

New Imperfect High Threshold and Slice Secret Sharing Schemes

Andrej Bogdanov*

Yanbo Chen†

Abstract

We present new constructions of $(n - t + 1)$ -out-of- n threshold and slice secret sharing schemes for short secrets with imperfect secrecy for constant t and growing n . Our main results are

1. an ϵ -differentially secret, $(n - t + 1)$ -out-of- n bit secret sharing scheme with share alphabet size $\tilde{O}_t(1/\epsilon^2)$. In contrast, perfect secrecy requires share alphabet size $n - t + 1$ (Bogdanov, Guo, and Komargodski, ToC 2020). For $t > 3$, the security of our scheme relies on a conjecture about factorizations of univariate polynomials in finite fields.
2. an ϵ -statistically secret bit secret sharing scheme for every $(n - 2)$ -out-of- n slice with share size $O(\log n \log 1/\epsilon)$. Previous constructions with comparable parameters achieved only computational security (Beimel, Farràs, Lasri, and Nir, TCC 2024).

Variants of our first construction produce plain and anonymous $(n - 1)$ -out-of- n weak bit secret sharing schemes with optimal share size. We discuss the relevance of these schemes to secret sharing duality.

1 Introduction

Threshold secret sharing schemes [Bla79, Sha79] for n parties require share alphabet of size $\Omega(n)$ even when the secret is a single bit [KN90, BGK20]. Can more compact schemes be obtained by weakening the requirements?

A common relaxation is to allow unqualified parties a small advantage in guessing the secret. In this work we consider computationally unbounded distinguishers. A scheme is ϵ -statistically secret if no unqualified set can distinguish secrets with advantage ϵ .

A stronger solution concept is ϵ -differential secrecy: Any possible view of the unqualified parties (not merely a probable view) cannot tell secrets apart with advantage more than ϵ . We adopt the following definition. For every pair of secrets 0, 1, every unqualified set S receiving shares x_S , and every possible share sequence a_S ,

$$\left| \frac{\Pr(x_S = a_S|0) - \Pr(x_S = a_S|1)}{\Pr(x_S = a_S|0) + \Pr(x_S = a_S|1)} \right| \leq \epsilon, \quad (1)$$

where $\Pr(\cdot|\sigma)$ is the probability of a given share sequence for secret σ .

If a scheme is ϵ -differentially secret, it is ϵ -statistically secret, as the statistical distance between the views is at most

$$\frac{1}{2} \sum_{a_S} |\Pr(x_S = a_S|0) - \Pr(x_S = a_S|1)| \leq \epsilon \cdot \frac{1}{2} \left(\sum_{a_S} \Pr(x_S = a_S|0) + \sum_{a_S} \Pr(x_S = a_S|1) \right) = \epsilon.$$

The converse does not hold. Under statistical secrecy, there may be unlucky outcomes of sharing that allow unqualified parties to perfectly recover the secret. We'll see an example in Section 3.

Zero-differential secrecy is the same as perfect secrecy. If we merely require that ϵ is strictly less than 1 we obtain Beimel and Franklin's definition of weak secrecy [BF07]: For every unqualified set, the sets of possible share sequences of 0 and 1 are identical.¹

Beimel and Franklin show that weak secrecy can overcome information-theoretic barriers for perfect secrecy: There is a weakly secret 2-out-of- n bit secret sharing scheme with share alphabet size 4. Perfect

*University of Ottawa. abogdano@uottawa.ca.

†University of Ottawa. ychen918@uottawa.ca.

¹More generally, Beimel et al.'s notion of p -weak secrecy [BIKL25] is equivalent to $|(1-p)/(1+p)|$ -differential secrecy.

secrecy requires alphabet size n . Beimel, Ishai, Kushilevitz, and Li [BIKL25] further show weak bit secret sharing for any access structure where the unqualified sets have size at most t and the share alphabet has size $2^{t+O(\log t)}$.

A curious phenomenon in secret sharing is that in many instances, the best known share size for a given access structure is the same or close to that of its dual [Csi21]. The dual $A^\dagger$ of access structure A is obtained by taking the complements of all qualified/unqualified sets and switching their roles.

It is known that *linear* secret sharing is closed under duality [Gál95, Lok09]: If A admits a linear scheme with a given (total or maximum) share size, so does $A^\dagger$. It is not known whether duality extends to general (non-linear) schemes. Duality, however, appears to have surprising predictive power. It has led to new schemes for access structures $A^\dagger$ matching or approximating the best known constructions for A [Bog23, BFLN24].

Threshold access structures

The dual of a threshold access structure is a threshold access structure. Specifically, $(t\text{-out-of-}n)^\dagger$ is $(n - t + 1)\text{-out-of-}n$. Beimel et al. [BIKL25] ask whether their alphabet size upper bounds for weak $t\text{-out-of-}n$ secret sharing for $t \ll n$ can be matched for the corresponding dual access structures. We show that this is the case when $t = 2$:

Theorem 1. *For every n , there exists a $\Omega(1)$ -differentially secret $(n - 1)\text{-out-of-}n$ bit secret sharing scheme with share alphabet size 4.*

Differential secrecy can be amplified (see Proposition 18). In particular, Theorem 1 implies ϵ -differential secrecy for $(n - 1)\text{-out-of-}n$ bit sharing with share alphabet size $1/\epsilon^{O(1)}$. Our main result is a conditional extension to other high thresholds:

Theorem 2. *For every t, ϵ , and sufficiently large n there exists an anonymous, ϵ -differentially secret $(n - t + 1)\text{-out-of-}n$ Σ -secret sharing scheme with share alphabet size $O_t(|\Sigma|/\epsilon^2 \cdot \log(|\Sigma|/\epsilon))$, assuming Conjecture 12.*

Conjecture 12 is a statement about statistics of factorizations of low-degree univariate polynomials. We prove it for degree 2 and 3, making Theorem 2 unconditional for $(n - 1)\text{-out-of-}n$ and $(n - 2)\text{-out-of-}n$ secret sharing.

An equivalent form of Theorem 2 is that for fixed t , the share size for a k -bit secret is $k + 2 \log(1/\epsilon)$ plus some lower-order terms. This is close to best possible.

Proposition 3. *In an ϵ -differentially secret scheme, the share size of every relevant party must be at least the secret size minus $\log(1 + \epsilon)/(1 - \epsilon)$.*

Party i is relevant if there exists an unqualified set S such that $S \cup \{i\}$ is qualified.

One consequence of Theorem 1 or 2 is that duality does not hold for statistical bit secret sharing. As differential secrecy implies statistical secrecy, there exist $(n - 1)\text{-out-of-}n$ ϵ -statistically secure schemes with share alphabet of size $\tilde{O}(1/\epsilon^2)$. In contrast, even a $\frac{1}{3}$ -statistically secure 2-out-of- n scheme requires alphabet size $\frac{4}{9}n$.

Proposition 4. *ϵ -statistically secure 2-out-of- n bit secret sharing requires share alphabet size $(1 - \epsilon)^2 n$.*

Csirmaz [Csi21] refutes a related approximate variant of duality. He constructs an almost ideal access structure A with 174 parties whose dual $A^\dagger$ is not almost ideal. An access structure is almost ideal if it admits a sequence of schemes, indexed by secret size, whose shares and secrets are of equal size (the ideal part), and whose secrecy and reconstruction errors converge to zero (the almost part).² In particular, for any fixed secrecy and reconstruction errors and sufficiently long secrets, A admits secret sharing with smaller share size than $A^\dagger$ does.

Csirmaz's counterexample to approximate duality is incomparable to ours. Theorem 2 is not interesting when the secret alphabet size exceeds the number of parties. All threshold access structures are in fact ideal. For the same reason, Csirmaz's result has no relevance to threshold secret sharing duality with short secrets and many parties.

²Csirmaz plausibly conjectures that the reconstruction error can be eliminated.

The top slice

In a t -slice (partial) access structure both the minimal qualified sets and the maximal unqualified sets have size t .³ Recent years have seen a flurry of improvements in the share size when $t \ll n/2$ [BP18, LVW17, AA20, LV18, ABF⁺19].

Beimel, Farràs, Lasri, and Nir [BFLN24] asked whether related improvements can be obtained for high slices. Their main results are a perfect scheme for all $(n - t + 1)$ slices with share size $tn \exp(\tilde{O}(\sqrt{t \log n}))$, and an ϵ -computationally secure scheme with share size $O(t^2(\log n)(\log 1/\epsilon))$.

We show that the same *statistical* secrecy can be achieved for the top $(n - 2)$ -slice:

Theorem 5. *For every $(n - 2)$ -out-of- n slice there exists a ϵ -statistically secret bit secret sharing scheme with share alphabet size $n^{O(\log 1/\epsilon)}$.*

For comparison, 2-slices have perfectly secret bit sharing with share size $\exp \tilde{O}(\sqrt{\log n})$ [LVW17]. We are not aware of any improvement if the secrecy requirement is relaxed to statistical.

Anonymous schemes

A scheme is anonymous if the distribution on shares induced by every secret is permutation-invariant. Optimizing constants in the proof of Theorem 2 we show that

Corollary 6. *For every n , there exists an anonymous $(n - 1)$ -out-of- n bit secret sharing scheme with alphabet size 5, perfect reconstruction, and 15/17-differential secrecy.*

The alphabet size is best possible.

Proposition 7. *For every $n \geq 5$, anonymous $(n - 1)$ -out-of- n weak bit secret sharing requires share alphabet of size 5.*

One consequence of Proposition 7 is that the maximum share size for anonymous weak bit secret sharing is not closed under duality: 2-out-of-5 bit secret sharing allows share alphabet size 4 [BF07], but 4-out-of-5 requires share alphabet size 5.

Techniques

All our schemes have a similar structure. Shares are sampled uniformly conditioned on certain relations so that all qualified sets can recover the missing ones (possibly with some extra public information). In contrast, the views of unqualified sets are consistent with several possible choices of the missing shares. The secret is determined by a coloring H of the possible share sequences. Secrecy amounts to equidistribution of the coloring among possible choices of the missing shares. To illustrate these ideas, we prove Theorem 1.

$(n - 1)$ -out-of- n scheme

Share: Given a secret $\sigma \in \{0, 1\}$, sample x and y uniformly from $\mathbb{F}_2^n$ conditioned on

$$\begin{aligned} x_1 + \dots + x_n &= 0 \\ y_1 + \dots + y_n &= 0, \end{aligned} \tag{2}$$

and

$$H(x, y) := \sum_{i < j} x_i y_j = \sigma. \tag{3}$$

Assign share x_i, y_i to party i .

Reconstruct: Given shares $(x_k, y_k : k \neq i)$, plug them into (2), solve for x_i, y_i , and output $H(x, y)$.

³The total definition requires in addition that all sets of size strictly greater than t can reconstruct, and all sets of size strictly less than t are subject to secrecy. This can be enforced by composing with threshold schemes.

Proof of Theorem 1. Assuming the conditioning is well-defined (a solution to the constraints exists), reconstruction is immediate. As for secrecy, the probability that parties $S = [n] \setminus \{i, j\}$, where $i < j$, observe share sequence $a_S \in (\mathbb{F}_2^{n-2})^2$ (for secret σ) is

$$\Pr((xy)_S = a_S | \sigma) = \frac{\#\{xy: (xy)_S = a_S \text{ and } xy \in S_\sigma\}}{\sum_{a'_S} \#\{xy: (xy)_S = a'_S \text{ and } xy \in S_\sigma\}},$$

where S_σ is the set of solutions to (2) and (3). Under the restriction $(xy)_S = a_S$, after performing the linear substitution (2), equation (3) reduces to $(x_i + b)(y_j + c) = d$, where b, c , and d depend on a_S and σ only. This equation has either one (when $d = 1$) or three (when $d = 0$) solutions (and in particular the conditioning is well-defined), so

$$\frac{1}{2^{2(n-2)} \cdot 3} \leq \Pr((xy)_S = a_S | \sigma) \leq \frac{3}{2^{2(n-2)} \cdot 1}$$

for every a_S and σ . Plugging these bounds into (1) yields 4/5-differential secrecy. $\square$

Concrete lower bounds Proposition 7 is proved by reducing to the case $n = 5$. The lower bound is then established by computer search.

The existence of a *weak* bit secret sharing scheme for any access structure with share alphabet of size q reduces to the satisfiability of a constraint satisfaction problem with q^n variables taking 3 possible values (in time proportional to the size of the output): Nondeterministically label each of the q^n share sequences as “possible shares of 0”, “possible shares of 1”, or “not possible”, and verify the secrecy and reconstruction constraints.

For anonymous schemes, the number of admissible share sequences drops to $\binom{n+q-1}{q-1}$. When $n = 5$ and $q = 4$, there are only 56 variables. This presents no challenge for modern SAT solvers.

2 High threshold secret sharing

The elementary symmetric polynomials are given by $s_d(x_1, \dots, x_n) = \sum_S \prod_{i \in S} x_i$, where S ranges over all d -subsets of $[n]$. We view them as polynomials over the finite field $\mathbb{F}_q$.

Fix a function $H: \mathbb{F}_q \rightarrow \Sigma$, where Σ is the set of possible secrets.

$(n - t + 1)$ -out-of- n scheme

Share: Given a secret $\sigma \in \Sigma$, sample x uniformly from $\mathbb{F}_q^n$ conditioned on

$$s_d(x) = 0, \quad 1 \leq d < t, \tag{4}$$

$$H(s_t(x)) = \sigma. \tag{5}$$

Assign share x_i to party i .

Reconstruct: Given shares $a_R = (a_i: i \in R)$, $|R| = n - t + 1$, recover the remaining shares by solving the system

$$s_d(a_R, x_{\bar{R}}) = 0, \quad 1 \leq d < t \tag{6}$$

and output $H(s_t(a_R, x_{\bar{R}}))$.

A (monic) univariate polynomial p *splits completely* if all its roots are in the base field $\mathbb{F}_q$, with possible repetitions. The effect of repeated roots will be measured by the number of distinct *root arrangements*:

$$\text{ra}(p) = \begin{cases} |\{(a_{\pi(1)}, \dots, a_{\pi(t)}): \pi \in \text{Sym}_t\}|, & \text{if } p(z) = (z - a_1) \cdots (z - a_t) \text{ with } a_1, \dots, a_t \in \mathbb{F}_q, \\ 0, & \text{if } p \text{ does not split completely.} \end{cases}$$

The secrecy of the scheme is then captured by the following notion: H is ϵ -good if for every monic univariate degree- t polynomial p and every σ ,

$$\left| \frac{\mathbb{E} \text{ra}(p(z) + b) \cdot 1(H(b) = \sigma)}{t!} - \text{sp}(p) \cdot \frac{1}{|\Sigma|} \right| \leq \frac{\epsilon \cdot \text{sp}(p)}{|\Sigma|} \quad (7)$$

The expectation is over a random choice of b in $\mathbb{F}_q$. The (weighted) *splitting probability* of p is

$$\text{sp}(p) = \frac{\mathbb{E} \text{ra}(p(z) + b)}{t!}.$$

To make sense of these definitions, it is useful to pretend that the roots of $p(z) + b$ are always distinct. When q is sufficiently large, this is typically the case as repeated roots can occur for at most $t - 1$ values of b . The normalized root arrangement count can then be replaced by the indicator for the event that $p(z) + b$ splits completely. Goodness of H posits that the portion of shifts b for which $p(z) + b$ splits completely is approximately equidistributed among the fibers of H .

Claim 8. *Assuming H is ϵ -good, for every S of size $n - t$, every $a_S \in \mathbb{F}_q^S$, and σ , the number of solutions to (4)-(5) with $x_S = a_S$ is within $(1 \pm \epsilon) \text{sp}(p) \cdot t!q/|\Sigma|$ for some degree- t polynomial p that depends on a_S only.*

In particular, if $\epsilon < 1$ for every p , there is at least one solution to (4)-(5) for every σ , and the sharing is well-defined.

Claim 9. *Assuming (4) has at least one solution for every σ , for every possible a_R there is an $x_{\bar{R}}$ that is unique up to permuting $\bar{R}$ that solves (6).*

The key to both claims is the definition of the polynomial p . The construction of this polynomial relies on the closure of the elementary symmetric polynomials under restriction. Specifically, viewing $\mathbb{F}_q[x_{\bar{S}}]$ as an $\mathbb{F}_q$ -vector space, we have

Claim 10. *For every S , a_S , and $d \leq |\bar{S}|$, $s_d(a_S, x_{\bar{S}}) - s_d(x_{\bar{S}})$ is in the span of $1, s_1(x_{\bar{S}}), \dots, s_{d-1}(x_{\bar{S}})$.*

Proof. When $|S| = 1$, $s_d(a_S, x_{\bar{S}}) = s_d(x_{\bar{S}}) + a_S \cdot s_{d-1}(x_{\bar{S}})$. The claim now follows by induction on the size of S . $\square$

Proof of Claim 9. By Claim 10, for every fixed a_R the system (6) can be transformed into a (possibly inhomogeneous) lower triangular linear system in $s_1(x_{\bar{R}}), \dots, s_{t-1}(x_{\bar{R}})$. The diagonal coefficients are all 1. The remaining coefficients depend on a_R only. Solving this system we obtain an equivalent formulation of (6) given by

$$s_d(x_{\bar{R}}) = p_d, \quad 1 \leq d < t, \quad (8)$$

where $p_1, \dots, p_{t-1} \in \mathbb{F}_q$ depend on a_R only.

By Vieta's formulas, system (8) solves to $x_{\bar{R}} = a_{\bar{R}}$ if and only if the polynomial

$$p(z) = z^{t-1} - p_1 z^{t-2} + \dots + (-1)^{t-1} p_{t-1}$$

factors as $\prod_{i \in \bar{R}} (z - a_i)$. By uniqueness of factorizations there can be at most one such solution up to permuting $\bar{R}$. By assumption, a_R extends to at least one solution, so $a_{\bar{R}}$ is unique. $\square$

Proof of Claim 8. By the same argument as in the proof of Claim 9, for every fixed a_S and b the system

$$\begin{aligned} s_d(a_S, x_{\bar{S}}) &= 0, & 1 \leq d < t, \\ s_t(a_S, x_{\bar{S}}) &= b \end{aligned} \quad (9)$$

is equivalent to

$$\begin{aligned} s_d(x_{\bar{S}}) &= p_d, & 1 \leq d < t, \\ s_t(x_{\bar{S}}) &= p_t + b, \end{aligned} \quad (10)$$

where $p_1, \dots, p_t \in \mathbb{F}_q$ depend on a_S only.

By Vieta's formulas, each solution $x_{\bar{S}} = a_{\bar{S}}$ to system (10) corresponds to an ordered factorization $\prod_{i \in \bar{S}} (z + a_i)$ of the polynomial $p(z) + b$, where

$$p(z) = z^t + p_1 z^{t-1} + \dots + p_t$$

The number of solutions therefore equals the number of distinct arrangements of the roots of $p(z) + b$ if this polynomial splits completely, and zero otherwise. This is precisely $\text{ra}(p(z) + b)$.

As (9) and (10) are equivalent, the number of completions $x_{\bar{S}}$ that solve (9) for some b satisfying $H(b) = \sigma$ equals $\text{ra}(p(z) + b) \cdot 1(H(b) = \sigma)$ summed over all b . As H is ϵ -good, this number is within $1 \pm \epsilon$ of $\text{sp}(p) \cdot t!q/|\Sigma|$. $\square$

Claim 11. *Assuming $\epsilon < 1$ and H is ϵ -good, this scheme admits perfect reconstruction and $2\epsilon/(1 + \epsilon^2)$ -differential secrecy.*

Proof. As $\epsilon < 1$, by Claim 9 every possible view of $n - t + 1$ parties admits unique recovery of the remaining shares. By (5) the output of the reconstruction algorithm must equal the secret σ .

As for secrecy, for a given secret σ , the probability of observing shares a_S is proportional to the fraction of the set S_σ of solutions to (4)-(5):

$$\Pr(x_S = a_S | \sigma) = \frac{\#\{x: x_S = a_S \text{ and } x \in S_\sigma\}}{\sum_{a'_S} \#\{x: x_S = a'_S \text{ and } x \in S_\sigma\}}$$

Fix two secrets 0 and 1. The likelihood ratio is

$$\ell = \frac{\Pr(x_S = a_S | 0)}{\Pr(x_S = a_S | 1)} = \frac{\#\{x: x_S = a_S \text{ and } x \in S_0\}}{\#\{x: x_S = a_S \text{ and } x \in S_1\}} \cdot \frac{\sum_{a'_S} \#\{x: x_S = a'_S \text{ and } x \in S_1\}}{\sum_{a'_S} \#\{x: x_S = a'_S \text{ and } x \in S_0\}}.$$

By Claim 8, each of these ratios is between $(1 - \epsilon)/(1 + \epsilon)$ and its reciprocal. The differential secrecy parameter is

$$\frac{|1 - \ell|}{1 + \ell} \leq \frac{1 - ((1 - \epsilon)/(1 + \epsilon))^2}{1 + ((1 - \epsilon)/(1 + \epsilon))^2} = \frac{2\epsilon}{1 + \epsilon^2}. \quad \square$$

2.1 Splitting probability of polynomials

To gain some intuition let us start with the splitting probability of quadratics over $\mathbb{F}_q$ with $q > 2$ prime. After a possible shift of input and output, every (monic) quadratic reduces to x^2 . The splitting probability of x^2 then equals the (weighted) fraction of quadratic residues, which is exactly $1/2$. The reason is that non-zero roots of any quadratic residue come in pairs $\{x, -x\}$.

More generally, the splitting probability of x^d over $\mathbb{F}_q$ is $1/d - 1/dq + 1/d!q$ if d divides $q - 1$, and $1/d!q$ otherwise. The reason is that the function x^d is regular over the multiplicative group $\mathbb{F}_q^*$. Unless this function is d -to-1, $x^d - b$ splits only when $b = 0$, which accounts for the $1/d!q$ term. The function is d -to-1 if and only if $\mathbb{F}_q$ contains all d -th roots of unity. As the roots are a subgroup, d must divide $q - 1$. When d does divide $q - 1$, the roots are the powers of $g^{(q-1)/d}$ where g is the generator of $\mathbb{F}_q^*$. In that case there are $(q - 1)/d$ power- d residues in $\mathbb{F}_q^*$, accounting for the $1/d$ factor.

This example shows that when d does not divide $q - 1$, there are example polynomials with splitting probability as small as $1/q$. In these cases, the goodness assumption in Claim 8 cannot be enforced. This appears to be the only obstruction to a universal bound independent of q on the splitting probability:

Conjecture 12. *For every $d \geq 1$ and all sufficiently large primes q such that d divides $q - 1$, the splitting probability of every monic degree- d polynomial p over $\mathbb{F}_q$ is lower bounded by some constant $c_d > 0$ that depends on d only.*

We prove Conjecture 12 in the cases $d = 2$ and $d = 3$ and justify its plausibility in general in Appendix A. Using the Chebotarev density theorem for function fields [Ros13, Theorem 9.13B], Conjecture 12 reduces to Conjecture 23, a purely algebraic proposition about constants in Galois extensions of the field of $\mathbb{F}_q$ -rational functions. A proof or a counterexample may well exist in the relevant literature.

We do not know of an effective bound on c_d , though we suspect that $c_d \geq 1/(d!)^{O(1)}$ can be deduced from known results if the conjecture is true. The constant c_d cannot be larger than $1/d!$ asymptotically.

The reason is that the number of $\mathbb{F}_q$ -polynomials that split into distinct factors is $\binom{q}{d}$, which is less than a $1/d!$ -fraction of the total. As q tends to infinity, distinct factors are typical. Thus in expectation over a random choice of polynomial, the probability of a complete split tends to

$$\frac{\text{number of splits into distinct factors}}{\text{number of monic degree-}d \text{ polynomials}} = \frac{\binom{q}{d}}{q^d} \rightarrow \frac{1}{d!}.$$

2.2 Existence of good colorings

Conjecture 12 provides a universal lower bound on the splitting probability. We now turn to the study of good colorings H for which the splitting probabilities are approximately equidistributed among the fibers of H .

Let us again start with the case of quadratics ($d = 2$) over the field $\mathbb{F}_5$ specifically. The splitting probability of every quadratic p is $1/2$. The set of values $b \in \mathbb{F}_5$ for which $p(z) + b$ splits is a shift of the quadratic residues $-1, 0, 1$, namely is a sequence of three consecutive elements. The coloring $H(0) = H(2) = 1$, and $H(1) = H(3) = H(4) = 0$, for instance, ensures that there is at least one complete split of $p(z) + b$ for both $\sigma = 0$ and $\sigma = 1$. This coloring is therefore $3/5$ -good. Applying Claim 11, this proves Corollary 6.

Asymptotically, approximate equidistribution can be enforced by picking H at random.

Proof of Theorem 2. Let H be a random coloring. Every $b \in \mathbb{F}_q$ is assigned a value $H(b)$ in Σ uniformly and independently. For a fixed p and σ , the random variable

$$Z = \sum_b \frac{\text{ra}(p(z) + b)}{t!} \cdot 1(H(b) = \sigma)$$

is a sum of independent, bounded samples. Its expectation is $\mathbb{E} Z = \text{sp}(p) \cdot q/|\Sigma|$. The number of nonzero terms in the sum is at most $\sum_b \text{ra}(p(z) + b) = t! \text{sp}(p)q$, and each has variance at most $1/|\Sigma|$. The variance of Z is therefore at most $t! \text{sp}(p) \cdot q/|\Sigma|$. By Bernstein's inequality, the probability that (7) is violated for this p and σ is at most $2 \exp(-\Omega(\epsilon^2 \text{sp}(p)q/t!|\Sigma|))$.

By Conjecture 12 this is at most $2 \exp(-c_t \epsilon^2 q/|\Sigma|)$ for a constant $c_t > 0$ that depends on t only. By a union bound, an ϵ -good H exists provided that $4|\Sigma|q^t \cdot \exp(-c_t \epsilon^2 q/|\Sigma|) < 1$. This is satisfied for $q = O_t(|\Sigma|/\epsilon^2 \cdot \log(|\Sigma|/\epsilon))$, provided a prime number congruent to 1 modulo t of this magnitude exists. That is guaranteed by Dirichlet's density theorem for arithmetic progressions. Reconstruction and secrecy then follow from Claim 11. $\square$

3 Statistically secure sharing for the top slice

An $(n - 2)$ -out-of- n slice can be described by a graph G on n vertices so that if ij is not an edge then all parties except i, j can reconstruct, and if ij is an edge then those parties cannot distinguish secrets. Let q be an odd prime power. Fix a function $H : \mathbb{F}_q^n \rightarrow \{0, 1\}$.

$(n - 2)$ -out-of- n slice scheme

Share: Given a secret $\sigma \in \{0, 1\}$, sample x uniformly from $\mathbb{F}_q^n$ conditioned on $x_1 + \dots + x_n = 0$ and $H(x) = \sigma$. Assign share x_i to party i . Publish $\sum_{ij \in G} x_i x_j$.

Reconstruct: Given shares $(x_k : k \neq i, j)$, plug them in the above equations and solve for x_i, x_j . If the solution is unique, output $H(x)$.

We say G is nondegenerate if no two vertices have the same set of neighbors. Let A^c denote the complement of set A .

Proposition 13. *Assume G is nondegenerate and n is sufficiently large. There exists H so that for every $ij \notin G$, parties in $\{i, j\}^c$ can reconstruct except with probability $O(1/q)$, and for every $ij \in G$, parties in $\{i, j\}^c$ have distinguishing advantage $O(1/q)$.*

Proof. By default, we consider x to be uniformly distributed over $\mathbb{F}_q^n$ conditioned on $x_1 + \dots + x_n = 0$. Since there are q^{n-1} such x , by a Chernoff bound, a random coloring H is likely to be almost balanced for sufficiently large n . Specifically,

Claim 14. *Except with probability less than $1/3$ over the choice of H ,*

$$\left| \Pr_x(H(x) = 0) - \frac{1}{2} \right| < \frac{1}{2q}. \quad (11)$$

We first consider the reconstruction error. Fix any $ij \notin G$. Consider the view of the qualified set $S = \{i, j\}^{\mathbb{G}}$, which consists of the shares x_S of those parties and the published value $\sum_{ij \in G} x_i x_j$. From the constraint $x_1 + \dots + x_n = 0$, these parties learn the value of $x_i + x_j$. From the published value $\sum_{ij \in G} x_i x_j$, these parties learn the value $a_i x_i + a_j x_j$, where $a_i = a_i(x_S) = \sum_{k \in N(i)} x_k$, $a_j = a_j(x_S) = \sum_{k \in N(j)} x_k$, and $N(i)$, $N(j)$ are the sets of neighbors of i , j , respectively. The reconstruction fails only if $a_i = a_j$.

Claim 15. *For any $ij \notin G$, except with probability less than $1/3n^2$ over the choice of H ,*

$$\left| \Pr_x(a_i = a_j \wedge H(x) = 0) - \frac{1}{2q} \right| < \frac{1}{2q^2}, \quad (12)$$

Proof. Since G is nondegenerate, we have $N(i) \neq N(j)$, so for any $ij \notin G$, $\Pr_x[a_i = a_j] = 1/q$. Namely, there are q^{n-2} share sequences x such that $a_i = a_j$. By a Chernoff bound, (12) fails to hold with probability at most $2 \cdot \exp(-q^{n-4}/6)$, which is smaller than $1/3n^2$ for sufficiently large n . $\square$

We now consider the distinguishing advantage. Fix any $ij \in G$. Consider the view of the parties in the unqualified set $S = \{i, j\}^{\mathbb{G}}$. From the constraint $x_1 + \dots + x_n = 0$, these parties learn that $x_i + x_j$ equals $b = b(x_S) = -\sum_{k \in S} x_k$. From the published value $\sum_{ij \in G} x_i x_j$, these parties learn $x_i x_j + cx_i + dx_j$, where $c = c(x_S) = \sum_{k \in N(i) - \{j\}} x_k$ and $d = d(x_S) = \sum_{k \in N(j) - \{i\}} x_k$. Substituting $b - x_i$ for x_j , the parties learn the value of the following quadratic function of x_i :

$$x_i(b - x_i) + cx_i + d(b - x_i) = -x_i^2 + (b + c - d)x_i + bd,$$

where b, c, d are constants that depend on x_S only. Unless $x_i = (b + c - d)/2$, the view of S is consistent with two distinct share sequences x and x' . Let M be the involution that swaps these two sequences. (Specifically, $M(x)_i = b + c - d - x_i$, $M(x)_j = b + d - c - x_j$, and $M(x)_S = x_S$.) There is a reasonable chance that x and $M(x)$ arise from different secrets:

Claim 16. *For any $ij \in G$, except with probability less than $1/3n^2$ over the choice of H ,*

$$\left| \Pr_x(H(x) \neq H(M(x))) - \frac{1}{2} \right| < \frac{1}{q}, \quad (13)$$

Proof. Unless $x_i = (b + c - d)/2$, there are $q^{n-2} \cdot (q-1)/2$ pairs of x, x' such that $x \neq x'$ and $x' = M(x)$. The x and x' in each such pair are assigned different colors independently with probability $1/2$. The event $H(x) \neq H(M(x))$ means that x belongs to such a pair where $M(x)$ has a different color. The claim considers the probability over H such that at most a $\frac{1}{2}(1 - \frac{1}{q-1})$ or at least a $\frac{1}{2}(1 + \frac{1}{q-1})$ fraction of these pairs are colored differently. By a Chernoff bound, (13) fails to hold with probability at most $2 \cdot \exp(-q^{n-2}/6(q-1))$, which is smaller than $1/3n^2$ for sufficiently large n . $\square$

By a union bound, there exists an H for which all of (11), (12), and (13) hold simultaneously for all relevant pairs ij . We fix this H .

From (11) and (12), for any $ij \notin G$, the reconstruction error for the parties in $\{i, j\}^{\mathbb{G}}$ is

$$\Pr(a_i = a_j \mid H(x) = \sigma) = \frac{\Pr(a_i = a_j \wedge H(x) = \sigma)}{\Pr(H(x) = \sigma)} < \frac{1/2q + 1/2q^2}{1/2 - 1/2q} = O(1/q).$$

For secrecy, let V be the possible views of $S = \{i, j\}^{\mathbb{G}}$ that are consistent with both secrets 0 and 1. For each view $v \in V$ and each secret $\sigma \in \{0, 1\}$, there is exactly one share sequence x that is compatible

with v and σ . Owing to the imbalance (11) in H , these share sequences are not equally likely given σ , but their likelihood ratio is close to one:

$$\left| 1 - \frac{\Pr(v|1)}{\Pr(v|0)} \right| \leq \left| 1 - \frac{1 + 1/q}{1 - 1/q} \right| = \frac{2}{q - 1}.$$

The contribution of V to the statistical distance between S -views of 0 and 1 is bounded by

$$\frac{1}{2} \sum_{v \in V} |\Pr(v|0) - \Pr(v|1)| = \frac{1}{2} \sum_{v \in V} \Pr(v|0) \cdot \left| 1 - \frac{\Pr(v|1)}{\Pr(v|0)} \right| \leq \frac{1}{q - 1}.$$

The contribution of $V^{\mathcal{C}}$ is at most

$$\frac{1}{2} \sum_{v \notin V} (\Pr(v|0) + \Pr(v|1)) = \frac{\Pr(V^{\mathcal{C}}|0) + \Pr(V^{\mathcal{C}}|1)}{2}. \quad (14)$$

Instead of upper bounding $\Pr(V^{\mathcal{C}}|\sigma)$, we lower bound $\Pr(V|\sigma)$. From Bayes' rule, (11) and (13),

$$\begin{aligned} \Pr(V|\sigma) &= \Pr(H(x) \neq H(M(x)) | H(x) = \sigma) \\ &= \frac{\Pr(H(x) = \sigma | H(x) \neq H(M(x))) \cdot \Pr(H(x) \neq H(M(x)))}{\Pr(H(x) = \sigma)} \\ &> \Pr(H(x) = \sigma | H(x) \neq H(M(x))) \cdot \frac{1/2 - 1/q}{1/2 + 1/2q}. \end{aligned}$$

By symmetry, conditioned on $H(x) \neq H(M(x))$, the events $H(x) = \sigma$ and $H(M(x)) = \sigma$ are equally likely. As they are disjoint and span the sample space, each must have probability $1/2$. Therefore

$$\Pr(V|\sigma) > \frac{1}{2} \cdot \frac{1 - 2/q}{1 + 1/q} \geq \frac{1}{2} - \frac{3}{2q}.$$

Plugging into (14), the contribution of $V^{\mathcal{C}}$ is less than $1/2 + 3/2q$. The distinguishing advantage is therefore $1/2 + O(1/q)$ as desired. $\square$

The nondegeneracy assumption incurs no loss of generality.

Claim 17. *If there exists a scheme for every $(n - 2)$ -out-of- n nondegenerate slice in which every $(n - 1)$ -subset can also reconstruct, so does one with the same parameters without the nondegeneracy assumption.*

Proof. Call a vertex degenerate if it has the same neighbors as some other vertex in G . The claim is proved by induction on the number of degenerate vertices. For the base case, suppose there is a single pair of vertices i, j with the same set of neighbors. In particular, $\{i, j\}$ is not an edge so $\{i, j\}^{\mathcal{C}}$ is authorized.

Let S be the scheme for the nondegenerate slice obtained by merging i and j into a single party ij . The scheme S' assigns the same shares as S to all parties except i and j , and 2-out-of-2 shares of party ij 's share σ_{ij} in S to i and j in S' . Namely, i and j are given random shares that add up to σ_{ij} . The analysis splits into three cases depending on the identity of the missing parties $\{a', b'\}$ in S' . Let

$$a = \begin{cases} ij, & \text{if } a' = i \text{ or } a' = j, \\ a', & \text{otherwise,} \end{cases} \quad \text{and} \quad b = \begin{cases} ij, & \text{if } b' = i \text{ or } b' = j, \\ b', & \text{otherwise.} \end{cases}$$

We will show that $\{a', b'\}^{\mathcal{C}}$ qualifies in S' if and only if $\{a, b\}^{\mathcal{C}}$ qualifies in S . When $a' = b'$, this implies all sets missing one party qualify in S' as long as they qualify in S . When $a' \neq b'$, this means the graph specifying the qualified sets in S' is derived from that of S by cloning vertex ij into vertices i and j , as desired.

- If $\{a', b'\}$ and $\{i, j\}$ do not intersect, then parties $\{a', b'\}^{\mathcal{C}}$ in S' have the same information about the secret that parties $\{a, b\}^{\mathcal{C}}$ have in S . Specifically, $i, j \in \{a', b'\}^{\mathcal{C}}$ learn the share of $ij \in \{a, b\}^{\mathcal{C}}$ and nothing more. Therefore one set qualifies if and only if the other one does.

- If $\{a', b'\}$ and $\{i, j\}$ intersect, and $\{a, b\}^G$ can reconstruct in S , then $\{a', b'\}^G$ can reconstruct in S' . The reason is that ij belongs to $\{a, b\}$, so its share is not used for reconstruction in S . Therefore i 's and j 's shares are not needed for reconstruction in S' . All other shares are available.
- If $\{a', b'\}$ and $\{i, j\}$ intersect, and $\{a, b\}^G$ does not qualify in S , then $\{a', b'\}^G$ does not qualify in S' . The shares of $\{a', b'\}^G$ in S' can be simulated from those in $\{a, b\}^G$ in S . The reason is that at most one of i and j can appear in $\{a', b'\}^G$, so its share can be simulated by a random value. All other shares are inherited from S .

The inductive step is analogous. Merging i and j reduces the number of degenerate vertices by one, and the scheme is extended in the same manner. $\square$

Proof of Theorem 5. By Claim 17, we may assume without loss of generality that G is non-degenerate.

We modify the scheme from Proposition 13 to enforce perfect reconstruction: Sample x under the additional condition that for every $ij \notin G$, $[n] - \{i, j\}$ can reconstruct. By the union bound, for any $\sigma \in \{0, 1\}$, x satisfies this condition except with probability $O(n^2)/q$. Sampling x under this condition changes the view of any unqualified set by a statistical distance of at most $O(n^2/q)$, so the scheme is $2/3$ -statistically secret with $q = O(n^2)$. It can then be amplified to be ϵ -statistically secret with alphabet size $q^{O(\log 1/\epsilon)} = n^{O(\log 1/\epsilon)}$ as will be shown in Proposition 18 in Section 4. $\square$

4 Amplification

The following procedure amplifies secrecy for bit sharing schemes with perfect reconstruction.

Amplification

Input: A bit secret sharing scheme $(Share, Rec)$.

Share: Given secret σ , sample bits $\sigma_1, \dots, \sigma_t$ conditioned on $\sigma_1 + \dots + \sigma_t = \sigma$. Party j is assigned the sequence $(Share(\sigma_1)_j, \dots, Share(\sigma_t)_j)$ as its share.

Reconstruct: Given a qualified set R with shares $x_{ij} : i \in R, j \in [t]$, output $\sum_j Rec(x_{ij} : i \in R)$.

Amplification blows up share size by a factor of t . Equivalently, if $(Share, Rec)$ has share alphabet Σ , the amplified scheme has share alphabet Σ^t .

Amplification clearly preserves perfect reconstruction. As for secrecy,

Proposition 18. *If $(Share, Rec)$ is ϵ -statistically (resp. differentially) secret, then $Ampl$ is ϵ^t -statistically (resp., differentially) secret.*

Statistical secrecy amplification is essentially the polarization lemma of Sahai and Vadhan [SV99]. Both statistical and differential secrecy follow easily from the following alternative definitions. We assume σ is a random bit and x is the corresponding share sequence.

Claim 19. *A bit scheme is ϵ -statistically secret if and only if for every unqualified set S and every real-valued function A , $|\mathbb{E}[A(x_S) \cdot (-1)^\sigma]| \leq \epsilon \|A\|_\infty$.*

Proof. We use the notation $A \rightarrow S$ for an S -valued function A . The statistical secrecy equals

$$\begin{aligned}
\max_{A \rightarrow \{0,1\}} (\Pr(A(x_S)|\sigma = 0) - \Pr(A(x_S)|\sigma = 1)) &= 2 \max_{A \rightarrow \{0,1\}} \mathbb{E} A(x_S)(-1)^\sigma \\
&= 2 \max_{B \rightarrow \{-1,1\}} \mathbb{E} \frac{1 + B(x_S)}{2} \cdot (-1)^\sigma \\
&= \max_{B \rightarrow \{-1,1\}} \mathbb{E} B(x_S)(-1)^\sigma \\
&= \sup_{B \rightarrow [-1,1]} \mathbb{E} B(x_S)(-1)^\sigma \\
&= \sup_{B \rightarrow [-1,1]} |\mathbb{E} B(x_S)(-1)^\sigma| \\
&= \frac{1}{\|B\|_\infty} \sup_{B \rightarrow \mathbb{R}} \mathbb{E} B(x_S)(-1)^\sigma.
\end{aligned}$$

The fourth equality is true because a $[-1, 1]$ valued distinguisher can be realized as a mixture of $\{-1, 1\}$ -valued ones, and the optimum advantage is attained by one of the components of this mixture. $\square$

Claim 20. *A bit scheme is ϵ -differentially secret if and only if for every unqualified set S and every possible share sequence a_S ,*

$$\frac{|\mathbb{E} 1(x_S = a_S)(-1)^\sigma|}{\Pr(x_S = a_S)} \leq \epsilon.$$

Proof. After halving both numerator and denominator, the left-hand side of (1) becomes

$$\frac{|\mathbb{E} \mathbb{E}[1(x_S = a_S)(-1)^\sigma | \sigma]|}{\mathbb{E} \Pr(x_S = a_S | \sigma)}.$$

Apply the law of total expectation. $\square$

Proof of Proposition 18. When σ is a random bit, the partial views $x_{S1}, \dots, x_{St}$ corresponding to $\sigma_1, \dots, \sigma_t$ of every set S of parties, in particular an unqualified one, are (unconditionally) independent.

For statistical secrecy, we show that

$$|\mathbb{E} A \cdot (-1)^{\sigma_1 + \dots + \sigma_t}| \leq \epsilon^t \|A\|_\infty$$

for every A that is determined by $x_{S1}, \dots, x_{St}$. The proof is by induction on t . The base case $t = 1$ holds by assumption. For the inductive step, let $B = \mathbb{E}[A \cdot (-1)^{\sigma_2 + \dots + \sigma_t} | x_{S1}]$. By inductive hypothesis, $\|B\|_\infty \leq \epsilon^{t-1} \|A\|_\infty$. By assumption,

$$|\mathbb{E} A \cdot (-1)^{\sigma_1 + \dots + \sigma_t}| = |\mathbb{E} B \cdot (-1)^{\sigma_1}| \leq \epsilon \|B\|_\infty \leq \epsilon^t \|A\|_\infty.$$

For differential secrecy, by independence both numerator and denominator factor out:

$$\frac{|\mathbb{E} 1(x_S = a_S)(-1)^\sigma|}{\Pr(x_S = a_S)} = \frac{\prod_{j=1}^t |\mathbb{E} 1(x_{Sj} = a_{Sj})(-1)^{\sigma_j}|}{\prod_{j=1}^t \Pr(x_{Sj} = a_{Sj})} \leq \epsilon^t. \quad \square$$

5 Lower bounds

We prove the lower bounds Proposition 3, 4, and 7.

Proof of Proposition 3. The proof is an extension of Csirmaz's argument [Csi97]. Assume the secret σ is uniform. Let a be the share of the relevant party i and s be the shares of the unqualified set that becomes qualified upon including party i . By perfect reconstruction, $H(\sigma as) = H(as)$. H denotes entropy.

If secrecy were perfect, σ and s would be independent, and we would also have $H(\sigma s) = H(\sigma) + H(s)$. Imperfection entails a small correction:

$$H(\sigma s) \geq H(\sigma) + H(s) - \log \frac{1 + \epsilon}{1 - \epsilon}. \quad (15)$$

The reason is that by the uniformity of σ , the mutual information $I(\sigma : s) = H(\sigma) + H(s) - H(\sigma s)$ equals the average KL-divergence $\mathbb{E} KL(\sigma | s, \sigma)$. By Bayes' rule and ϵ -differential secrecy, the maximum divergence is at most

$$\log \frac{\Pr(\sigma | s)}{\Pr(\sigma)} = \log \frac{\Pr(s | \sigma)}{\mathbb{E} \Pr(s | \sigma')} \leq \max_{\sigma, \sigma'} \log \frac{\Pr(s | \sigma)}{\Pr(s | \sigma')}.$$

By (1),

$$\left| \frac{1 - \Pr(s | \sigma) / \Pr(s | \sigma')}{1 + \Pr(s | \sigma) / \Pr(s | \sigma')} \right| \leq \epsilon,$$

from where we obtain (15) with some algebra. The proof is completed using Shannon's entropy inequalities:

$$H(a) \geq H(a | s) = H(as) - H(s) = H(\sigma as) - H(s) \geq H(\sigma s) - H(s) \geq H(\sigma) - \log \frac{1 + \epsilon}{1 - \epsilon}. \quad \square$$

Proof of Proposition 4. The proof is an extension of Kilian and Nisan’s argument. Let $X_1 \cdots X_n$ and $Y_1 \cdots Y_n$ be independent share sequences of 0 and 1. If q is the share alphabet size, then

$$1 \geq \mathbb{E} \#\{i : X_i = Y_i\} = \sum_{i=1}^n \Pr(X_i = Y_i),$$

by the reconstruction requirement and linearity of expectation, respectively. By secrecy, X_i and Y_i are ϵ -close, so there exist identically distributed random variables X'_i, Y'_i that independently equal X_i, Y_i , respectively, with probability $1 - \epsilon$ each. Then $\Pr(X'_i = Y'_i)$ is a collision probability for a random variable of support size at most q , so it is lower bounded by $1/q$. By the law of total probability,

$$\begin{aligned} \Pr(X_i = Y_i) &\geq \Pr(X_i = Y_i \mid X_i = X'_i, Y_i = Y'_i) \cdot \Pr(X_i = X'_i, Y_i = Y'_i) \\ &= \Pr(X'_i = Y'_i) \cdot \Pr(X_i = X'_i) \cdot \Pr(Y_i = Y'_i) \\ &\geq \frac{1}{q} \cdot (1 - \epsilon) \cdot (1 - \epsilon), \end{aligned}$$

from where $q \geq (1 - \epsilon)^2 n$. $\square$

Proof of Proposition 7. The maximum share size of a weak $(k - 1)$ -out-of- $(n - 1)$ scheme P is no worse than that of a weak k -out-of- n scheme P' . Given P' , P is the scheme obtained by fixing some possible share of party n , call it 0, and assigning the remaining shares in any way consistent with this fixing. As $\{n\}$ is an unauthorized set for P' , there must exist at least one possible share sequence for every secret.

Reconstruction in P for authorized set R works by simulating P' on authorized set $R \cup \{n\}$ with party n receiving share 0. If any $k - 1$ shares a_R in P were consistent with two secrets, then so would be the shares $a_R 0$ for P' , violating reconstructibility.

As for secrecy, any possible view a_S of an unauthorized set S in P must be consistent with some shares of all secrets because $S \cup \{n\}$ is unauthorized for P' and the view $a_S 0$ is possible.

In particular, a lower bound on the maximum share size for 4-out-of-5 bit secret sharing extends to a lower bound for $(n - 1)$ -out-of- n for all $n \geq 5$. A lower bound of 5 is proved by computer search as described in the introduction.⁴ $\square$

6 Conclusion and open problems

The unreasonable effectiveness of linear secret sharing is powered by the equivalence of linear and statistical (in)dependence. Nonlinear proposals are in general more difficult to analyze. Our results suggest that “low-degree” nonlinear constructions sometimes exhibit approximate statistical independence.

For our top slice scheme, the guarantees of Proposition 13 are quite weak. Reconstruction usually works, while secrecy holds only about half the time. Nevertheless, these can be amplified to obtain bona fide statistical security with reasonable information ratio. It would be interesting to study if stronger amplification (e.g., composition with a threshold scheme) can lead to further improvements in share size.

Another natural direction to explore is extensions to other high slices. Disregarding technical obstacles, one would expect that generalizing our scheme in Section 3 to $(n - t + 1)$ -out-of- n slices over a sufficiently large alphabet should allow qualified sets to rule out all but $t - 1$ possible secrets, while unqualified sets would usually be consistent with t possible secrets. Can this weak gap be amplified into a statistical secrecy-reconstruction gap?

Reasoning about unrealistically weak notions of secrecy, coupled with the support of AI-powered coding tools, opens up new opportunities. Our Proposition 7 was discovered by computer search, leading to a matching construction (Corollary 6). Computer search further showed that the optimal alphabet size is exactly 4 for plain (not anonymous) 3-out-of-6 weak bit secret sharing, and is at least 4 for 4-out-of-6. Ruling out alphabet size 4 for 4-out-of-6 weak bit schemes would prove that weak bit secret sharing is not closed under duality. This appears to be beyond the capabilities of unoptimized computer search on standard equipment.

⁴<https://github.com/yanbo-chen/search-for-weak-threshold-secret-sharing>. The code was generated by Google Gemini.

Acknowledgements

We thank Amos Beimel for raising the question of duality in weak secret sharing and for his pivotal help in reconstructing some forgotten insights from his work with Franklin [BF07]. We also thank the TCC 2026 reviewers for pointing out two errors in a previous version of the manuscript.

References

- [AA20] Benny Applebaum and Barak Arkis. On the power of amortization in secret sharing: d -Uniform secret sharing and CDS with constant information rate. *ACM Trans. Comput. Theory*, 12(4):24:1–24:21, 2020.
- [ABF⁺19] Benny Applebaum, Amos Beimel, Oriol Farràs, Oded Nir, and Naty Peter. Secret-sharing schemes for general and uniform access structures. In *Advances in Cryptology, EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 11478 of *Lecture Notes in Computer Science*, pages 441–471, Cham, 2019. Springer-Verlag.
- [BF07] Amos Beimel and Matthew Franklin. Weakly-private secret sharing schemes. In Salil P. Vadhan, editor, *Proceedings of the 4th Conference on Theory of Cryptography*, volume 4392 of *Lecture Notes in Computer Science*, pages 253–272, Berlin, Heidelberg, 2007. Springer-Verlag.
- [BFLN24] Amos Beimel, Oriol Farràs, Or Lasri, and Oded Nir. Secret-sharing schemes for high slices. In *Theory of Cryptography: 22nd International Conference, TCC 2024*, volume 15367 of *Lecture Notes in Computer Science*, pages 581–613, Cham, 2024. Springer-Verlag.
- [BGK20] Andrej Bogdanov, Siyao Guo, and Ilan Komargodski. Threshold secret sharing requires a linear-size alphabet. *Theory of Computing*, 16(2):1–18, 2020.
- [BIKL25] Amos Beimel, Yuval Ishai, Eyal Kushilevitz, and Hanjun Li. Cryptography with weak privacy. In Benny Applebaum and Huijia Lin, editors, *Theory of Cryptography - 23rd International Conference, TCC 2025, Aarhus, Denmark, December 1-5, 2025, Proceedings, Part IV*, volume 16271 of *Lecture Notes in Computer Science*, pages 415–447. Springer, 2025.
- [Bla79] G. R. Blakley. Safeguarding cryptographic keys. In *Proceedings of the 1979 AFIPS National Computer Conference*, pages 313–317, Montvale, NJ, USA, 1979. AFIPS Press.
- [Bog23] Andrej Bogdanov. Csirmaz’s duality conjecture and threshold secret sharing. In Kai-Min Chung, editor, *4th Conference on Information-Theoretic Cryptography (ITC 2023)*, volume 267 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:6, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [BP18] Amos Beimel and Naty Peter. Optimal linear multiparty conditional disclosure of secrets protocols. In *Proceedings of the 24th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*, volume 11274 of *Lecture Notes in Computer Science*, pages 332–362, Cham, 2018. Springer-Verlag.
- [Car53] L. Carlitz. Permutations in a finite field. *Proceedings of the American Mathematical Society*, 4(4):538, 1953.
- [Csi97] László Csirmaz. The size of a share must be large. *Journal of Cryptology*, 10(4):223–231, 1997.
- [Csi21] László Csirmaz. Secret sharing and duality. *J. Math. Cryptol.*, 15(1):157–173, 2021.
- [Gál95] Anna Gál. *Combinatorial Methods in Boolean Function Complexity*. PhD thesis, University of Chicago, 1995.
- [KN90] Joe Kilian and Noam Nisan, 1990. Unpublished.
- [Lok09] Satyanarayana V. Lokam. *Complexity Lower Bounds Using Linear Algebra*. Now Publishers Inc., Hanover, MA, USA, 2009.

- [LV18] Tianren Liu and Vinod Vaikuntanathan. Breaking the circuit-size barrier in secret sharing. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 699–708. Association for Computing Machinery, 2018.
- [LVW17] Tianren Liu, Vinod Vaikuntanathan, and Hoeteck Wee. Conditional disclosure of secrets via non-linear reconstruction. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 758–790. Springer, 2017.
- [Ros13] Michael Rosen. *Number Theory in Function Fields*, volume 210 of *Graduate Texts in Mathematics*. Springer, 2013.
- [Sha79] Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, Nov. 1979.
- [SV99] Amit Sahai and Salil Vadhan. Manipulating statistical difference. In Panos Pardalos, Sanguthevar Rajasekaran, and José Rolim, editors, *Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997)*, volume 43 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 251–270. American Mathematical Society, 1999.

A Discussion of Conjecture 12

We begin by proving Conjecture 12 for degrees 2 and 3. We then discuss its relation to permutation polynomials and to Chebotarev’s density theorem for function fields.

It will be more convenient to discuss the unweighted splitting probability, namely the fraction of shifts b for which $p(z) - b$ splits completely, disregarding the effect of repeated roots. The weighting does not affect the validity of the conjecture because repeated roots of $p(z) - b$ must also be roots of the derivative $p'(z)$, so they can occur for at most $d - 1$ values of b .

A.1 Special cases

Proof of Conjecture 12 for $d = 2$. By a linear shift, the splitting probability of every degree-2 polynomial modulo q is the same as that of x^2 (assuming $q \neq 2$). The polynomial $x^2 - b$ splits if and only if b is a quadratic residue, which happens with probability $1/2 + 1/2q \geq 1/2$. $\square$

Proof of Conjecture 12 for $d = 3$. By a linear shift, we may assume the degree-3 polynomial is of the form $p(x) = x^3 - ax$. We will argue that for a random choice of x_0 , $p(x) - p(x_0)$ splits completely with probability at least $1/2 - 1/2q$. As each value $p(x_0)$ can be taken at most three times, it follows that $p(x) - b$ splits completely with probability at least $1/6 - 1/6q$.

The polynomial $p(x) - p(x_0)$ factors as

$$p(x) - p(x_0) = (x - x_0)(x^2 + x_0x + x_0^2 - a).$$

The remaining two roots are in $\mathbb{F}_q$ if and only if the discriminant $x_0^2 - 4(x_0^2 - a) = 4a - 3x_0^2$ is a quadratic residue. By Gauss’s quadratic reciprocity law, -3 is a quadratic residue modulo q when $q \equiv 1 \pmod{3}$:

$$\left(\frac{-3}{q}\right) = \left(\frac{-1}{q}\right)\left(\frac{3}{q}\right) = (-1)^{(q-1)/2} \cdot (-1)^{(q-1)/2}\left(\frac{q}{3}\right) = \left(\frac{1}{3}\right) = 1.$$

Every x_0 of the form $t/2 + 2at^{-1}/3$ is then a good choice because

$$4a - 3\left(\frac{t}{2} + \frac{2at^{-1}}{3}\right)^2 = \left(\frac{\sqrt{-3} \cdot t}{2} - \frac{2\sqrt{-3} \cdot at^{-1}}{3}\right)^2.$$

As t ranges over $\mathbb{F}_q^*$, $t/2 + 2at^{-1}/3$ takes every value at most twice, resulting in at least $(q - 1)/2$ good choices for x_0 . $\square$

A.2 Permutation polynomials

Conjecture 12 in particular postulates, for every polynomial p , the *existence* of a value taken d times by p , provided d divides $q - 1$ (and q is sufficiently large). As we will see shortly, these two statements—the existence and positive density of complete splits—are equivalent, provided factorizations with repeated roots are discounted.

An even weaker consequence is that some value should be taken at least *twice* by p . This was proved by Carlitz [Car53]. We reproduce his argument.

Proposition 21. *If $d > 1$ divides $q - 1$ and q is prime, every degree- d polynomial over $\mathbb{F}_q$ takes some value at least twice.*

Proof. Let g be a generator of $\mathbb{F}_q^*$. For every $t \leq q - 2$,

$$\sum x^t = \sum_{i=0}^{q-1} g^{it} = (g^{t(q-1)} - 1)/(g^t - 1) = 0,$$

so $\sum r(x) = 0$ for every r of degree at most $q - 2$. All sums range over $x \in \mathbb{F}_q$ by default.

Assume p takes every value exactly once. We may assume p is monic. Then $\sum p(x)^{(q-1)/d} = \sum x^{(q-1)/d} = 0$. On the other hand, $p(x)^{(q-1)/d}$ is a monic polynomial of degree $q - 1$, so $\sum p(x)^{(q-1)/d} = \sum x^{q-1} + r(x)$, where r has degree at most $q - 2$. The sum over r vanishes and the remaining sum evaluates to -1 , so $0 = -1$, a contradiction. $\square$

A.3 The Chebotarev density theorem for function fields

Chebotarev’s density theorem for function fields describes, among other things, the asymptotic density of factorization types of $p(x) - b$ over $\mathbb{F}_q$ for large q . To illustrate it, take $p(x) = x^4 + x$ and $q = 101$. Here are some example factorizations into irreducibles:

$$\begin{aligned} x^4 + x &= x(x + 1)(x^2 - x + 1) \\ x^4 + x + 1 &= x^4 + x + 1 \\ x^4 + x + 2 &= (x + 73)(x^3 + 28x^2 + 77x + 36) \\ x^4 + x + 23 &= (x + 3)(x + 6)(x + 17)(x + 75). \end{aligned}$$

A *factorization type* is an integer partition $\lambda = (d_1, \dots, d_f)$ of the degree d . A factorization into irreducibles is of type λ if there are f factors of respective degrees d_i . For example, the above four factorizations are of types $(1, 1, 2)$, (4) , $(1, 3)$, and $(1, 1, 1, 1)$, respectively. Here are the statistics of the factorization types of $x^4 + x + b$ over a random choice of b :

factorization type	$(1, 1, 1, 1)$	$(1, 1, 2)$	$(1, 3)$	$(2, 2)$	(4)
probability	$4/101$	$26/101$	$34/101$	$12/101$	$25/101$

How do these statistics compare to those of a completely random monic degree- d polynomial? The distribution of factorization types is known to converge to the distribution of cycle lengths of a random permutation on d elements [Ros13, Chapter 2]. When $d = 4$ those statistics are

cycle type	$(1, 1, 1, 1)$	$(1, 1, 2)$	$(1, 3)$	$(2, 2)$	(4)
probability	$1/24$	$1/4$	$1/3$	$1/8$	$1/4$

Chebotarev’s density theorem explains this match. Before we state it let us look at the statistics of another example $p(x) = x^4 + x^2$:

factorization type	$(1, 1, 1, 1)$	$(1, 1, 2)$	$(1, 3)$	$(2, 2)$	(4)
probability	$13/101$	$25/101$	0	$38/101$	$25/101$

There is a reason the factorization type $(1, 3)$ never occurs. The complete factorization of $x^4 + x^2 - b$ can be calculated by first factoring $y^2 + y - b$. If this polynomial is irreducible, then $x^4 + x^2 - b$ does not have any root in $\mathbb{F}_q$. Otherwise, $x^4 + x^2 - b$ decomposes as $(x^2 - y_1)(x^2 - y_2)$ for some $y_1, y_2 \in \mathbb{F}_q$. It further splits into types $(2, 2)$, $(1, 1, 2)$, or $(1, 1, 1, 1)$ depending on the quadratic residuosity of y_1 and y_2 . The factorization type $(1, 3)$ is not among the possibilities.

In general, the possible factorization types of $p(x) - b$ are determined by the Galois group of the function field L obtained by extending the field of rational functions $\mathbb{F}_q(t)$ with the d roots of the polynomial $p(x) - t$. Here t is a transcendental (symbolic) element (and so will be the roots). The Galois group of $L/\mathbb{F}_q(t)$ is the group of automorphisms of L that leave $\mathbb{F}_q(t)$ fixed.

When $p(x) = x^4 + x$, the Galois group $\text{Gal}(L/\mathbb{F}_q(t))$ is the full symmetric group S_4 . All permutations of the roots are automorphisms. In contrast, when $p(x) = x^4 + x^2$, $\text{Gal}(L/\mathbb{F}_q(t))$ is generated by the swap of the two roots of $x^2 - y_1$ with the two roots of $x^2 - y_2$ and the swaps of the square roots $-\sqrt{y_i}$ with $\sqrt{y_i}$ for $i = 1, 2$. This is the wreath product of C_2 with itself. In cycle notation, it consists of the eight permutations

$$\epsilon, (12), (34), (12)(34), (13)(24), (1423), (1324), (14)(23)$$

The cycle type statistics are

cycle type	(1, 1, 1, 1)	(1, 1, 2)	(1, 3)	(2, 2)	(4)
probability	1/8	1/4	0	3/8	1/4

The cycle type statistics and factorization type statistics align very well again. Chebotarev's density theorem for function fields explains this phenomenon. Here is a special case of [Ros13, Theorem 9.13B]:

Theorem 22. *Assume L is a geometric extension of $\mathbb{F}_q(t)$. The number of complete factorizations of type λ of $p(x) - b$, excluding those in which $p(x) - b$ has repeated roots over the algebraic closure of $\mathbb{F}_q$, is within $O(\sqrt{q})$ of $p(\lambda) \cdot q$, where $p(\lambda)$ is the fraction of cycles of type λ in $\text{Gal}(L/\mathbb{F}_q(t))$.*

As the cycle type $(1, \dots, 1)$ has probability $1/|G|$ in every Galois group G , Theorem 22 implies that $p(x) - b$ splits completely with probability at least $1/|G| \geq 1/d!$ minus a vanishing lower order term $O(q^{-1/2})$. The factorizations with repeated roots do not affect the asymptotics as there can be at most $d - 1$ of those over the choice of b . So why isn't Theorem 22 sufficient to justify Conjecture 12?

One technical issue in the statement is that the big-oh constant in the error term may in general depend on p itself, and not only on its degree d . A closer inspection of the proof indicates that this is not the case. The constant only depends on $|G| \leq d!$ and the *genus* of $L/\mathbb{F}_q$. The genus can be bounded by a function of d only via the Riemann-Hurwitz formula.

The remaining obstacle is the assumption that the extension is geometric. An extension is geometric if every constant in L is in the base field $\mathbb{F}_q$. A constant is an element that is algebraic over $\mathbb{F}_q$.

Here is an example of a Galois extension that is not geometric. When $p(x) = x^3$ and $q = 5$, L contains the three roots $x_1(t), x_2(t), x_3(t)$ of $x^3 - t = 0$. Although x_1, x_2 , and x_3 are transcendental, x_1/x_2 is not: It is the root of $y^3 - 1 = 0$. The extension is not geometric because the element x_1/x_2 is not in $\mathbb{F}_5$. One way to see this is to apply the Frobenius automorphism $z \rightarrow z^5$ to it. This map fixes all elements of $\mathbb{F}_5$ but moves all other algebraic elements. In particular, $(x_1/x_2)^5 = x_1^5/x_2^5 = x_1^2t/x_2^2t = (x_1/x_2)^2 \neq x_1/x_2$, so $x_1/x_2 \notin \mathbb{F}_5$.

The Chebotarev density theorem should in fact not apply in this example, as x^3 is a permutation polynomial over $\mathbb{F}_q$ when $q \not\equiv 1 \pmod{3}$. The factorization type $(1, 1, 1)$ never occurs. We conjecture that no such counterexamples exist when d divides $q - 1$:

Conjecture 23. *When d divides $q - 1$ and q is prime, all extensions L of $\mathbb{F}_q(t)$ obtained by adjoining the roots of $p(x) - t$, where p is monic of degree d , are geometric.*

Conjecture 12 follows from Conjecture 23.