \\ &= m \left(\frac{|R|}{n}\right)^b (1 + o(1)) \\ &= n(1 + o(1)). \end{aligned}$$

Since ρ_n approaches a positive constant, we have $\mu_b(R) \ge n - |R|$ for all sufficiently large n.

Every residual in this class is constant and therefore has exact affine agreement $q_b(R, \mathbf{u}) = 1$. We may consequently select $k_b = n - |R|$ and $k_r = 0$ for all $r \ne b$, leaving no free variables. The cost of this configuration is

$$\begin{aligned} C^\star(n, w; \gamma) &\le \frac{1}{p_{\text{guess}}(R, \mathbf{u})} \\ &= 2^{|R| + O(\log n)} \\ &= 2^{(2^{-(s-1)/\gamma} + o(1))n}. \end{aligned}$$

The bound holds uniformly over all $w \in \mathcal{W}_n$ and is independent of the choice of M_b. □

A.5 Proof of Theorem 2

Proof: Sufficiency. Assume that $\mathrm{adv}_{\theta b}(M_b) \to 1/2$ for every fixed $\theta \in (0,1)$. For any fixed $\gamma > 0$ and $\theta \in (0,1)$, Lemma 4 gives

$$\log_2 C^\star(n, w; \gamma) \le (1 - \theta + o(1))n$$

uniformly over $w \in \mathcal{W}_n$. Since this holds for every fixed $\theta < 1$, letting $\theta \uparrow 1$ gives

$$\sup_{w \in \mathcal{W}_n} \frac{\log_2 C^\star(n, w; \gamma)}{n} \longrightarrow 0.$$

Here $\log_2 C^\star \ge 0$ follows from the cost definition. Thus $C^\star(n, w; \gamma) = 2^{o(n)}$ uniformly over $w \in \mathcal{W}_n$, proving sufficiency.

Necessity. Suppose the advantage condition fails. Then there exist constants $\theta \in (0,1)$ and $\delta > 0$, and a sequence $b_i \to \infty$ such that

$$\mathrm{adv}_{\theta b_i}(M_{b_i}) \le \frac{1}{2} - \delta. \tag{17}$$

Fix any $\gamma > 0$, and choose n_i to be the even integer nearest to $2^{(b_i - 1/2)/\gamma}$. Then $b_i = \gamma \log_2 n_i$ for all sufficiently large i. Set $w_i = n_i/2$.

Assume, for contradiction, that the cost is subexponential uniformly over typical seed weights. Along this sequence, choose feasible parameters $(R, \mathbf{u}, \mathbf{k})$ attaining $C^\star$. We then have $\log_2 C(R, \mathbf{u}, \mathbf{k}) = o(n)$. Suppress the index i and recall that $k = \sum_r k_r$. Write $t = |R|$, $n' = n - t$, and $w' = w - \|\mathbf{u}\|$.

By the definition of the cost,

$$\begin{aligned} \log_2 C(R, \mathbf{u}, \mathbf{k}) = g(n', w') + (n' - k) \\ + \sum_r k_r \log_2 q_r(R, \mathbf{u})^{-1}. \end{aligned} \tag{18}$$

All three terms on the right are nonnegative. Hence $g(n', w') = o(n)$ and $n' - k = o(n)$. Since $g(n', w') \ge t - O(\log n)$ for $w = n/2$, we obtain

$$t = o(n), \qquad n' = n - o(n), \qquad k = n - o(n). \tag{19}$$

We first show that equations with many fixed nonlinear inputs are too few. Set

$$\lambda \triangleq \frac{1 - \theta}{2}, \qquad r_0 \triangleq \lambda b.$$

By the union bound,

$$\begin{aligned} \sum_{r \ge r_0} \mu_r &\le m \binom{b}{r_0} (t/n)^{r_0} \\ &\le m 2^b (t/n)^{\lambda b} \\ &= o(n). \end{aligned} \tag{20}$$

The last equality follows from $t/n = o(1)$, $m = n^s$, and $b = \Theta(\log n)$.

We next bound $q_r(R, \mathbf{u})$ for the remaining overlap classes. For $r < r_0$, let

$$b' \triangleq b - r > (1 - \lambda)b, \qquad b_0 \triangleq \theta b.$$

In the random pairing of Lemma 2, let κ count the pairs containing exactly one of the b' selected positions.

A selected position is not counted in κ only if it is unpaired or its partner is also selected. Since $w = n/2$, we have

$$|n' - 2w'| = |2\|\mathbf{u}\| - t| \le t.$$

It follows that

$$\mathbb{E}[b' - \kappa] \le \frac{b't}{n'} + \frac{b'(b'-1)}{n'-1} = o(b),$$

uniformly over $r < r_0$. Since $b' - b_0 \ge \lambda b - 1$, Markov's inequality gives

$$\Pr[\kappa < b_0] = o(1).$$

By the monotonicity of bit-fixing affine advantage and (17),

$$\begin{aligned}\Phi(n', w', b') &= \mathbb{E}[\mathrm{adv}_\kappa(M_b)] \\ &\le \mathrm{adv}_{b_0}(M_b) + \frac{1}{2}\Pr[\kappa < b_0] \\ &\le \frac{1}{2} - \delta + o(1).\end{aligned}$$

Therefore, by Lemma 2, uniformly over $r < r_0$,

$$q_r(R, \mathbf{u}) \le 1 - \delta + o(1) \le 1 - \frac{\delta}{2}$$

for sufficiently large n.

By (19) and (20),

$$\sum_{r<r_0} k_r \ge k - \sum_{r \ge r_0} \mu_r = n - o(n).$$

Thus, (18) gives

$$\begin{aligned}\log_2 C(R, \mathbf{u}, \mathbf{k}) &\ge \sum_{r<r_0} k_r \log_2 q_r(R, \mathbf{u})^{-1} \\ &\ge (n - o(n)) \log_2 \left(\frac{1}{1 - \delta/2}\right) \\ &= \Omega(n),\end{aligned}$$

contradicting $\log_2 C(R, \mathbf{u}, \mathbf{k}) = o(n)$.

This proves necessity. Since the argument holds for every fixed $\gamma > 0$, the claimed characterization follows. □

A.6 Proof of Corollary 2

Proof. Fix $\gamma > 0$ and an integer r satisfying the hypothesis, and choose a global guess $(R, \mathbf{u})$ with $|R| = r$. For every $w \in \mathcal{W}_n$, its success probability satisfies

$$p_{\mathrm{guess}}(R, \mathbf{u}) = \frac{(w)_{\|\mathbf{u}\|}(n-w)_{r-\|\mathbf{u}\|}}{(n)_r} = 2^{-r}(1 + o(1)),$$

where $(a)_j$ denotes the falling factorial.

The expected capacity of overlap class r is

$$\begin{aligned}\mu_r(R) &= m\frac{\binom{n-r}{b-r}}{\binom{n}{b}} \\ &= m\frac{(b)_r}{(n)_r} = \Theta\left(n^{s-r}(\log n)^r\right).\end{aligned}$$

Since $r \le s-1$, we have $\mu_r(R) \ge n-r$ for all sufficiently large n.

Let $n' = n-r$, $w' = w - \|\mathbf{u}\|$, and $d = b-r$. For every $\mathbf{z} \in \{0,1\}^d$, the fixed-weight distribution satisfies

$$\begin{aligned}\Pr[Z_d^{(R,\mathbf{u})} = \mathbf{z}] &= \frac{(w')_{\|\mathbf{z}\|}(n'-w')_{d-\|\mathbf{z}\|}}{(n')_d} \\ &= 2^{-d}\left(1 + O\left(d\sqrt{\frac{\ln n}{n}} + \frac{d^2}{n}\right)\right) \\ &= 2^{-d}(1+o(1)).\end{aligned}$$

The approximation is uniform over $\mathbf{z}$, $\mathbf{u}$, and $w \in \mathcal{W}_n$, since r is fixed and $d = O(\log n)$.

Write

$$\varepsilon_b \triangleq \frac{1}{2} - \mathrm{adv}_{b-r}(M_b) = O\left(\frac{\log n}{n}\right).$$

By the definition of bit-fixing affine advantage, there exist a residual $M' \in \mathcal{R}_d(M_b)$ and an affine function $f \in \mathcal{A}_d$ whose disagreement probability under uniform inputs is ε_b. The multiplicative distribution bound therefore gives

$$q_r(R,\mathbf{u}) \ge 1 - (1+o(1))\varepsilon_b \ge 1 - K\frac{\log n}{n}$$

for some constant $K > 0$ and all sufficiently large n.

We may thus choose $k_r = n-r$ and $k_j = 0$ for $j \ne r$, satisfying the capacity constraints. Substituting this configuration into the cost formula yields

$$\begin{aligned}C^\star(n,w;\gamma) &\le \frac{1}{p_{\mathrm{guess}}(R,\mathbf{u})q_r(R,\mathbf{u})^{n-r}} \\ &\le 2^{r+o(1)}\left(1 - K\frac{\log n}{n}\right)^{-(n-r)} \\ &= \exp(O(\log n)) = n^{O(1)}.\end{aligned}$$

All bounds hold uniformly over $w \in \mathcal{W}_n$, completing the proof. □

A.7 Proof of Corollary 3

Proof. Fix any constant $\theta \in (3/4, 1)$, and let $b' = \lceil \theta b \rceil$ and $r_0 = b - b'$. Choose t to be the even integer nearest to $2(1-\theta)n$, fix a set $R \subseteq [n]$ of size t, and write $n' = n - t$. Then

$$\frac{t}{n} = 2(1-\theta) + O(1/n), \qquad \frac{n'}{n} \longrightarrow 2\theta - 1 > 0.$$

We first show that there are sufficiently many equations with at least r_0 fixed nonlinear inputs. Let $\mathcal{G}$ denote the set of all such feasible overlap classes, and put

$$\mu_{\mathcal{G}} = \sum_{r \in \mathcal{G}} \mu_r(R).$$

The hypergeometric overlap has mean $(2(1-\theta)+o(1))b$ and variance at most b. Since $r_0 = (1-\theta)b + O(1)$ and $b \to \infty$, Chebyshev's inequality gives

$$\mu_{\mathcal{G}} = m\big(1 - O(1/b)\big) \gg n.$$

We may therefore select

$$k_r = \frac{n'\mu_r(R)}{\mu_{\mathcal{G}}} \quad \text{for } r \in \mathcal{G}, \qquad k_r = 0 \quad \text{otherwise.}$$

These counts satisfy $0 \le k_r \le \mu_r(R)$ and $\sum_r k_r = n'$.

Write

$$\varepsilon_b = \frac{1}{2} - \mathrm{adv}_{b'}(M_b) \longrightarrow 0.$$

For every $r \in \mathcal{G}$, we have $b - r \le b'$. By monotonicity of bit-fixing affine advantage, there exist a residual M'_r and an affine function f_r satisfying

$$q(M'_r, f_r) \ge 1 - \varepsilon_b.$$

For $h \in [n'+1]$, let $e_r(h)$ denote their disagreement probability under the $(b-r)$-coordinate marginal of a uniformly random n'-bit vector of weight h, and define

$$E(h) = \frac{1}{n'} \sum_{r \in \mathcal{G}} k_r e_r(h).$$

If $H \sim \mathrm{Bin}(n', 1/2)$, averaging the fixed-weight marginals over H gives uniform inputs. Consequently,

$$\mathbb{E}[E(H)] = \frac{1}{n'} \sum_{r \in \mathcal{G}} k_r \big(1 - q(M'_r, f_r)\big) \le \varepsilon_b.$$

Since $\Pr[H \in W_{n'}] \ge 1 - 2/(n')^2$, there exists a common weight $h \in W_{n'}$ such that

$$E(h) \le \frac{\varepsilon_b}{\Pr[H \in W_{n'}]} = (1 + o(1))\varepsilon_b = o(1).$$

For every $w \in W_n$, choose a global guess (R, u) with $\|u\| = w - h$. This guess is feasible for all sufficiently large n, because

$$w - h = \frac{t}{2} + O(\sqrt{n \ln n}) \in [0, t].$$

Thus $w' = h$. Stirling's formula gives

$$g(n', w') = \log_2 \binom{n}{w} - \log_2 \binom{n'}{w'} = t + O(\log n),$$

uniformly over $w \in W_n$.

By the maximization defining $q_r(R, u)$,

$$q_r(R, u) \ge 1 - e_r(h).$$

Moreover, $q_r(R,u) \geq 1/2$, since an affine function and its complement are both allowed. Using $-\log_2 q \leq 2(1-q)$ for $q \in [1/2,1]$, we obtain

$$\begin{aligned}\log_2 C^\star(n,w) &\leq g(n',w') + \sum_{r\in\mathcal{G}} k_r \log_2 q_r(R,u)^{-1} \\ &\leq t + O(\log n) + 2\sum_{r\in\mathcal{G}} k_r e_r(h) \\ &= t + O(\log n) + 2n'E(h) \\ &= 2(1-\theta)n + o(n).\end{aligned}$$

All estimates are uniform over $w \in W_n$. Since this holds for every fixed $\theta \in (3/4,1)$, letting $\theta \uparrow 1$ gives

$$\sup_{w\in W_n} \frac{\log_2 C^\star(n,w)}{n} \longrightarrow 0,$$

where nonnegativity follows from $C^\star \geq 1$. Therefore, $C^\star(n,w) = 2^{o(n)}$ uniformly over $w \in W_n$.

□

B Proofs of the MOD12 Results

B.1 Bit-Fixing Affine Advantage

We establish upper bounds on the bit-fixing affine advantage of modular predicates by analyzing the affine agreement of their residual functions under uniform inputs.

For $t \geq 2$ and $S_0 \subseteq \{0,\ldots,t-1\}$, let $g = g_{t,S_0}$ denote the signed weight profile defined by

$$g(w) = \begin{cases} 1, & \text{if } w \bmod t \in S_0, \\ -1, & \text{otherwise.}\end{cases}$$

Recall that its Fourier coefficients are

$$\widehat{g}_\nu = \frac{1}{t}\sum_{w=0}^{t-1} g(w)\omega^{-\nu w}, \qquad \omega = \exp(2\pi \mathrm{i}/t), \qquad 0 \leq \nu < t.$$

Throughout this subsection, agreement probabilities are evaluated under independent uniform inputs.

Lemma 5 (Fourier bound for modular bit-fixing affine advantage)**.** *Let $t \geq 2$ and $S_0 \subseteq \{0,\ldots,t-1\}$ satisfy $4 \mid t$ and*

$$|S_0 \cap 2\mathbb{Z}| = |S_0 \cap (2\mathbb{Z}+1)| = \frac{t}{4}. \tag{21}$$

For $b \geq 1$ and $0 \leq r < b$, let M' be any residual of $M_b = \mathsf{MOD}_b^{t,S_0}$ obtained by fixing r inputs, and let $b' = b - r$. For any affine function of the form

$$f(\mathbf{z}) = c \oplus \bigoplus_{j\in A} z_j, \qquad A \subseteq [b'], \quad c \in \mathbb{F}_2,$$

we have

$$\left| q(M',f) - \frac{1}{2} \right| \leq \frac{1}{2} \sum_{\substack{0\leq\nu<t \\ \nu\neq 0,t/2}} |\widehat{g}_\nu| \left|\sin\frac{\pi\nu}{t}\right|^{|A|} \left|\cos\frac{\pi\nu}{t}\right|^{b'-|A|}. \tag{22}$$

Consequently, for every $1 \le b' \le b$,

$$\mathrm{adv}_{b'}(M_b) \le \min\left\{\frac{1}{2}, \frac{\sqrt{t-2}}{2}\left(\cos\frac{\pi}{t}\right)^{b'}\right\}. \tag{23}$$

Proof. Write $M' = M_{b|T,\mathbf{v}}$ with $|T| = r$; the number of fixed inputs equal to one is $\|\mathbf{v}\|$. For the remaining inputs,

$$M'(\mathbf{z}) = 0 \quad \Longleftrightarrow \quad (\|\mathbf{v}\| + \|\mathbf{z}\|) \bmod t \in S_0.$$

Hence,

$$(-1)^{M'(\mathbf{z})} = g(\|\mathbf{v}\| + \|\mathbf{z}\|).$$

By condition (21),

$$\widehat{g}_0 = \frac{1}{t}\sum_{w=0}^{t-1} g(w) = 0, \qquad \widehat{g}_{t/2} = \frac{1}{t}\sum_{w=0}^{t-1}(-1)^w g(w) = 0.$$

Therefore,

$$g(w) = \sum_{\substack{0\le\nu<t\\ \nu\neq 0,t/2}} \widehat{g}_\nu \omega^{\nu w}.$$

Let $\mathbf{Z} = (Z_0, \ldots, Z_{b'-1})$ be uniformly distributed over $\{0,1\}^{b'}$. By the definitions of g and f,

$$\begin{aligned} 2q(M', f) - 1 &= \mathbb{E}\left[(-1)^{M'(\mathbf{Z})\oplus f(\mathbf{Z})}\right] \\ &= (-1)^c \mathbb{E}\left[g\left(\|\mathbf{v}\| + \sum_{j\in[b']} Z_j\right)(-1)^{\sum_{j\in A} Z_j}\right]. \end{aligned}$$

Since each Z_j is a uniform bit,

$$\mathbb{E}[\omega^{\nu Z_j}] = \frac{1+\omega^\nu}{2}, \qquad \mathbb{E}[(-1)^{Z_j}\omega^{\nu Z_j}] = \frac{1-\omega^\nu}{2}.$$

By independence of the inputs, substituting the Fourier expansion of g gives

$$2q(M', f) - 1 = (-1)^c \sum_{\substack{0\le\nu<t\\ \nu\neq 0,t/2}} \widehat{g}_\nu \omega^{\nu\|\mathbf{v}\|} \left(\frac{1-\omega^\nu}{2}\right)^{|A|} \left(\frac{1+\omega^\nu}{2}\right)^{b'-|A|}.$$

Taking absolute values and using $|(-1)^c| = |\omega^{\nu\|\mathbf{v}\|}| = 1$, the triangle inequality gives

$$\left|q(M', f) - \frac{1}{2}\right| \le \frac{1}{2} \sum_{\substack{0\le\nu<t\\ \nu\neq 0,t/2}} |\widehat{g}_\nu| \left|\frac{1-\omega^\nu}{2}\right|^{|A|} \left|\frac{1+\omega^\nu}{2}\right|^{b'-|A|}.$$

Since

$$\left|\frac{1-\omega^\nu}{2}\right| = \left|\sin\frac{\pi\nu}{t}\right|, \qquad \left|\frac{1+\omega^\nu}{2}\right| = \left|\cos\frac{\pi\nu}{t}\right|,$$

we obtain (22).

Finally, for every $\nu \notin \{0, t/2\}$ with $0 \le \nu < t$,

$$\left|\sin \frac{\pi\nu}{t}\right| \le \cos \frac{\pi}{t}, \qquad \left|\cos \frac{\pi\nu}{t}\right| \le \cos \frac{\pi}{t}.$$

Therefore,

$$\left|q(M', f) - \frac{1}{2}\right| \le \frac{1}{2}\left(\cos \frac{\pi}{t}\right)^{b'} \sum_{\substack{0 \le \nu < t \\ \nu \ne 0, t/2}} |\widehat{g}_\nu|.$$

By Parseval's identity and $|g(w)| = 1$,

$$\sum_{\substack{0 \le \nu < t \\ \nu \ne 0, t/2}} |\widehat{g}_\nu|^2 = \sum_{\nu=0}^{t-1} |\widehat{g}_\nu|^2 = \frac{1}{t}\sum_{w=0}^{t-1} |g(w)|^2 = 1.$$

The Cauchy–Schwarz inequality gives

$$\sum_{\substack{0 \le \nu < t \\ \nu \ne 0, t/2}} |\widehat{g}_\nu| \le \sqrt{(t-2) \sum_{\substack{0 \le \nu < t \\ \nu \ne 0, t/2}} |\widehat{g}_\nu|^2} = \sqrt{t-2}.$$

Consequently, for every residual M' and affine function f,

$$\left|q(M', f) - \frac{1}{2}\right| \le \min\left\{\frac{1}{2}, \frac{\sqrt{t-2}}{2}\left(\cos \frac{\pi}{t}\right)^{b'}\right\}.$$

Since $f \oplus 1$ is also affine, the bit-fixing affine advantage satisfies

$$\mathrm{adv}_{b'}(M_b) = \max_{\substack{M' \in \mathcal{R}_{b'}(M_b) \\ f \in \mathcal{A}_{b'}}} \left|q(M', f) - \frac{1}{2}\right|.$$

Taking the maximum over all residuals and affine functions proves (23). □

B.2 Properties of Modulus Functions

Proof of Theorem 3. We first prove necessity. Assume that $\mathrm{adv}_{b'}(\mathsf{MOD}_b^{t,S_0}) \to 0$ as $b' \to \infty$, uniformly over all $b \ge b'$. In particular, taking $b = b'$ gives $\mathrm{adv}_{b'}(\mathsf{MOD}_{b'}^{t,S_0}) \to 0$.

Let $W_b \triangleq \|U_b\|$ denote the Hamming weight of a uniformly random b-bit vector. We first show that $W_b \bmod t$ converges to the uniform distribution over $\{0, \ldots, t-1\}$, without imposing any condition on S_0 or the parity of t. For $0 \le j < t$, define

$$p_b(j) \triangleq \Pr[W_b \bmod t = j].$$

Viewing p_b as a t-periodic function, its Fourier coefficients satisfy

$$\begin{aligned}
\widehat{p}_b(\nu) &= \frac{1}{t}\sum_{j=0}^{t-1} p_b(j)\omega^{-\nu j} \\
&= \frac{1}{t}\mathbb{E}[\omega^{-\nu W_b}] \\
&= \frac{1}{t}\left(\frac{1+\omega^{-\nu}}{2}\right)^b,
\end{aligned}$$

where $\omega = \exp(2\pi\mathrm{i}/t)$ and $0 \le \nu < t$. The last equality follows from the independence of the input bits.

By Fourier inversion,

$$\begin{aligned} p_b(j) &= \sum_{\nu=0}^{t-1} \widehat{p}_b(\nu)\omega^{\nu j} \\ &= \frac{1}{t} + \frac{1}{t}\sum_{\nu=1}^{t-1} \omega^{\nu j}\left(\frac{1+\omega^{-\nu}}{2}\right)^b. \end{aligned}$$

For every $1 \le \nu < t$,

$$\left|\frac{1+\omega^{-\nu}}{2}\right| = \left|\cos\frac{\pi\nu}{t}\right| < 1.$$

Since t is fixed, every term in the sum tends to zero as $b \to \infty$. Therefore,

$$p_b(j) \longrightarrow \frac{1}{t} \qquad (0 \le j < t).$$

Consequently,

$$\Pr[M_b(U_b) = 0] = \sum_{j \in S_0} p_b(j) \longrightarrow \frac{|S_0|}{t}.$$

Since constant functions are affine, $\mathrm{adv}_b(M_b) \to 0$ requires

$$|S_0| = \frac{t}{2}.$$

In particular, t must be even.

Let

$$e = |S_0 \cap 2\mathbb{Z}|, \qquad o = |S_0 \cap (2\mathbb{Z}+1)|.$$

Consider the affine function

$$f(\mathbf{x}) = \bigoplus_{j \in [b]} x_j.$$

Since t is even, the Hamming weight of an input and its residue modulo t have the same parity. Thus, f equals zero on even residues and one on odd residues, while M_b equals zero exactly on S_0. The two functions therefore agree on the e even residues in S_0 and the $t/2 - o$ odd residues outside S_0. Since the residue distribution converges to uniformity,

$$q(M_b, f) \longrightarrow \frac{e + t/2 - o}{t} = \frac{1}{2} + \frac{e-o}{t}.$$

The complement $f \oplus 1$ is also affine, so

$$\max\{q(M_b, f), q(M_b, f \oplus 1)\} \longrightarrow \frac{1}{2} + \frac{|e-o|}{t}.$$

Therefore, $\mathrm{adv}_b(M_b) \to 0$ requires $e = o$. Together with $e + o = t/2$, this yields

$$e = o = \frac{t}{4}.$$

Since e and o are integers, $4 \mid t$. This proves necessity.

Conversely, suppose that $4 \mid t$ and condition (21) holds. Let $b \geq b'$ and let $M' \in \mathcal{R}_{b'}(M_b)$ be an arbitrary residual. For every affine function f,

$$q(M', f) + q(M', f \oplus 1) = 1,$$

so the maximum affine agreement probability is at least $1/2$. By Lemma 5,

$$0 \leq \max_{f \in \mathcal{A}_{b'}} q(M', f) - \frac{1}{2} \leq \frac{\sqrt{t-2}}{2} \left(\cos \frac{\pi}{t}\right)^{b'}.$$

Since t is fixed and $\cos(\pi/t) < 1$, this bound tends to zero as $b' \to \infty$, uniformly over $b \geq b'$ and all residuals $M' \in \mathcal{R}_{b'}(M_b)$. Taking the maximum over these residuals gives $\mathrm{adv}_{b'}(M_b) \to 0$ uniformly, proving sufficiency. □

Lemma 6 (Degree bound for power-of-two moduli)**.** *Let $t = 2^\ell$ with $\ell \geq 1$. For every residue set $S_0 \subseteq \{0, \ldots, t-1\}$ and every $b \geq 1$,*

$$\deg(\mathsf{MOD}_b^{t,S_0}) \leq \min\{b, t-1\}.$$

Proof. Since the signed weight profile g_{t,S_0} has period t, we have $\lambda_j(t) = \lambda_j(0)$ for every $j \geq 0$. By Lemma 7 with $k = \ell$,

$$\lambda_{j+t}(0) = \lambda_j(0) + \lambda_j(t) = 0 \qquad (j \geq 0).$$

Thus, all ANF coefficients of degree at least t vanish. Since the algebraic degree is also at most b, the result follows. □

Proof of Theorem 6. Fix $1 \leq b' \leq b$ and an arbitrary residual $M' \in \mathcal{R}_{b'}(\mathsf{MOD12}_b)$. Let $g = g_{12,S_0}$ be the signed weight profile of $\mathsf{MOD12}_b$. For $S_0 = \{0, 1, 3, 4, 6, 9\}$, we have

$$\begin{aligned}&(g(0), \ldots, g(11))\\ &\quad = (1, 1, -1, 1, 1, -1, 1, -1, -1, 1, -1, -1).\end{aligned}$$

Substituting these values into the definition of the Fourier coefficients gives

$$|\widehat{g}_\nu| = \begin{cases} 0, & \nu \in \{0, 2, 6, 10\}, \\ \dfrac{\sqrt{2}}{6}, & \nu \in \{1, 3, 5, 7, 9, 11\}, \\ \dfrac{1}{\sqrt{3}}, & \nu \in \{4, 8\}. \end{cases}$$

Consider any affine function of the form

$$f(\mathbf{z}) = c \oplus \bigoplus_{j \in A} z_j, \qquad A \subseteq [b'], \quad c \in \mathbb{F}_2.$$

Set

$$\rho = \cos \frac{\pi}{12}, \qquad \eta = \sin \frac{\pi}{12}.$$

Substituting these Fourier coefficients into (22) and combining the terms for ν and $12 - \nu$, we obtain

$$\begin{aligned}\left| q(M', f) - \frac{1}{2} \right| \leq{}& \frac{\sqrt{2}}{6} \left(\rho^{b'-|A|} \eta^{|A|} + \eta^{b'-|A|} \rho^{|A|} + 2^{-b'/2} \right) \\ &+ \frac{1}{\sqrt{3}} \left(\frac{1}{2}\right)^{b'-|A|} \left(\frac{\sqrt{3}}{2}\right)^{|A|}.\end{aligned} \tag{24}$$

To obtain a bound independent of $|A|$, we apply the weighted arithmetic–geometric mean inequality:

$$\rho^{b'-|A|}\eta^{|A|} \le \left(1 - \frac{|A|}{b'}\right)\rho^{b'} + \frac{|A|}{b'}\eta^{b'},$$
$$\eta^{b'-|A|}\rho^{|A|} \le \left(1 - \frac{|A|}{b'}\right)\eta^{b'} + \frac{|A|}{b'}\rho^{b'}.$$

Adding these inequalities yields

$$\rho^{b'-|A|}\eta^{|A|} + \eta^{b'-|A|}\rho^{|A|} \le \rho^{b'} + \eta^{b'}.$$

Moreover, since $1/2 \le \sqrt{3}/2$,

$$\left(\frac{1}{2}\right)^{b'-|A|}\left(\frac{\sqrt{3}}{2}\right)^{|A|} \le \left(\frac{\sqrt{3}}{2}\right)^{b'}.$$

Substituting these bounds into (24) gives

$$\left|q(M', f) - \frac{1}{2}\right| \le B_{12}(b').$$

Since the bound holds for every residual $M' \in \mathcal{R}_{b'}(\mathsf{MOD12}_b)$ and affine function $f \in \mathcal{A}_{b'}$, taking their maximum gives

$$\mathrm{adv}_{b'}(\mathsf{MOD12}_b) \le B_{12}(b').$$

Combining this with the trivial bound $\mathrm{adv}_{b'}(\mathsf{MOD12}_b) \le 1/2$ proves (7). □

B.3 Bit-Fixing Algebraic Degree

Let $t \ge 2$ and $S_0 \subseteq \{0, \dots, t-1\}$, and consider $M_b = \mathsf{MOD}_b^{t,S_0}$ with signed weight profile $g = g_{t,S_0}$. For integers $j, w \ge 0$, define

$$\lambda_j(w) \triangleq \sum_{h=0}^{j} \binom{j}{h} \frac{1 - g(w+h)}{2} \pmod 2.$$

For any constraint $(T, \mathbf{v}_T)$, let $b' = b - |T|$. The residual $M_{b|T,\mathbf{v}_T}$ is symmetric, and $\lambda_j(\|\mathbf{v}_T\|)$ is the common ANF coefficient of all degree-j monomials, for $0 \le j \le b'$. All sums involving these coefficients are evaluated over $\mathbb{F}_2$.

Lemma 7 (Binomial transform identity)**.** *For any modular function $M_b = \mathsf{MOD}_b^{t,S_0}$ and all integers $j, k, w \ge 0$,*

$$\lambda_{j+2^k}(w) = \lambda_j(w) + \lambda_j(w + 2^k).$$

The identity also holds for the binomial transform of any binary sequence.

Proof. Over $\mathbb{F}_2$, we have

$$(1+X)^{2^k} = 1 + X^{2^k}.$$

Consequently,

$$(1+X)^{j+2^k} = (1+X)^j(1+X^{2^k}).$$

Comparing coefficients of X^h gives

$$\binom{j+2^k}{h} = \binom{j}{h} + \binom{j}{h-2^k} \pmod 2,$$

where binomial coefficients outside their valid range are taken to be zero.

Substituting this identity into the definition of $\lambda_{j+2^k}(w)$ yields

$$\begin{aligned}\lambda_{j+2^k}(w) &= \sum_{h=0}^{j+2^k} \binom{j+2^k}{h} \frac{1-g(w+h)}{2} \\ &= \sum_{h=0}^{j} \binom{j}{h} \frac{1-g(w+h)}{2} + \sum_{h=0}^{j} \binom{j}{h} \frac{1-g(w+h+2^k)}{2} \\ &= \lambda_j(w) + \lambda_j(w+2^k),\end{aligned}$$

where all equalities are over $\mathbb{F}_2$. The same argument applies to any binary sequence. □

Lemma 8 (Structure of ANF coefficients)**.** *For* $\mathsf{MOD12}_b$*, the coefficients* $\lambda_j(w)$ *satisfy*

$$\lambda_{j+12}(w) = \lambda_j(w), \qquad \lambda_j(w+4) = \lambda_{j+8}(w)$$

for every $w \geq 0$ *and* $j \geq 4$*. Moreover, the sequence* $(\lambda_j(w))_{j\geq 4}$ *contains no six consecutive zeros, and at least one of* $\lambda_4(w), \lambda_5(w), \lambda_6(w)$ *is nonzero.*

Proof. Let $g = g_{12,S_0}$ be the signed weight profile of $\mathsf{MOD12}_b$. Since g has period 12, we have $\lambda_j(w+12) = \lambda_j(w)$ for all $j, w \geq 0$. By Lemma 7,

$$\begin{aligned}\lambda_{j+16}(w) &= \lambda_j(w) + \lambda_j(w+16) \\ &= \lambda_j(w) + \lambda_j(w+4) \\ &= \lambda_{j+4}(w).\end{aligned}$$

Replacing $j+4$ by j gives

$$\lambda_{j+12}(w) = \lambda_j(w) \qquad (j \geq 4).$$

For the second identity, define the binary sequence

$$\gamma_h \triangleq \frac{1-g(h)}{2} \oplus \frac{1-g(h+4)}{2} \oplus \frac{1-g(h+8)}{2}.$$

Since g has period 12, the sequence γ_h has period 4.

Applying Lemma 7 to the binary sequence γ_h, and writing $j = \ell + 4$, we obtain

$$\begin{aligned}\sum_{h=0}^{j} \binom{j}{h} \gamma_{w+h} &= \sum_{h=0}^{\ell} \binom{\ell}{h} \left(\gamma_{w+h} + \gamma_{w+h+4}\right) \\ &= 0,\end{aligned}$$

where the last equality follows from the period 4 of γ_h. Expanding the definition of γ_h gives

$$\lambda_j(w) + \lambda_j(w+4) + \lambda_j(w+8) = 0 \qquad (j \geq 4).$$

On the other hand, Lemma 7 gives

$$\lambda_{j+8}(w) = \lambda_j(w) + \lambda_j(w+8).$$

Combining these equalities yields

$$\lambda_j(w+4) = \lambda_{j+8}(w) \qquad (j \geq 4).$$

The two identities imply that it suffices to examine the following four periodic coefficient sequences. The remaining values of w mod 12 correspond to cyclic shifts by four or eight positions:

w	$(\lambda_4(w), \ldots, \lambda_{15}(w))$
0	$(0,1,1,1,1,1,1,0,1,0,0,1)$
1	$(1,0,0,0,0,0,1,1,1,0,1,1)$
2	$(1,0,0,0,0,1,0,0,1,1,0,0)$
3	$(1,0,0,0,1,1,0,1,0,1,0,1)$

Inspection of these sequences, including across their period boundaries, shows that no six consecutive coefficients are zero. The table and the corresponding cyclic shifts also show that at least one of $\lambda_4(w), \lambda_5(w), \lambda_6(w)$ is nonzero for every w. □

Proof of Theorem 4. Fix an arbitrary residual $M' \in \mathcal{R}_{b'}(\mathsf{MOD12}_b)$ and write $M' = M_{b|T,\mathbf{v}_T}$ with $|T| = b - b'$. By the definition of $\lambda_j(\|\mathbf{v}_T\|)$, it is the common coefficient of all degree-j monomials in the ANF of $M_{b|T,\mathbf{v}_T}$, for $0 \leq j \leq b'$.

For $b' \geq 9$, Lemma 8 ensures that at least one of $\lambda_{b'-5}(\|\mathbf{v}_T\|), \ldots, \lambda_{b'}(\|\mathbf{v}_T\|)$ is nonzero, giving

$$\deg(M_{b|T,\mathbf{v}_T}) \geq b' - 5.$$

For $6 \leq b' \leq 8$, the second assertion of Lemma 8 ensures that at least one of $\lambda_4(\|\mathbf{v}_T\|)$, $\lambda_5(\|\mathbf{v}_T\|)$, and $\lambda_6(\|\mathbf{v}_T\|)$ is nonzero. Therefore,

$$\deg(M_{b|T,\mathbf{v}_T}) \geq 4 \geq b' - 5.$$

The upper bound $\deg(M') \leq b'$ follows because M' has b' inputs. □

B.4 Bit-Fixing Algebraic Immunity

Lemma 9. *Let H be a symmetric Boolean polynomial on d variables with degree at most e. Let $r_0 \in \{0,1,2\}$ and suppose that $r_0 + 3e \leq d$. If*

$$H(\mathbf{z}) = 0 \quad \textit{whenever} \quad \|\mathbf{z}\| \in \{r_0, r_0+3, \ldots, r_0+3e\},$$

then $H \equiv 0$.

Proof. Since H is symmetric, all degree-j monomials in its ANF have the same coefficient, denoted by $c_j \in \mathbb{F}_2$. For an input $\mathbf{z}$ of Hamming weight w, we have

$$H(\mathbf{z}) = \sum_{j=0}^{e} \binom{w}{j} c_j \pmod 2.$$

By assumption, $H(\mathbf{z}) = 0$ whenever $\|\mathbf{z}\| = r_0 + 3i$, for $0 \leq i \leq e$. Therefore,

$$\sum_{j=0}^{e} \binom{r_0+3i}{j} c_j = 0 \quad \text{in } \mathbb{F}_2, \qquad 0 \leq i \leq e.$$

The determinant of the corresponding integer matrix is

$$\det\left[\binom{r_0+3i}{j}\right]_{0\le i,j\le e} = \frac{\prod_{0\le i<j\le e} 3(j-i)}{\prod_{j=0}^{e} j!} = 3^{e(e+1)/2}.$$

The first equality follows from the Vandermonde determinant formula; see also [BL13, Lemma 5.1].

Since the determinant is odd, the matrix is invertible over $\mathbb{F}_2$. Hence $c_0 = \cdots = c_e = 0$, implying $H \equiv 0$. □

Lemma 10 (Algebraic immunity lower bound)**.** *Let P_d be a symmetric Boolean function on $d \ge 6$ inputs. Suppose that there exist distinct residues $\alpha, \beta \in \{0,1,2\}$ such that*

$$P_d(\mathbf{z}) = 0 \quad \text{if } \|\mathbf{z}\| \equiv \alpha \pmod 3,$$
$$P_d(\mathbf{z}) = 1 \quad \text{if } \|\mathbf{z}\| \equiv \beta \pmod 3.$$

Then

$$\mathrm{AI}(P_d) \ge \left\lfloor \frac{d+1}{3} \right\rfloor.$$

Proof. Suppose, for contradiction, that

$$\delta = \mathrm{AI}(P_d) < \left\lfloor \frac{d+1}{3} \right\rfloor.$$

Then there exists $\varepsilon \in \{0,1\}$ such that $P_d \oplus \varepsilon$ has a nonzero annihilator of degree δ.

Since P_d is symmetric, the annihilators of $P_d \oplus \varepsilon$ form a permutation-invariant ideal generated by $P_d \oplus \varepsilon \oplus 1$. By the symmetrization lemma [BL13, Lemma 4.1], there exists a nonzero annihilator of the form

$$H(z_{2\ell}, \ldots, z_{d-1}) \prod_{j\in[\ell]} (z_{2j} \oplus z_{2j+1}),$$

after a permutation of the inputs, where $0 \le \ell \le \delta$ and H is a nonzero symmetric polynomial on $d' = d - 2\ell$ variables with degree at most $e = \delta - \ell$.

For each $j \in [\ell]$, fix $z_{2j} \ne z_{2j+1}$. Then every XOR factor equals one, so the annihilator reduces to H. Moreover, each pair contributes exactly one to the Hamming weight. Thus, if the remaining inputs have Hamming weight w, the full input has Hamming weight $w + \ell$.

By assumption, $P_d \oplus \varepsilon$ equals one at every weight congruent to β modulo 3 when $\varepsilon = 0$, and to α modulo 3 when $\varepsilon = 1$. Since the polynomial is an annihilator of $P_d \oplus \varepsilon$, H must vanish at every weight congruent to r_0 modulo 3, where

$$r_0 \equiv \begin{cases} \beta - \ell, & \text{if } \varepsilon = 0, \\ \alpha - \ell, & \text{if } \varepsilon = 1 \end{cases} \pmod 3, \qquad r_0 \in \{0,1,2\}.$$

The assumption on δ implies $d \ge 3\delta + 2$. Consequently,

$$d' = d - 2\ell \ge 3\delta + 2 - 2\ell = 3e + \ell + 2 \ge 3e + 2.$$

Since $r_0 \le 2$, all weights $r_0, r_0 + 3, \ldots, r_0 + 3e$ lie between 0 and d'.

By Lemma 9, we have $H \equiv 0$, contradicting the choice of H. Therefore,

$$\mathrm{AI}(P_d) \geq \left\lfloor \frac{d+1}{3} \right\rfloor.$$

□

Proof of Theorem 5. Fix an arbitrary residual $M' \in \mathcal{R}_{b'}(\mathsf{MOD12}_b)$ and write $M' = M_{b|T,\mathbf{v}_T}$ with $|T| = b - b'$. By the definition of S_0, $\mathsf{MOD12}_b$ equals zero at every Hamming weight congruent to 0 modulo 3 and one at every Hamming weight congruent to 2 modulo 3. Consequently,

$$M_{b|T,\mathbf{v}_T}(\mathbf{z}) = 0 \quad \text{if } \|\mathbf{z}\| \equiv -\|\mathbf{v}_T\| \pmod 3,$$
$$M_{b|T,\mathbf{v}_T}(\mathbf{z}) = 1 \quad \text{if } \|\mathbf{z}\| \equiv 2 - \|\mathbf{v}_T\| \pmod 3.$$

Since these are distinct residue classes, Lemma 10 gives

$$\mathrm{AI}(M_{b|T,\mathbf{v}_T}) \geq \left\lfloor \frac{b'+1}{3} \right\rfloor.$$

The asymptotic claim follows since the algebraic immunity of any Boolean function on b' inputs is at most b'. □

References

[ABBS25] Amit Agarwal, Carsten Baum, Lennart Braun, and Peter Scholl. Low-Bandwidth Mixed Arithmetic in VOLE-Based ZK from Low-Degree PRGs. In *EUROCRYPT (4) 2025*, pages 396–426, 2025.

[ADI+17] Benny Applebaum, Ivan Damgård, Yuval Ishai, Michael Nielsen, and Lior Zichron. Secure Arithmetic Computation with Constant Computational Overhead. In *CRYPTO (1) 2017*, pages 223–254, 2017.

[AK23] Benny Applebaum and Niv Konstantini. Actively Secure Arithmetic Computation and VOLE with Constant Computational Overhead. In *EUROCRYPT (2) 2023*, pages 190–219, 2023.

[AL16] Benny Applebaum and Shachar Lovett. Algebraic attacks against random local functions and their countermeasures. In *STOC 2016*, pages 1087–1100, 2016.

[App16] Benny Applebaum. Cryptographic Hardness of Random Local Functions - Survey. *Computational Complexity*, 25(3):667–722, 2016.

[BCG+17] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, and Michele Orrù. Homomorphic Secret Sharing: Optimizations and Applications. In *CCS 2017*, pages 2105–2122, 2017.

[BCM+24] Dung Bui, Geoffroy Couteau, Pierre Meyer, Alain Passelègue, and Mahshid Riahinia. Fast Public-Key Silent OT and More from Constrained Naor-Reingold. In *EUROCRYPT (6) 2024*, pages 88–118, 2024.

[BCPR25] Christina Boura, Geoffroy Couteau, Léo Perrin, and Yann Rotella. Sok: On shallow weak prfs: A common symmetric building block for mpc protocols. *IACR Transactions on Symmetric Cryptology*, 2025(3):289–336, 2025.

[BL13] Chris Beck and Yuan Li. Represent MOD function by low degree polynomial with unbounded one-sided error, 2013. arXiv preprint arXiv:1304.0713.

[CDD+24] Geoffroy Couteau, Lalita Devadas, Srinivas Devadas, Alexander Koch, and Sacha Servan-Schreiber. QuietOT: Lightweight oblivious transfer with a public-key setup. In *Advances in Cryptology – ASIACRYPT 2024, Part II*, volume 15485 of *LNCS*, pages 197–231. Springer, 2024.

[CDM+18] Geoffroy Couteau, Aurélien Dupin, Pierrick Méaux, Mélissa Rossi, and Yann Rotella. On the Concrete Security of Goldreich's Pseudorandom Generator. In *ASIACRYPT (2) 2018*, pages 96–124, 2018.

[CV05] Anne Canteaut and Marion Videau. Symmetric boolean functions. *IEEE Transactions on Information Theory*, 51(8):2791–2811, 2005.

[DMÜ26] Antonio Giulio D'Antona, Pierrick Méaux, and Akin Ünal. On the common bias of majorities: Poly-time attacks on THR-XOR PRGs. Cryptology ePrint Archive, Paper 2026/880, 2026.

[EKM17] Andre Esser, Robert Kübler, and Alexander May. LPN Decoded. In *CRYPTO (2) 2017*, pages 486–514, 2017.

[FLLL26] Ximing Fu, Mo Li, Shihan Lyu, and Chuanyi Liu. Attacks on goldreich's pseudorandom generators by grouping and solving. In *Advances in Cryptology – EUROCRYPT 2026, Part VI*, volume 16546 of *LNCS*, pages 3–32. Springer, 2026. Earlier version: IACR ePrint 2024/1594.

[Gol00] Oded Goldreich. Candidate One-Way Functions Based on Expander Graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(90), 2000.

[GRR+16] Lorenzo Grassi, Christian Rechberger, Dragos Rotaru, Peter Scholl, and Nigel P. Smart. MPC-Friendly Symmetric Key Primitives. In *CCS 2016*, pages 430–443, 2016.

[IKOS08] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with Constant Computational Overhead. In *STOC 2008*, pages 433–442, 2008.

[LLMF26] Shihan Lyu, Mo LI, Chuhan Ma, and Ximing Fu. Improved cryptanalysis of local PRGs. Cryptology ePrint Archive, Paper 2026/2240, 2026.

[LV17] Alex Lombardi and Vinod Vaikuntanathan. Limits on the Locality of Pseudorandom Generators and Applications to Indistinguishability Obfuscation. In *TCC (1) 2017*, pages 119–137, 2017.

[LWYY25] Hanlin Liu, Xiao Wang, Kang Yang, and Yu Yu. BitGC: Garbled circuits with 1 bit per gate. In *Advances in Cryptology – EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 437–466. Springer, 2025.

[MMT11] Alexander May, Alexander Meurer, and Enrico Thomae. Decoding Random Linear Codes in $\tilde{\mathcal{O}}(2^{0.054n})$. In *ASIACRYPT 2011*, pages 107–124, 2011.

[OW14] Ryan O'Donnell and David Witmer. Goldreich's PRG: Evidence for Near-Optimal Polynomial Stretch. In *CCC 2014*, pages 1–12, 2014.

[WWLM26] Tianning Wang, Haoyang Wang, Fukang Liu, and Willi Meier. Cryptanalysis of goldreich's PRGs with MAJ–XOR predicates. Cryptology ePrint Archive, Paper 2026/2156, 2026.

[YGJL22] Jing Yang, Qian Guo, Thomas Johansson, and Michael Lentmaier. Revisiting the Concrete Security of Goldreich's Pseudorandom Generator. *IEEE Trans. Inf. Theory*, 68(2):1329–1354, 2022.

[ZLY+26] Wenhao Zhang, Hanlin Liu, Kang Yang, Wen jie Lu, Yu Yu, Xiao Wang, and Chenkai Weng. BitGC made (more) efficient. In *Advances in Cryptology – EUROCRYPT 2026, Part III*, volume 16543 of *LNCS*, pages 214–245. Springer, 2026.