Fully Anonymous Perfect Threshold Secret Sharing: Explicit One-Bit Construction and Rate Amplification

Chongxu Ren and Hongbo Yu

Abstract. Fully anonymous perfect threshold secret sharing combines exactly uniform unauthorized shares with reconstruction from unlabeled share values. Con, Li, and Mazor proved that short-share schemes exist for every threshold, leaving efficient general constructions and constant rate open. We address both questions. First, we give an explicit one-bit scheme with $(2t-3)\lceil \log_2 n \rceil$-bit shares, expected-polynomial-time sharing, and deterministic polynomial-time reconstruction. The key is an exact sampler for localized affine relations, including collision patterns. Second, an affine-hull compiler turns any one-bit primitive with alphabet $\{0,1\}^b$ into a long-secret scheme with additive share overhead $b(t-1)\lceil \log_2(n+1) \rceil$, up to field-alignment padding. Authorized payloads determine an affine space; a short anonymously shared certificate identifies the secret within it. Instantiation gives $L + O(t^2 \log^2 n)$-bit shares and rate $1/2$ at $L = \Theta(t^2 \log^2 n)$. For fixed $2 \le t < n$, the rate approaches one, which the classical bound of Kishimoto et al. shows is unattainable at any finite nontrivial secret alphabet. Finally, we prove the universal bound $|\Omega| \ge (n-t+1)(n-t+2)$ for $3 \le t \le n-2$. Together with Con–Li–Mazor's specialized threshold-three construction, this establishes the exact fixed-length optimum of $2m$ share bits for $(t,n) = (3, 2^m)$, $m \ge 3$. All guarantees are information-theoretic and exact. The technical work is done by GPT 6 Astra.

Keywords: Fully anonymous secret sharing · Perfect reconstruction · Exact sampling · Rate amplification · Share-size lower bounds

1 Introduction

Threshold secret sharing allows any t of n participants to recover a secret while fewer than t learn nothing [17,3]. Standard labeled schemes supply the decoder with participant identities. In Shamir's scheme, it needs the evaluation points as well as the polynomial values. Removing this metadata raises a different question: can share values hide both the secret and identifying information, yet still permit exact reconstruction without labels?

Fully anonymous perfect secret sharing requires precisely this. For every fixed secret, any $t-1$ shares must be independent and uniform over a common alphabet, whereas any t shares must reconstruct from values alone, in any order [2,6]. The requirements meet at adjacent coalition sizes; there is no gap between privacy and reconstruction. Anonymous sharing is motivated in part by

abuse-resistant location tracking [8]. Here we study the underlying information-theoretic primitive.

A basic obstacle is that exact uniformity includes collisions. Replacing Shamir's public evaluation points by fresh random points can conceal fixed identities, but repeated points may prevent interpolation. Sampling points without replacement avoids that failure while destroying product uniformity [8,6]. Thus collisions and other degeneracies must be handled by the construction, not removed as rare events.

Con, Li, and Mazor [6] resolved existence in this exact model. For one-bit secrets they give an explicit threshold-three scheme with $2\lceil \log_2 n \rceil$-bit shares and an existential $O(t \log n)$-bit construction for every threshold. Their explicit arbitrary-access-structure construction instead uses $\binom{n}{t}^2$ bits in its threshold specialization. The remaining questions are whether short shares can be sampled efficiently for general t, whether longer secrets can achieve constant rate, and when the available share lengths are optimal.

1.1 Our results

We give an efficient short-share construction, a compiler with overhead independent of secret length, and a universal lower bound. Table 1 summarizes the construction parameters.

An exact polynomial-time realization. Theorem 1 gives, for every $2 \le t \le n$, a one-bit scheme with $b = (2t-3)\lceil \log_2 n \rceil$ share bits. Preprocessing and reconstruction are deterministic polynomial-time; sharing is expected-polynomial-time and terminates almost surely, with t part of the input. The advance over the general existence theorem of [6] is algorithmic, not an asymptotic improvement on $O(t \log n)$. Their specialized threshold-three construction remains smaller.

Table 1. Constructions with exact share anonymity and perfect reconstruction. Here $m = \lceil \log_2 n \rceil$, $\nu = \lceil \log_2(n+1) \rceil$, and compiler lengths are field-aligned: $L \in \nu\mathbb{N}$.

Result	Secret bits	Share bits	Dealer guarantee
Prior, general t [6]	1	$O(t \log n)$	Existential
Prior, $t = 3$ [6]	1	$2m$	Explicit
Theorem 1	1	$(2t-3)m$	Expected polynomial
Theorem 2	L	$L + b(t-1)\nu$	From base scheme

High rate from a short certificate. Theorem 2 uses any one-bit primitive with alphabet $\{0,1\}^b$ to share L bits with

$$\ell_L = L + H, \qquad H = b(t-1)\lceil \log_2(n+1) \rceil,$$

for field-aligned L; other lengths require fewer than $\lceil \log_2(n+1) \rceil$ padding bits. Our primitive gives $H = O(t^2 \log^2 n)$, hence rate $1/2$ at $L = \Theta(t^2 \log^2 n)$ and rate tending to one for fixed (n, t). This also yields the $O(nt \log^2 t)$ constant-rate scale, with the sharper $O(\log^2 n)$ length when t is fixed (Section 4.4). The existence conclusion already follows by instantiating the compiler with the prior existential primitive; our exact sampler additionally makes it efficient. This separates the answer to the constant-rate question of [6, Section 1.4] from the algorithmic contribution.

Lower bounds and the scope of optimality. For every scheme with at least two possible secrets, Theorem 3 proves

$$|\Omega| \geq (n-t+1)(n-t+2), \qquad 3 \leq t \leq n-2.$$

Combined with the specialized Con–Li–Mazor upper bound, this gives $\ell_{\min}(3, 2^m) = 2m$ for $m \geq 3$. The optimum concerns fixed-length *bits*, not exact alphabet size or the general-threshold construction. For long secrets, the classical bound of Kishimoto et al. [10, Theorem 3.1], $|\Omega| \geq 1+(n-t+1)(|\Sigma|-1)$, implies that for $2 \leq t < n$ rate one is a nonattained supremum. Thus the compiler has optimal leading dependence on L, but its additive overhead remains unmatched.

1.2 Technical overview

From global correction to local relation sampling. Our one-bit construction follows the polynomial-evaluation and marginal-matching strategy of Con–Li–Mazor. In their general proof, the secret-one distribution is obtained through a linear system indexed by forbidden global outputs and a cylinder-sparsity argument for nonnegativity [6, Section 4.2]. That representation does not supply a polynomial-time sampler. We instead sample an affine-relation space supported on at most $t-1$ positions. Möbius inversion prescribes its probabilities; positivity and grouping by only $O(t^2)$ support-size/dimension types make the law efficiently samplable. A moment-curve quotient realizes exactly these local relations. Affine randomization then matches the full unauthorized-tuple distribution, while a rank test separates the two secrets. Section 3 follows this order: probability law, algorithms, reconstruction, anonymity, and implementation.

Recover an affine space rather than the missing labels. For long secrets, vector-valued Shamir payloads give every authorized set the same affine space containing the secret. Its dimension is at most $t-1$, regardless of the payload length. A canonical row-echelon basis selects pivot coordinates that identify a point in this space. The dealer shares only the secret's values at these pivots using the one-bit primitive; the receiver derives the pivot indices from the payloads. Correctness therefore needs no interpolation-label search. For anonymity, the proof conditions on all payload randomness: the certificate is fixed, but its unauthorized shares remain product-uniform. This is what makes a payload-dependent certificate safe to compose with the payload.

Separate the overlap condition from its moment consequences. After conditioning on $t-3$ shares, we work with a $(3, n-t+3)$ scheme. Outputs for different secrets can share at most two multiset entries. Rather than average over both outputs immediately, we fix one and obtain an inequality for each supported output. Pairwise uniformity then forces an all-distinct output, which yields the lower bound. The five-participant boundary requires a separate finite argument.

1.3 Related work

Anonymous reconstruction and share-size bounds. Early work studied combinatorial threshold schemes and strongly ideal sharing [18,15]. Blundo and Stinson [4] developed anonymous reconstruction and its share-size bounds; Kishimoto et al. [10] strengthened the bounds, and Deng, Guo, and Liu [7] gave design-based constructions. These works concern reconstruction from values alone, which does not itself impose product-uniform unauthorized shares. We use the classical secret-alphabet bound as the long-secret converse and prove our quadratic bound using the stronger uniformity requirement. For the broader secret-sharing literature, see Beimel [1].

Share anonymity and applications. Guillermo, Martin, and O'Keefe [9] and Paskin-Cherniavsky and Olimid [14] studied additional anonymity and unpredictability notions. Bishop et al. [2] organized these guarantees into a hierarchy and studied general access structures, reconstruction error, and computational relaxations. Eldridge et al. [8] also consider mixed-dealer reconstruction for offline finding. We adopt the exact-uniformity requirement of [2,6]; the single-execution scope of our guarantees is specified in Section 2.1.

Anonymous Shamir sharing and permutation channels. Con [5] chooses Reed–Solomon evaluation points so that unlabeled values determine the polynomial, obtaining privacy threshold $k-1$ and reconstruction threshold $2k-1$. Zhang, Sun, Zhang, and Ge [19] obtain polynomial-size alphabets when $k = Rn$ and an additional εn reconstruction gap is allowed, for fixed $R, \varepsilon > 0$ with $2R + \varepsilon < 1$. These schemes already permit rate one with a threshold gap. Our compiler instead uses arbitrary distinct nonzero evaluation points and a shared certificate to retain adjacent thresholds; it does not improve their permutation-error-correcting codes.

The multiset viewpoint of Kovačević and Tan [12] also clarifies our converse. Cross-secret outputs must be separated after order is forgotten, but outputs for the *same* secret need not be separated. Moreover, a sharing scheme must assign probabilities that give uniform unauthorized marginals. Section 5.2 states the precise distance interpretation before imposing these moment constraints; a multiset-code packing bound alone does not supply them.

Subspace invariants and share compression. Kötter and Kschischang [11] use subspaces as information carriers when linear network transformations are unknown to the receiver. Our compiler uses the same change-of-basis invariant, but for a different purpose: its random affine space does not encode a unique secret

by itself; the certificate selects the secret point. This analogy explains the linear-algebra step, not its security. Krawczyk's share-compression construction [13] instead combines encryption, secret sharing, and information dispersal. Our overhead reduction uses no encryption and preserves exact information-theoretic anonymity.

Organization. Section 2 fixes the model. Sections 3, 4, and 5 develop the one-bit scheme, compiler, and lower bounds, respectively. Section 6 identifies the remaining quantitative gaps. Appendix A contains the finite sampler weights and detailed operation counts.

2 Model and notation

2.1 Syntax and security requirements

Write $[n] = \{1, \dots, n\}$ and take all logarithms to base two. For a finite secret alphabet Σ and share alphabet Ω, a threshold scheme consists of

$$\mathsf{Share} : \Sigma \to \Omega^n, \qquad \mathsf{Recon} : \Omega^t \to \Sigma.$$

The dealer may use participant labels internally. The deterministic decoder is a total function, receives only share values, and is invariant under their permutation. For one-bit schemes, $\Sigma = \{0, 1\}$.

Definition 1 (Fully anonymous perfect threshold sharing). *Let $2 \le t \le n$ and let Σ be a finite secret alphabet. The pair $(\mathsf{Share}, \mathsf{Recon})$ is a fully anonymous perfect (t, n) threshold scheme if the following two properties hold.*

(i) ***Share anonymity.*** *For every secret $s \in \Sigma$, every ordered sequence of $t-1$ distinct participants $(i_1, \dots, i_{t-1})$, and every $(x_1, \dots, x_{t-1}) \in \Omega^{t-1}$,*

$$\Pr[(W_{i_1}, \dots, W_{i_{t-1}}) = (x_1, \dots, x_{t-1}) \mid s] = |\Omega|^{-(t-1)}.$$

(ii) ***Perfect anonymous reconstruction.*** *For every ordered sequence of t distinct participants, every secret s, and every output in the support of $\mathsf{Share}(s)$,*

$$\mathsf{Recon}(W_{i_1}, \dots, W_{i_t}) = s.$$

For more than t shares, any t of them may be used.

Uniformity on $t-1$ coordinates implies uniformity on smaller sets by marginalization, and perfect secrecy because the law is independent of s.

Scope of the guarantees. Correctness concerns honest shares from distinct participants in one execution; their values need not be distinct. Total permutation invariance also specifies decoder behavior on other inputs, but does not guarantee recovery from modified shares or mixtures of executions. All privacy and correctness claims are exact, with no computational assumptions.

We use $V = |\Omega|$ for alphabet size. A fixed-length injective binary encoding requires $\ell \geq \lceil \log_2 V \rceil$ bits. Such an encoding need not use every binary string of length ℓ. This distinction matters for the six-symbol scheme in Section 5.4; the compiler instead assumes a primitive whose alphabet is exactly $\{0,1\}^b$. For a uniform secret, the information rate is $\log_2 |\Sigma| / \log_2 |\Omega|$.

Our algorithms use fair random bits and exact arithmetic. Public preprocessing is charged once for fixed parameters. A sharing algorithm with rejection loops must terminate almost surely, and its stated cost is an expectation. No loop is truncated and no correctness, anonymity, or abort error is introduced. The case $t = 1$ is handled by giving every participant the secret; the nontrivial constructions below assume $t \geq 2$.

2.2 Affine relations

For a prime power q, let $\mathbb{F}_q$ denote the field of order q. The groups $\mathrm{GL}_D(\mathbb{F}_q)$ and $\mathrm{AGL}_D(\mathbb{F}_q)$ consist of invertible linear and affine maps on $\mathbb{F}_q^D$, respectively. For an ordered tuple $x = (x_i)_{i \in A}$ in a vector space, define its affine-relation space by

$$K(x) = \left\{ \lambda \in \mathbb{F}_q^A : \sum_{i \in A} \lambda_i = 0, \quad \sum_{i \in A} \lambda_i x_i = 0 \right\}.$$

Equivalently, if C_x is the matrix whose ith column is $(1, x_i)^{\mathsf{T}}$, then $K(x) = \ker C_x$. A coordinate i occurs nontrivially in a relation exactly when the corresponding column lies in the span of the other columns.

For $A \subseteq [n]$, identify relations on A with vectors in $\mathbb{F}_q^n$ by extending them by zero, and write

$$H_A = \{\lambda \in \mathbb{F}_q^n : \mathrm{supp}(\lambda) \subseteq A,\ \sum_{i \in A} \lambda_i = 0\}.$$

For a subspace R, $\mathrm{supp}\, R$ is the union of the supports of its vectors. The Gaussian binomial coefficient $G(a, j) = \begin{bmatrix} a \\ j \end{bmatrix}_q$ counts the j-dimensional subspaces of $\mathbb{F}_q^a$; set it to zero unless $0 \leq j \leq a$, including when $a < 0$. Appendix A.1 derives the identities used below.

3 An explicit one-bit construction

The construction separates two tasks: realizing a prescribed pattern of affine relations and sampling that pattern with exactly the right law. The former gives reconstruction; together they give anonymity.

Theorem 1 (Explicit general-threshold construction). *For every $2 \le t \le n$, there is a fully anonymous perfect one-bit (t, n) threshold scheme with share alphabet $\mathbb{F}_q^D$, where*

$$q = 2^{\lceil \log_2 n \rceil}, \qquad k = t - 1, \qquad D = 2k - 1 = 2t - 3.$$

Each share has exactly $Dm = (2t-3)\lceil \log_2 n \rceil$ bits, where $m = \lceil \log_2 n \rceil$. For a supplied polynomial-basis representation of $\mathbb{F}_q$, deterministic preprocessing takes $O(k^5)$ integer operations on $O(k^2 m)$-bit operands. After preprocessing, sharing takes expected $O(nk^2 + k^3)$ field operations, plus exact integer sampling, and reconstruction takes $O(t^3)$ field operations. Using schoolbook binary arithmetic, the complete bit costs can be bounded by

$$\begin{aligned} \textit{Preprocessing: } & O(t^9 m^2), \\ \textit{Sharing: } & \mathbb{E}[T_{\mathsf{Share}}] = O((nt^2 + t^3)m^2), \\ \textit{Reconstruction: } & T_{\mathsf{Recon}} = O(t^3 m^2). \end{aligned}$$

A deterministic field construction from n adds $O(nm^3)$ bit operations once. Sharing terminates almost surely and has no correctness or anonymity error.

Use the theorem's parameters throughout this section, and write $Q = q^D$. Fix distinct public points $\alpha_1, \ldots, \alpha_n \in \mathbb{F}_q$. The exact type probabilities are precomputed once for these parameters.

3.1 The relation-space law

The secret-one dealer samples a relation space from

$$\mathcal{P}_k = \{R \le H_{[n]} : |\operatorname{supp} R| \le k\}.$$

For $R \in \mathcal{P}_k$, assign the per-subspace probability

$$\pi(R) = \sum_{\substack{S \in \mathcal{P}_k \\ R \subseteq S}} (-1)^{\dim S - \dim R} q^{\binom{\dim S - \dim R}{2}} Q^{-\dim S}. \tag{1}$$

The purpose of this law is to match independent uniform points. A d-dimensional space of affine relations is satisfied by D independent uniform coordinate rows with probability $q^{-Dd} = Q^{-d}$. Möbius inversion on the subspace lattice [16] prescribes the masses in (1) to obtain these containment probabilities. Positivity is the additional requirement that turns the inversion formula into a sampling distribution.

Lemma 1 (Möbius law). *The values in* (1) *form a probability distribution on* $\mathcal{P}_k$ *and for every subspace* $T \le H_{[n]}$ *with* $|\operatorname{supp} T| \le k$,

$$\Pr[T \subseteq R] = Q^{-\dim T}. \tag{2}$$

Proof (of Lemma 1). By Appendix A.1, summing (1) over $R \supseteq T$ and interchanging the finite sums cancels every term indexed by $S \neq T$. Hence

$$\sum_{R \supseteq T} \pi(R) = Q^{-\dim T}.$$

Taking $T = 0$ gives normalization. It remains to prove positivity.

For R of dimension r, let $N_j(R)$ be the number of superspaces in $\mathcal{P}_k$ of dimension $r + j$. For $j \ge 1$, every such superspace lies in H_C for some k-set C, so

$$N_j(R) \le \binom{n}{k} G(k-1-r, j) \le 2n^k q^{j(k-1-r-j)}.$$

Consequently, using $D = 2k - 1$,

$$\left| \frac{\pi(R)}{Q^{-r}} - 1 \right| \le 2n^k \sum_{j \ge 1} q^{-j(k+r)-j(j+1)/2} \le \frac{2n^k}{q^{k+1}-1} < 1$$

for $k \ge 2$. To justify the geometric-series bound, for $j \ge 1$ we have $j(k+r)+j(j+1)/2 \ge j(k+1)$. The final strict inequality follows from $2n^k \le 2q^k < q^{k+1} - 1$ since $q \ge 4$. Hence every mass is positive. For $k = 1$, $\mathcal{P}_k = \{0\}$ and the conclusion is immediate.

The sampler groups spaces by support size and dimension rather than enumerating $\mathcal{P}_k$; see Section 3.5.

3.2 Sharing and reconstruction

Our construction shares the polynomial-evaluation and exact marginal-matching strategy of Con–Li–Mazor [6, Sections 3–4].

Secret zero: $\mathsf{Share}(0)$. Sample D independent uniform polynomials $f_1, \dots, f_D \in \mathbb{F}_q[X]_{\le k-1}$. Give participant i the vector

$$W_i = (f_1(\alpha_i), \dots, f_D(\alpha_i)) \in \mathbb{F}_q^D.$$

Secret one: $\mathsf{Share}(1)$. Sample R according to (1), put $B = \operatorname{supp} R$, $b = |B|$, and set $h = k + b$. Define moment-curve vectors

$$u_i = (\alpha_i, \alpha_i^2, \dots, \alpha_i^h) \in \mathbb{F}_q^h.$$

Let

$$\mathcal{L} = \left\{ \sum_{i \in B} \rho_i u_i : \rho \in R \right\}.$$

Vandermonde independence proves that the map from R to $\mathcal{L}$ is injective (Lemma 2). Compute any surjective linear map $P : \mathbb{F}_q^h \to \mathbb{F}_q^{h-\dim R}$ with kernel $\mathcal{L}$, and pad Pu_i with zero coordinates to obtain $v_i \in \mathbb{F}_q^D$. The inequality

$$h - \dim R \le 2k - 1 = D$$

holds because $b \le k$ and a nonzero R has dimension at least one; for $R = 0$, $h = k \le D$.

A concrete choice for P is the matrix whose rows are a basis of $\mathcal{L}^\perp$ under the standard nondegenerate dot product; Gaussian elimination computes it. Finally sample $M \leftarrow \mathrm{GL}_D(\mathbb{F}_q)$ and $z \leftarrow \mathbb{F}_q^D$ independently and uniformly, and give participant i the share

$$W_i = Mv_i + z. \tag{3}$$

To sample M, draw all D^2 entries uniformly and repeat until its rank is D. All field draws use exactly m fair bits. Only the W_i values are distributed; B, R, P, M, z are private dealer randomness or private intermediate state.

Value-only reconstruction: Recon. Given t share vectors $x_1, \ldots, x_t \in \mathbb{F}_q^D$, form

$$C = \begin{pmatrix} 1 & \cdots & 1 \\ x_1 & \cdots & x_t \end{pmatrix}.$$

Output zero if every column of C lies in the span of the other columns, and output one otherwise. This is a rank or nullspace computation and is manifestly invariant under permuting the share values.

3.3 Perfect anonymous reconstruction

The quotient must create exactly the sampled relations, not additional relations on authorized tuples. The following identity supplies that property.

Lemma 2 (Localized relations). *For every sampled R, every $A \subseteq [n]$ with $|A| \le k + 1$, and the vectors v_i of Section 3.2,*

$$K((v_i)_{i\in A}) = R \cap H_A. \tag{4}$$

Proof. The augmented vectors $(1, u_i)$ are columns of a Vandermonde matrix. For any $a \le h+1$ distinct indices, its first a rows have determinant $\prod_{i<j}(\alpha_j - \alpha_i) \neq 0$, so those columns are independent. In particular, if a vector of $R \le H_B$ maps to zero in $\mathcal{L}$, it gives a relation among the at most b columns indexed by B, and must be zero. This proves $\dim \mathcal{L} = \dim R$ and the existence of the prescribed quotient. Suppose $\lambda \in K((v_i)_{i\in A})$. Since P has kernel $\mathcal{L}$, there is $\rho \in R$ with

$$\sum_{i\in A} \lambda_i u_i = \sum_{i\in B} \rho_i u_i.$$

Both λ and ρ have coordinate sum zero. Therefore $\lambda - \rho$ is an affine relation among the augmented moment-curve columns indexed by $A \cup B$. This set has size at most $k + 1 + b = h + 1$, so independence gives $\lambda = \rho$. Thus $\lambda \in R \cap H_A$. The reverse inclusion follows directly from $\ker P = \mathcal{L}$.

For secret zero, on any $(k+1)$-tuple of distinct evaluation points, the Lagrange relation

$$\lambda_j = \left(\prod_{\ell \neq j} (\alpha_j - \alpha_\ell) \right)^{-1}$$

is nonzero in every coordinate and annihilates every degree-$k-1$ polynomial, including the constant polynomial. To see the identity directly, write any polynomial of degree at most $k-1$ in its Lagrange interpolation form on $k+1$ points and compare the coefficient of X^k; that coefficient on the left is zero and on the right is $\sum_j \lambda_j f(\alpha_j)$. Here the α_j denote the selected evaluation points. Hence every augmented column lies in the span of the others and reconstruction outputs zero.

For secret one, let A be any $k+1$ participants. Since $|\operatorname{supp} R| \leq k$, choose $i \in A \setminus \operatorname{supp} R$. By (4), every relation in $K((v_j)_{j \in A})$ has zero ith coordinate. Affine transformations preserve relations, so the same holds for the shares. The ith column is not in the span of the others, and reconstruction outputs one on every support point.

3.4 Share anonymity

Matching relation spaces is not by itself enough to match share distributions. We also need uniformity among tuples with the same relations.

Lemma 3 (Affine orbits). *Two ordered tuples in $(\mathbb{F}_q^D)^a$ have the same affine-relation space if and only if one is obtained from the other by an element of $\mathrm{AGL}_D(\mathbb{F}_q)$. Consequently, uniform affine randomization is uniform on the set of tuples with a fixed affine-relation space.*

Proof. If x and y have the same relation space, define on the affine span of x,

$$\phi\left(\sum_i c_i x_i \right) = \sum_i c_i y_i, \qquad \sum_i c_i = 1.$$

Two affine representations differ by an element of the common relation space, so ϕ is well-defined. Its linear part is an isomorphism between the translation spaces, and extending bases gives an element of $\mathrm{AGL}_D(\mathbb{F}_q)$ mapping every x_i to y_i. The converse is immediate. A uniform group element pushes forward to a uniform orbit because all points in an orbit have equal-size stabilizers.

Secret zero. For any k distinct evaluation points, the $k \times k$ evaluation matrix for a degree-at-most-$k-1$ polynomial is invertible. The k evaluations are therefore independent uniform field elements. The D coordinate polynomials are independent, so every unauthorized tuple is product-uniform in $\mathbb{F}_q^D$.

Secret one. We prove exact uniformity of every unauthorized tuple. Fix A with $|A| \le k$. If X_i are independent uniform points in $\mathbb{F}_q^D$, then for every $T \le H_A$ of dimension r,

$$\Pr[T \subseteq K((X_i)_{i\in A})] = q^{-Dr} = Q^{-r}. \tag{5}$$

Indeed, each of the D coordinate rows must satisfy r independent linear equations. Lemma 1 gives the same containment moments for $R \cap H_A$. Since the containment transform on the finite subspace lattice is invertible, the two relation spaces have the same distribution.

By Lemma 2, the candidate tuple has relation space $R \cap H_A$. Conditioned on this space, uniform M, z make it uniform on its affine orbit by Lemma 3. The independent uniform tuple X_A is also uniform on the same orbit when conditioned on its relation space. Their conditional orbit laws and their relation-space laws agree, so the share tuples are identically distributed. This includes repeated points and every lower-rank affine configuration.

3.5 Exact implementation and running time

There are only $O(k^2)$ nonzero types $(b, r) = (|\operatorname{supp} R|, \dim R)$, and every space of a type has the same mass. The sampler draws a type from precomputed integer weights, chooses a uniform support, and samples a full-support basis by rejection. Each subspace has equally many ordered bases, so the accepted space is uniform within its type. The basis loop accepts with probability at least $1/4$. Appendix A.1 derives the weights with common denominator Q^{k-1}; Appendix A.2 gives the full algorithm and acceptance proof.

Example: the sampler at $k = 2$. Here $D = 3$ and $\mathcal{P}_2$ consists of the zero space and one line $H_{\{i,j\}} = \operatorname{span}(e_i - e_j)$ for each unordered pair of participants. Equation (1) gives each line mass q^{-3} and the zero space mass $1 - \binom{n}{2}q^{-3}$. Thus one either selects no relation or selects a uniform pair to coincide in the secret-one configuration. The latter branch is required to reproduce the collision probability of two independent uniform shares; suppressing it would destroy exact anonymity. This illustrates the general sampler's type, support, and basis stages, not an improvement over the specialized threshold-three construction.

Polynomial evaluation, the quotient map, and affine randomization take expected $O(nk^2 + k^3)$ field operations after preprocessing. The decoder computes one augmented nullspace in $O(t^3)$ operations and inspects its coordinate support. Appendix A.3 converts these counts to the bit bounds in Theorem 1, including exact integer sampling and field construction.

Proof (of Theorem 1). Lemma 1 supplies a probability law with the required containment moments. Sections 3.3 and 3.4 prove perfect anonymous reconstruction and exact share anonymity, respectively. Each share lies in $\mathbb{F}_q^D$ and therefore has exactly Dm bits. Lemma 8 and Appendix A.3 establish the exact sampler, almost-sure termination, and the stated operation counts.

4 Rate amplification by an affine-hull certificate

The compiler uses Shamir evaluations [17] for the payload and the one-bit scheme only for a certificate. The receiver first recovers an affine space containing the secret, then uses the certificate to select one point in that space. This avoids recovering the missing evaluation labels.

Definition 2 (One-bit auxiliary primitive). *Fix (t, n). A one-bit auxiliary primitive is a scheme satisfying Definition 1 with one-bit secrets and share alphabet $\Gamma = \{0,1\}^b$. Its algorithms are denoted by Share_1 and Recon_1. Independent invocations use independent randomness. Their messages may be functions of the secret and the dealer's payload randomness.*

Throughout this section, put

$$k = t - 1, \qquad \nu = \lceil \log_2(n+1) \rceil, \qquad K = \mathbb{F}_{2^\nu}. \tag{6}$$

Fix a polynomial-basis representation of K and distinct nonzero public points $\alpha_1, \ldots, \alpha_n \in K$. The field is independent of the secret length; Section 4.5 accounts for constructing it.

Theorem 2 (Affine-hull rate compiler). *Given a one-bit auxiliary primitive as in Definition 2, for every integer $d \geq 1$ the algorithms below form a fully anonymous perfect (t, n) threshold scheme with*

$$\Sigma = K^d, \qquad \Omega = K^d \times \Gamma^{k\nu}.$$

Writing $L = d\nu$, the secret length is L bits and the share length is

$$\ell_L = L + H, \qquad H = b(t-1)\lceil \log_2(n+1) \rceil. \tag{7}$$

The decoder is a total permutation-invariant function on Ω^t. Sharing and reconstruction each use $k\nu$ calls to the auxiliary primitive, plus $O(ndk + dk^2)$ and $O(dk^2)$ field operations, respectively.

The fixed k-entry certificate also covers $d < k$ and rank-deficient masks. No rank conditioning is used.

4.1 Vector payloads and the certificate

Sharing. For $s \in K^d$, sample independent uniform rows $R_1, \ldots, R_k \in K^d$, let $R \in K^{k \times d}$ be their matrix, and set

$$Y_i = s + \sum_{j=1}^{k} \alpha_i^j R_j \qquad (i \in [n]). \tag{8}$$

Compute the unique reduced row-echelon basis U of rowspan R, retaining only its $r = \operatorname{rank} R$ nonzero rows. Let $j_1 < \cdots < j_r$ be its pivot columns; thus $U_{a,j_c} = \mathbf{1}_{a=c}$. Define

$$C = (s_{j_1}, \ldots, s_{j_r}, 0, \ldots, 0) \in K^k. \tag{9}$$

Using the fixed field representation, encode C as $k\nu$ bits. For each bit C_ℓ, independently sample $(Z_{1,\ell}, \ldots, Z_{n,\ell}) \leftarrow \mathsf{Share}_1(C_\ell)$ and output

$$W_i = (Y_i, Z_{i,1}, \ldots, Z_{i,k\nu}).$$

Only the pivot-coordinate values are shared; the receiver derives the pivot indices and rank from the payloads.

Reconstruction. Given $t = k + 1$ complete shares, first apply Recon_1 componentwise to recover $C \in K^k$. Let the payloads, in their presented order, be $y_0, y_1, \ldots, y_k$. Compute the reduced row-echelon basis

$$U' = \mathrm{RREF}\left(\mathrm{rowspan}\begin{pmatrix} y_1 - y_0 \\ \vdots \\ y_k - y_0 \end{pmatrix}\right),$$

with nonzero rows $U'_1, \ldots, U'_{r'}$ and pivot columns $j'_1 < \cdots < j'_{r'}$. Output

$$\widehat{s} = y_0 - \sum_{a=1}^{r'} \big((y_0)_{j'_a} - C_a\big) U'_a. \tag{10}$$

Here $r' \le k$, the sum is empty at rank zero, and every certificate bit string encodes an element of K^k.

Reading the certificate. For example, consider $k = 2$ and a rank-one execution in K^3 with canonical basis $U = (0, 1, \gamma)$. Its pivot is the second coordinate, so $C = (s_2, 0)$. Any payload is $y = s + \lambda U$, and the decoder returns $y - (y_2 - C_1)U = s$. Every authorized set recovers U from its differences; the certificate need not say which coordinate is the pivot.

4.2 Correctness and permutation invariance

The receiver recovers the dealer's translation space even when the masks are dependent. As in subspace coding [11], an unknown invertible change of basis does not obscure the row space. Here this invariant locates an affine space rather than identifying the secret by itself.

Lemma 4 (Authorized payloads determine the same affine hull). *For every secret s, every mask matrix R, and every t distinct participants $i_0, \ldots, i_k$,*

$$\mathrm{aff}\{Y_{i_0}, \ldots, Y_{i_k}\} = s + \mathrm{rowspan}\, R. \tag{11}$$

Proof. The payload-difference matrix is DR, where

$$D_{a,j} = \alpha_{i_a}^j - \alpha_{i_0}^j \qquad (1 \le a, j \le k).$$

The Vandermonde matrix with rows $(1, \alpha_{i_a}, \ldots, \alpha_{i_a}^k)$, $0 \le a \le k$, is invertible. Subtracting its first row from the others and expanding along the first column shows that D is invertible. Hence $\mathrm{rowspan}(DR) = \mathrm{rowspan}\, R$, with no restriction on $\mathrm{rank}\, R$. Since $Y_{i_0} - s \in \mathrm{rowspan}\, R$, the corresponding affine spans also agree.

On honest inputs, Lemma 4 gives $U' = U$ and identical pivot columns for dealer and receiver. Write $y_0 - s = \lambda U$. The pivot identity gives $\lambda_a = (y_0)_{j_a} - s_{j_a}$, so (9) and (10) recover s. The argument includes every rank and every multiplicity pattern.

Behavior on arbitrary inputs. Equation (10) is the unique point of $\mathrm{aff}\{y_0, \ldots, y_k\}$ whose pivot coordinates are $C_1, \ldots, C_{r'}$: the pivot submatrix of U' is the identity, so these coordinates give a bijection from that affine space to $K^{r'}$. The affine space, its canonical basis, and the auxiliary decoder's output C are all permutation-invariant. Thus the decoder is total and symmetric even for inconsistent payloads and headers, including rank zero.

4.3 Exact share anonymity

Proof (of Theorem 2). Perfect reconstruction and total permutation invariance were proved above. For anonymity, fix a secret s and an ordered sequence I of $a \le k$ distinct participants. The map from the mask coefficients in one payload coordinate to their evaluations has rows $(\alpha_i, \alpha_i^2, \ldots, \alpha_i^k)$ for $i \in I$. Its first a columns form an invertible scaled Vandermonde matrix. Each coordinate therefore has uniform evaluation tuple in K^a, and different coordinates are independent. Consequently,

$$\Pr[Y_I = y_I \mid s] = |K|^{-da}. \tag{12}$$

The certificate depends on both s and R, so independence of headers and payloads must be proved conditionally. Fix the entire matrix R as well as s. The certificate is now fixed, and fresh auxiliary randomness gives, for every header tuple $z_I \in \Gamma^{k\nu a}$,

$$\Pr[Z_I = z_I \mid s, R] = 2^{-bk\nu a}. \tag{13}$$

This probability is independent of R and of the certificate's value. Summing over the masks consistent with a prescribed payload tuple gives

$$\begin{aligned}\Pr[Y_I = y_I, Z_I = z_I \mid s] &= \sum_R \Pr[R]\, \mathbf{1}\{Y_I(s,R) = y_I\} \Pr[Z_I = z_I \mid s, R] \\ &= |K|^{-da} 2^{-bk\nu a} = |\Omega|^{-a}.\end{aligned}$$

This is exact product-uniformity, including all rank-deficient outcomes. The alphabet sizes give (7). Horner evaluation uses $O(ndk)$ field operations; row reduction uses $O(dk^2)$; the remaining certificate and affine-projection work fits within these bounds. The $k\nu$ independent auxiliary invocations give the stated call counts.

4.4 Rates and the dependence on the threshold

Corollary 1 (Constant overhead and high rate). *For a primitive as in Definition 2, the compiler has rate*

$$\rho_L = \frac{L}{L+H}, \qquad H = bk\nu,$$

for $L = d\nu$. In particular, $L = H$ gives rate $1/2$, and rate at least $1-\varepsilon$ is obtained when

$$d \geq \max\left\{1, \left\lceil \frac{1-\varepsilon}{\varepsilon} bk \right\rceil \right\}, \qquad 0 < \varepsilon < 1.$$

With Theorem 1, the exact overhead is

$$H = (2t-3)(t-1)\lceil \log_2 n \rceil \lceil \log_2 (n+1) \rceil = O(t^2 \log^2 n). \tag{14}$$

Thus rate $1/2$ is attained at $L = \Theta(t^2 \log^2 n)$ with expected polynomial-time sharing and deterministic polynomial-time reconstruction. For each fixed (n, t), the rate tends to one as L grows.

Proof. The rate follows from Theorem 2. Choose $d = bk$ for rate $1/2$, and rearrange $d/(d+bk) \geq 1-\varepsilon$ for the general bound. Substitution of $b = (2t-3)\lceil \log_2 n \rceil$ proves (14). Here H is independent of L, so $\rho_L \to 1$ for fixed parameters.

Comparison of parameter scales. The threshold-dependent bound also implies $O(nt \log^2 t)$ secret length for constant rate, uniformly over $2 \leq t \leq n$; the more informative bound is $O(t^2 \log^2 n)$. Indeed, with $x = n/t \geq 1$ and $\log t \geq 1$,

$$\frac{t^2 \log^2 n}{nt \log^2 t} = \frac{1}{x}\left(1 + \frac{\log x}{\log t}\right)^2 \leq \frac{(1+\log x)^2}{x} = O(1).$$

For fixed t and growing n, the exact formula is only $O(\log^2 n)$, rather than linear in n. When $t = \Theta(n)$, the two displayed scales agree up to constant factors.

Arbitrary bit lengths. For an L-bit secret with $\nu \nmid L$, append zeros to length $L' = \nu\lceil L/\nu \rceil$, run the compiler with $d = L'/\nu$, and have the decoder return the first L bits. The share alphabet is the full binary cube of dimension $L' + H$, and the security proof holds for each fixed padded secret. Thus every $L \geq 1$ admits

$$\ell_L = L + H + \delta_L, \qquad 0 \leq \delta_L < \nu.$$

No restriction of the secret set is mistaken for a restriction of the uniform share alphabet.

Concrete overhead and choice of primitive. At $t = 3$, the specialized Con–Li–Mazor primitive (Proposition 1) has $b = 2\lceil \log_2 n \rceil$, giving $H = 4\lceil \log_2 n \rceil \lceil \log_2 (n+1) \rceil$. Table 2 uses this smaller primitive where available. The entries are exact storage calculations, not implementation benchmarks.

An existential primitive [6, Theorem 1.4] gives the same asymptotic overhead without the algorithmic guarantee. At $t = n$, additive sharing already has rate one and no compiler is needed. The converse for $t < n$ appears in Section 5.5.

Table 2. Fixed per-share overhead H of the affine-hull compiler. A secret of H bits has rate 1/2; larger aligned secrets have rate $L/(L+H)$. Only the threshold-three row uses the specialized prior primitive.

n	t	Base share bits b	Certificate bits $k\nu$	Overhead bits H
30	20	185	95	17,575
1024	16	290	165	47,850
1024	3	20	22	440

4.5 Uniform setup and implementation cost

Let T_{share} and T_{recon} be bit costs for one auxiliary invocation, with sharing cost an expectation when appropriate. Given a polynomial-basis representation of K, schoolbook field arithmetic gives

$$T'_{\text{share}} = k\nu T_{\text{share}} + O((ndk + dk^2)\nu^2),$$
$$T'_{\text{recon}} = k\nu T_{\text{recon}} + O(dk^2\nu^2).$$

These bounds include generating the uniform mask and reading and writing the payloads. In particular, for fixed (n, t) both costs are linear in the secret length, apart from the fixed auxiliary cost.

For a uniform construction from parameters alone, set $K = \mathbb{F}_2[X]/(g)$, where g is the first monic irreducible degree-ν polynomial in a fixed order. The exact test in Appendix A.3 takes $O(\nu^3)$ bit operations per candidate. There are at most $2^\nu < 2(n+1)$ candidates, giving deterministic cost $O(n\nu^3)$. The first n nonzero coefficient strings give the evaluation points. Since $\nu = O(\log n)$, field setup is independent of L.

Using Theorem 1 and writing $\mu = \lceil \log_2 n \rceil$, the online bounds become

$$T'_{\text{share}} = O\big(k\nu(nt^2 + t^3)\mu^2 + (ndk + dk^2)\nu^2\big),$$
$$T'_{\text{recon}} = O\big(k\nu t^3\mu^2 + dk^2\nu^2\big).$$

The reusable type table costs $O(t^9\mu^2)$, and constructing both required fields adds $O(n\nu^3)$; if their degrees agree, one representation suffices. Setup is reused across all certificate bits and executions. Almost-sure termination follows from the corresponding guarantee of the primitive; sharing costs remain expectations.

5 Universal lower bounds and threshold-three optimality

The first bound uses only two secrets and yields threshold-three bit optimality. The classical bound in Section 5.5 retains the secret-alphabet size and supplies the compiler's converse. Both apply to arbitrary schemes in Definition 1, without linearity or computational assumptions; write $V = |\Omega|$.

Theorem 3 (Universal alphabet lower bound). *Every fully anonymous perfect (t, n) threshold scheme with at least two possible secrets satisfies*

$$V \geq (n - t + 1)(n - t + 2) \qquad (3 \leq t \leq n - 2). \tag{15}$$

For $t = n - 1 \geq 3$, it satisfies $V \geq 5$.

The proof reduces to threshold three and then uses an inequality that holds for each supported output. Before developing it, we state the matching consequence and identify which upper construction supplies it.

5.1 Where the upper and lower bounds match

The relevant upper bound is the specialized construction of Con, Li, and Mazor [6, Theorem 1.3], not a dimension reduction of Theorem 1.

Proposition 1 (Con–Li–Mazor threshold-three upper bound). *For every $n \geq 4$, there is an explicit fully anonymous perfect one-bit $(3, n)$ scheme with share alphabet $\mathbb{F}_q^2$, where $q = 2^{\lceil \log_2 n \rceil}$. Each share has exactly $2\lceil \log_2 n \rceil$ bits. Reconstruction uses a constant number of field operations; exact sharing has expected polynomial cost.*

Proof. For q participants, the construction of [6, Section 3] uses random affine lines for secret zero and random affine images of the parabola (α, α^2), with a random-matching branch, for secret one. Propositions 3.4 and 3.5 there prove perfect reconstruction and exact uniformity of every ordered pair. The dealer's field draws, uniform matching, invertible-matrix draw, and coin of bias $(q-1)/q^2$ have exact samplers with expected polynomial cost. Restricting to any fixed n participants preserves both guarantees. The decoder tests only equality and collinearity, using a constant number of field operations.

Corollary 2 (Exact fixed-length bit optimality). *Let $V_{\min}(3, n)$ be the minimum share-alphabet size and let $\ell_{\min}(3, n)$ be the minimum fixed-length binary encoding length, over all fully anonymous perfect one-bit $(3, n)$ schemes. For $n \geq 5$,*

$$(n - 1)(n - 2) \leq V_{\min}(3, n) \leq 2^{2\lceil \log_2 n \rceil}, \tag{16}$$

$$\lceil \log_2((n - 1)(n - 2)) \rceil \leq \ell_{\min}(3, n) \leq 2\lceil \log_2 n \rceil. \tag{17}$$

In particular, for every integer $m \geq 3$,

$$\ell_{\min}(3, 2^m) = 2m. \tag{18}$$

More generally, the upper and lower bit bounds coincide whenever

$$(n-1)(n-2) > 2^{2\lceil \log_2 n \rceil - 1}.$$

For arbitrary growing n, $V_{\min}(3,n) = \Theta(n^2)$ and $\ell_{\min}(3,n) = 2\log_2 n + O(1)$.

Proof. Combine Theorem 3 with Proposition 1 and round binary encoding lengths up to integers. For $n = 2^m \geq 8$,

$$2^{2m-1} < (2^m - 1)(2^m - 2) \leq 2^{2m}.$$

The lower bound therefore forces at least $2m$ bits and the upper construction uses exactly $2m$. The same rounding argument proves the more general criterion. Finally, $2^{\lceil \log_2 n \rceil} < 2n$ gives the asymptotic bounds.

At powers of two, the alphabet interval remains $[n^2 - 3n + 2, n^2]$. Its ratio tends to one, but exact alphabet optimality does not follow. The bit comparison is illustrated in Table 3.

Table 3. Threshold-three bounds. The upper bound is Proposition 1; matching bit columns certify fixed-length bit optimality.

n	Alphabet lower bound	Bit lower bound	Bit upper bound
5	12	4	6
6	20	5	6
8	42	6	6
13	132	8	8
16	210	8	8
1024	1,045,506	20	20

Table 3 also shows why the strengthened bound can matter beyond powers of two. At $n = 13$, $(n-1)(n-2) = 132$ forces eight bits, whereas the intermediate bound $(n-2)^2 = 121$ would force only seven.

5.2 Threshold reduction and cross-secret overlap

We begin the proof of Theorem 3 with a conditioning step that preserves the alphabet and the exact model.

Lemma 5 (Threshold reduction). *A fully anonymous perfect (t,n) scheme with $t \geq 3$ induces a fully anonymous perfect $(3,N)$ scheme over the same alphabet, where $N = n - t + 3$.*

Proof. Fix an ordered set J of $t-3$ participants and a vector $c \in \Omega^{t-3}$. For each secret s, the event $W_J = c$ has probability $V^{-(t-3)} > 0$. Condition the sharing distribution on this event and retain the remaining N shares. For any

two remaining distinct participants i, j and $x, y \in \Omega$, exact uniformity on $t-1$ coordinates gives

$$\Pr[W_i = x, W_j = y \mid W_J = c, s] = \frac{V^{-(t-1)}}{V^{-(t-3)}} = V^{-2}.$$

The new decoder prepends c to the three received values and invokes the original decoder. Every such input comes from t distinct original participants, so reconstruction is perfect and needs no identities. The case $t = 3$ uses the empty vector c.

Consider two supported outputs of the reduced scheme for different secrets. Let a_x, b_x be the multiplicities of $x \in \Omega$ in these outputs, so $\sum_x a_x = \sum_x b_x = N$. Their multiset overlap satisfies

$$\sum_x \min(a_x, b_x) \le 2. \tag{19}$$

Otherwise the two outputs would contain a common three-element submultiset. Ordering that multiset identically in the two outputs would give the deterministic anonymous decoder the same input for two different secrets, a contradiction. Repeated values still correspond to distinct participant positions.

Equivalently, in the multiset metric of Kovačević and Tan [12],

$$d_{\mathrm{ms}}(a, b) = \frac{1}{2} \sum_x |a_x - b_x| = N - \sum_x \min(a_x, b_x) \ge N - 2.$$

Only outputs for different secrets need this separation; the decoder need not distinguish two random sharings of one secret. Distance therefore captures reconstruction, while the exact-uniformity moments below impose the additional privacy constraint.

Lemma 6 (Refined overlap inequality). *For nonnegative integer multiplicities satisfying* (19)*, define*

$$Z = \sum_x a_x b_x, \qquad B = \sum_x a_x b_x (a_x + b_x - 1), \qquad Y = \sum_x a_x (a_x - 1) b_x (b_x - 1).$$

Then

$$Z^2 \le 2B - \frac{2Y}{N-1}. \tag{20}$$

If all multiplicities in both outputs are at most two, then also

$$Z^2 \le 2B - 2Y. \tag{21}$$

Proof. If $Y = 0$, every common symbol has minimum multiplicity one and there are at most two common symbols. For each such symbol, $a_x + b_x - 1 = a_x b_x$, so $B = \sum_x (a_x b_x)^2$. Cauchy–Schwarz gives $Z^2 \le 2B$, proving both assertions in this case.

If $Y > 0$, the overlap bound permits only one common symbol, whose multiplicities are $2, h$ in some order, with $2 \le h \le N$. Thus

$$Z = 2h, \qquad B = 2h(h+1), \qquad Y = 2h(h-1).$$

Consequently,

$$2B - \frac{2Y}{N-1} - Z^2 = \frac{4h(N-h)}{N-1} \ge 0.$$

If both outputs have multiplicities at most two, then $h = 2$ and $Z^2 = 2B - 2Y$.

5.3 A constraint on each output and the lower-bound proof

The key is to average the overlap inequality over one sharing distribution while retaining the other output as a fixed object.

Lemma 7 (A constraint on every supported output). *In a fully anonymous perfect $(3, N)$ scheme, fix any output for one secret and write*

$$A = \sum_x a_x(a_x - 1).$$

Then

$$V(N + A) \ge N(N-1)(N-2) + 2A. \tag{22}$$

In particular, if one supported output has all its shares distinct, then $V \ge (N-1)(N-2)$.

Proof. Sample an output for a different secret, keeping $(a_x)_x$ fixed. Single-share and pairwise uniformity of this second output give

$$\mathbb{E}[b_x] = \frac{N}{V}, \qquad \mathbb{E}[b_x(b_x-1)] = \frac{N(N-1)}{V^2}, \qquad \sum_x a_x^2 = N + A.$$

Writing $Z = \sum_{j=1}^{N} a_{W'_j}$, pairwise uniformity also evaluates the cross-terms of Z^2. Thus, with expectation only over the second output,

$$\mathbb{E}_b[Z^2] = \frac{N(N+A)}{V} + \frac{N^3(N-1)}{V^2}, \tag{23}$$

$$\mathbb{E}_b[B] = \frac{N(N+A)}{V} + \frac{N^2(N-1)}{V^2}, \tag{24}$$

$$\mathbb{E}_b[Y] = \frac{N(N-1)A}{V^2}. \tag{25}$$

For the second identity, expand $B = \sum_x a_x^2 b_x + \sum_x a_x b_x(b_x - 1)$. Taking expectations in (20), substituting these expressions, and multiplying by V^2/N proves (22). An all-distinct output has $A = 0$, giving the last assertion.

Proof (of Theorem 3). Use Lemma 5 and set $N = n - t + 3$. For a random output of either secret, pairwise uniformity gives

$$\mathbb{E}[A] = \frac{N(N-1)}{V}. \tag{26}$$

Averaging (22) over the first output and simplifying yields

$$V^2 - (N-1)(N-3)V - 2(N-1) \geq 0. \tag{27}$$

For $N \geq 5$, this excludes every positive integer $V \leq (N-1)(N-3)$, and hence implies $V \geq (N-2)^2$.

Suppose first that $N \geq 6$. Then $2(N-2)^2 > N(N-1)$, so (26) gives $\mathbb{E}[A] < 2$. Each $a_x(a_x - 1)$ is even, and A is nonnegative; therefore some supported output has $A = 0$. Lemma 7 now gives $V \geq (N-1)(N-2)$.

It remains to treat $N = 5$. The quadratic bound gives $V \geq 9$, so it suffices to exclude $V = 9, 10, 11$. If $V = 9$, inequality (22) gives $7A \geq 15$. Thus $A \geq 4$ for every output, contradicting $\mathbb{E}[A] = 20/9$. If $V = 10$, the same inequality excludes $A = 0$, while $\mathbb{E}[A] = 2$. Consequently every output for either secret has $A = 2$, that is, exactly one repeated pair and all other symbols distinct. All multiplicities are at most two, so (21) applies. Using (23)–(25) with $N = 5$ and $A = 2$ gives

$$\frac{35}{V} + \frac{500}{V^2} \leq 2\left(\frac{35}{V} + \frac{100}{V^2}\right) - \frac{80}{V^2},$$

which simplifies to $7V \geq 76$ and fails at $V = 10$. If $V = 11$, then $\mathbb{E}[A] = 20/11 < 2$, so an all-distinct output exists and Lemma 7 forces $V \geq 12$. This proves $V \geq (N-1)(N-2)$ also for $N = 5$. Substituting $N = n - t + 3$ proves (15).

Finally, $t = n - 1$ reduces to $N = 4$, where (27) reads $V^2 - 3V - 6 \geq 0$. Its smallest admissible positive integer is five.

Fixing one output improves the intermediate square bound by $n - t + 1$ alphabet symbols, using only pairwise uniformity after conditioning.

5.4 The four-participant boundary

The condition $m \geq 3$ in Corollary 2 cannot be dropped under the arbitrary-alphabet model.

Proposition 2 (A six-symbol $(3, 4)$ scheme). *There is a fully anonymous perfect one-bit $(3, 4)$ scheme over $\mathbb{Z}_6$. Consequently $5 \leq V_{\min}(3, 4) \leq 6$ and $\ell_{\min}(3, 4) = 3$.*

Proof. Let $T_0 = \{0, 0, 2, 5\}$ and $T_1 = \{0, 0, 1, 4\}$ in $\mathbb{Z}_6$. To share s, translate T_s by a uniform $u \in \mathbb{Z}_6$ and uniformly permute its four entry positions, counting repeated entries with multiplicity. In each base multiset, the twelve ordered pairs of distinct positions have each difference in $\mathbb{Z}_6$ exactly twice. The independent translation makes the first value uniform conditional on the chosen positions, so

every ordered pair of participants has probability $(2/12)(1/6) = 1/36$ of each prescribed pair of share values.

A repeated-value triple, after subtracting its repeated value, has singleton offset in $\{2, 5\}$ for secret zero and in $\{1, 4\}$ for secret one. Distinct-value triples are translates of $\{0, 2, 5\}$ or $\{0, 1, 4\}$, respectively. Their successive gaps around $\mathbb{Z}_6$ are cyclic rotations of $(2, 3, 1)$ or $(1, 3, 2)$, two disjoint classes. Thus no unordered triple can occur for both secrets. Testing membership in these two finite families reconstructs perfectly, with arbitrary output on unsupported inputs. The lower bound $V \geq 5$ from Theorem 3 and this six-symbol construction force exactly three fixed-length encoding bits.

The three-bit encoding is uniform on six codewords, not on all eight binary strings. It is therefore admissible in Definition 1, but is not directly a $b = 3$ auxiliary primitive for the compiler. The proposition does not decide whether five alphabet symbols suffice.

5.5 Dependence on the secret alphabet

Kishimoto, Okada, Kurosawa, and Ogata [10, Theorem 3.1] proved the following bound for $2 \leq t < n$. The direct proof below specializes to our exact-uniform model and also covers $t = n$.

Proposition 3 (Classical secret-alphabet lower bound). *Every fully anonymous perfect (t, n) scheme with $2 \leq t \leq n$ satisfies*

$$V \geq 1 + (n - t + 1)(M - 1), \qquad M = |\Sigma|. \tag{28}$$

Proof. Fix the shares of $t - 2$ participants to arbitrary values and condition on that event separately for each secret. Its probability is $V^{-(t-2)}$, and every remaining share is uniform by the original $(t - 1)$-wise uniformity. Appending the fixed values to reconstruction gives a fully anonymous perfect $(2, N)$ scheme on the same alphabets, where $N = n - t + 2$. The argument includes $t = 2$, with no conditioning.

Let $D : \Omega^2 \to \Sigma$ be its total symmetric decoder. Fix $a \in \Omega$, put $s_0 = D(a, a)$, and define

$$B_s(a) = \{x \in \Omega : D(a, x) = s\} \qquad (s \in \Sigma).$$

These sets partition Ω, and $a \in B_{s_0}(a)$. For a secret $s \neq s_0$, the symbol a occurs at most once in any supported output; otherwise (a, a) would reconstruct to the wrong secret. Individual uniformity therefore gives $\Pr[a \text{ occurs} \mid s] = N/V$. Whenever it occurs, all other $N - 1$ shares lie in $B_s(a)$. If C_s counts the shares in this class, then

$$\frac{N|B_s(a)|}{V} = \mathbb{E}[C_s \mid s] \geq (N - 1)\frac{N}{V}.$$

Hence $|B_s(a)| \geq N - 1$ for each $s \neq s_0$, while $|B_{s_0}(a)| \geq 1$. Summing the partition sizes proves $V \geq 1 + (M - 1)(N - 1)$.

Corollary 3 (A nonattained rate-one supremum). *For each fixed $2 \le t < n$, the supremum of the information rates of fully anonymous perfect (t, n) schemes is one, but no scheme with a finite secret alphabet of size greater than one attains it.*

Proof. For $t < n$ and $M > 1$, Proposition 3 implies $V > M$, so $\log M / \log V < 1$. Corollary 1 supplies a sequence of finite schemes with rates tending to one. At $t = n$, outside this assertion, additive sharing attains rate one.

For an L-bit secret, every fixed-length share encoding consequently obeys

$$\ell_L \ge \left\lceil \log_2\bigl(1 + (n - t + 1)(2^L - 1)\bigr) \right\rceil . \tag{29}$$

Before integer rounding, the necessary overhead tends to $\log_2(n - t + 1)$. Thus the compiler attains the optimal leading coefficient of L, but its $O(t^2 \log^2 n)$ additive overhead is not proved optimal.

6 Remaining quantitative gaps

The results leave two distinct share-size questions.

One-bit shares beyond threshold three. At $(n, t) = (1024, 16)$, the universal bound forces twenty fixed-length bits, whereas Theorem 1 uses 290. Reducing the latter within our construction requires preserving the local-relation identity, positivity, and an efficient exact sampler. A different construction need not obey these architectural constraints; the universal bound does not establish the necessity of an $O(t \log n)$ share length when t grows.

Additive overhead for long secrets. For the same parameters, the compiler adds 47,850 bits to an aligned payload, while the classical converse forces only ten additional encoding bits for $L \ge 2$. The coefficient of L is optimal, but the overhead's dependence on n, t is not. A multi-bit auxiliary scheme could reduce the cost of sharing the certificate; a different decoder could require a shorter certificate. These are separate targets from further improving the one-bit sampler.

The operation counts prove polynomial-time implementability, not practical performance. An implementation should measure reusable preprocessing separately from online cost and retain exact, untruncated sampling.

A Implementing the exact relation-space sampler

This appendix completes the algorithmic guarantee of Theorem 1: it gives the finite type weights, an exact sampler for (1), and the operation counts. Positivity, reconstruction, and anonymity are proved in Section 3. Throughout, $k = t - 1$, $m = \lceil \log_2 n \rceil$, $q = 2^m$, $D = 2k - 1$, and $Q = q^D$.

A.1 Counting the type masses

The Gaussian coefficients of Section 2.2 satisfy

$$G(a, j) = \prod_{i=0}^{j-1} \frac{q^a - q^i}{q^j - q^i}, \qquad G(a, j) = G(a-1, j) + q^{a-j} G(a-1, j-1).$$

The first formula counts ordered bases; the second separates subspaces inside a fixed hyperplane from the others. With $G(a, 0) = 1$, the recurrence gives

$$\sum_{j=0}^{a} G(a, j)(-1)^j q^{\binom{j}{2}} = \prod_{i=0}^{a-1} (1 - q^i) = \mathbf{1}_{a=0}.$$

Hence $\mu(U, S) = (-1)^j q^{\binom{j}{2}}$, where $j = \dim S - \dim U$, inverts containment sums on each subspace-lattice interval [16]. Every interval in $\mathcal{P}_k$ is a full such interval, which justifies the cancellation used in Lemma 1. For $q \geq 4$, the same product formula gives

$$G(a, j) \leq q^{j(a-j)} \prod_{\ell=1}^{j} (1 - q^{-\ell})^{-1} \leq \frac{q^{j(a-j)}}{1 - 1/(q-1)} < 2q^{j(a-j)},$$

using $\prod_\ell (1 - x_\ell) \geq 1 - \sum_\ell x_\ell$.

A nonzero relation space has type $(b, r) = (|\operatorname{supp} R|, \dim R)$ with $1 \leq r < b \leq k$. Inclusion–exclusion gives the number on a fixed support:

$$F(b, r) = \sum_{h=0}^{b-1} (-1)^h \binom{b}{h} G(b - h - 1, r). \tag{30}$$

For a fixed such R and a fixed c-set containing its support, the number of dimension-d superspaces with exactly that support is

$$E_{b,r}(c, d) = \sum_{h=0}^{c-b} (-1)^h \binom{c-b}{h} G(c - h - 1 - r, d - r). \tag{31}$$

Indeed, quotient by R and inclusion–exclude the $c - b$ new coordinates; omitting h of them leaves ambient dimension $c - h - 1 - r$. Thus the superspace counts

and the per-space masses in (1) are

$$N_{b,r,j} = \sum_{c=b}^{k} \binom{n-b}{c-b} E_{b,r}(c, r+j), \tag{32}$$

$$p_{b,r} = \sum_{j=0}^{k-1-r} (-1)^j q^{\binom{j}{2}} N_{b,r,j} Q^{-(r+j)}. \tag{33}$$

There are $\binom{n}{b}F(b,r)$ spaces of type (b,r), and the zero space has mass

$$p_0 = 1 - \sum_{b=2}^{k} \sum_{r=1}^{b-1} \binom{n}{b} F(b,r) p_{b,r}. \tag{34}$$

With $Z = Q^{k-1}$, store the nonnegative integer weights $w_0 = Zp_0$ and $w_{b,r} = Z\binom{n}{b}F(b,r)p_{b,r}$ and their cumulative sums. They sum to Z, a power of two. Lemma 1 proves their positivity.

A.2 Sampling without enumerating relation spaces

An exact uniform integer in a range of size a is obtained from $\lceil \log_2 a \rceil$ fair bits by rejecting values at least a; a singleton needs no bits. The following procedure returns $R = 0$ immediately when $k = 1$.

1. Draw $U \in \{0, \ldots, Z-1\}$ with $\log_2 Z$ fair bits and locate its cumulative-weight interval by binary search. Return $R = 0$ if the type is zero.
2. For type (b,r), select a uniform b-subset $B \subseteq [n]$ by b steps of partial Fisher–Yates shuffling of $(1, \ldots, n)$; sort the support.
3. Draw a uniform $r \times (b-1)$ matrix over $\mathbb{F}_q$ and append each row's negative sum to form an $r \times b$ matrix A. Repeat this step until $\operatorname{rank} A = r$ and no column is zero.
4. Return the row span of A embedded on B, storing the support and local basis rather than an n-column matrix.

Lemma 8 (Exact type sampler). *The procedure samples the law* (1)*. Conditional on* (b,r) *and* B*, its output is a uniform full-support* r*-subspace of* H_B*. The local-basis rejection loop accepts with probability at least* $1/4$*.*

Proof. The type and support probabilities are those counted above. Every ordered basis is equally likely before rejection, and every accepted subspace has exactly $|\operatorname{GL}_r(\mathbb{F}_q)|$ ordered bases. This proves the conditional law.

For $r \geq 2$, rank failure is at most $\sum_{i=0}^{r-1} q^{i-(b-1)} \leq 1/(q-1)$; a fixed column is zero with probability q^{-r}. Since $b \leq q-1$ and $q \geq 4$, total failure is at most $1/(q-1) + (q-1)q^{-r} \leq 1/3 + 1/4$. For $r = 1$, let N_b count vectors in H_B with no zero coordinate. The recurrence $N_1 = 0$, $N_b = (q-1)^{b-1} - N_{b-1}$ gives

$$N_b = \frac{(q-1)^b + (q-1)(-1)^b}{q}.$$

Its acceptance probability N_b/q^{b-1} is at least $1/3$ for even b, and at least $1/3 - (q-1)/q^3 > 1/4$ for odd $b \geq 3$, using $(1-1/q)^{q-1} > 1/3$. The latter inequality follows by reciprocating $(1+1/(q-1))^{q-1} < \sum_{j\geq 0} 1/j! < 3$.

A.3 Preprocessing and online cost

All bounds below count bit operations with schoolbook arithmetic and include exact integer sampling. They justify the costs stated in Theorem 1.

Field setup. In a polynomial-basis representation $\mathbb{F}_q = \mathbb{F}_2[X]/(g)$, addition costs $O(m)$, and multiplication, reduction, and extended-Euclidean inversion cost $O(m^2)$. To construct g deterministically, test at most $q < 2n$ monic degree-m candidates. Require $X^{2^m} = X \bmod g$ and $\gcd(g, X^{2^j} - X) = 1$ for $1 \leq j \leq \lfloor m/2 \rfloor$. The equality forces squarefreeness and factor degrees dividing m; any reducible candidate has a factor of degree at most $m/2$ and fails a gcd test. Successive squaring and the gcd tests cost $O(m^3)$ per candidate, hence $O(nm^3)$ once. For $m = 1$, take $g = X$. The first n coefficient strings supply the public evaluation points.

Type table. Compute Gaussian coefficients by the recurrence above, small binomials by Pascal's recurrence, and $\binom{n-b}{j}$ by successive exact divisions; no table with n rows is needed. Evaluate (30)–(34) with denominator Z, replacing $Q^{-(r+j)}$ by $Q^{k-1-r-j}$ when computing $Zp_{b,r}$. The four indices of $E_{b,r}(c,d)$ and its summation index give $O(k^5)$ integer operations. Gaussian exponents are $O(k^2)$, support choices are at most n^k, and Z has $D(k-1)m$ bits, so all operands and intermediate absolute sums have $O(k^2 m)$ bits. The bit cost is therefore $O(t^9 m^2)$. Only $O(k^2)$ cumulative weights, using $O(k^4 m)$ bits, need be retained. No table is needed for $k = 1$.

Online sharing and reconstruction. Writing $\beta = O(k^2 m)$, type selection and support sampling cost expected

$$O\big(\beta \log(k+1) + nm + km^2 + k\log(k+1)m\big)$$

bit operations: these cover the type comparisons, shuffle initialization, exact range draws, and sorting. Polynomial evaluation costs $O(nk^2)$ field operations. For secret one, the local basis, quotient, and invertible matrix cost expected $O(k^3)$, and forming all powers and transformed points costs $O(nk^2)$. The affine-matrix rejection accepts with probability at least $1 - 1/(q-1) \geq 2/3$ for $q \geq 4$; the only $q = 2$ case is $n = t = 2$, with acceptance $1/2$. Together with Lemma 8, this gives expected sharing cost $O((nt^2 + t^3)m^2)$, including the integer work above. Gaussian elimination on the $(D+1) \times t$ augmented matrix and inspection of its nullspace give reconstruction cost $O(t^3 m^2)$. All rejection loops preserve the exact conditional laws and terminate almost surely; no finite worst-case fair-bit bound or truncated sampler is claimed.

References

1. Beimel, A.: Secret-sharing schemes: A survey. In: Coding and Cryptology – IWCC 2011. Lecture Notes in Computer Science, vol. 6639, pp. 11–46. Springer, Berlin, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20901-7_2
2. Bishop, A., Green, M., Ishai, Y., Jain, A., Lou, P.: Fully anonymous secret sharing. In: Advances in Cryptology – CRYPTO 2025. Lecture Notes in Computer Science, vol. 16003, pp. 356–389. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01884-7_12, https://eprint.iacr.org/2025/1984
3. Blakley, G.R.: Safeguarding cryptographic keys. In: Proceedings of the 1979 AFIPS National Computer Conference. AFIPS Conference Proceedings, vol. 48, pp. 313–317. AFIPS Press (1979). https://doi.org/10.1109/AFIPS.1979.98
4. Blundo, C., Stinson, D.R.: Anonymous secret sharing schemes. Discrete Applied Mathematics **77**(1), 13–28 (1997). https://doi.org/10.1016/S0166-218X(97)89208-6
5. Con, R.: Anonymous Shamir's secret-sharing via Reed–Solomon codes against permutations, insertions, and deletions. IEEE Transactions on Information Theory **71**(12), 9534–9547 (2025). https://doi.org/10.1109/TIT.2025.3622889, https://arxiv.org/abs/2412.17003
6. Con, R., Li, R., Mazor, N.: Fully anonymous perfect secret-sharing. IACR Cryptology ePrint Archive, Report 2026/2145 (2026), https://eprint.iacr.org/2026/2145, revision of 23 September 2026
7. Deng, Y.p., Guo, L.f., Liu, M.l.: Constructions for anonymous secret sharing schemes using combinatorial designs. Acta Mathematicae Applicatae Sinica, English Series **23**(1), 67–78 (2007). https://doi.org/10.1007/s10255-006-0350-3
8. Eldridge, H., Beck, G., Green, M., Heninger, N., Jain, A.: Abuse-resistant location tracking: Balancing privacy and safety in the offline finding ecosystem. In: 33rd USENIX Security Symposium (USENIX Security 24). pp. 5431–5448. USENIX Association, Philadelphia, PA, USA (2024), https://www.usenix.org/conference/usenixsecurity24/presentation/eldridge
9. Guillermo, M., Martin, K.M., O'Keefe, C.M.: Providing anonymity in unconditionally secure secret sharing schemes. Designs, Codes and Cryptography **28**(3), 227–245 (2003). https://doi.org/10.1023/A:1024198519111
10. Kishimoto, W., Okada, K., Kurosawa, K., Ogata, W.: On the bound for anonymous secret sharing schemes. Discrete Applied Mathematics **121**(1–3), 193–202 (2002). https://doi.org/10.1016/S0166-218X(01)00236-0
11. Kötter, R., Kschischang, F.R.: Coding for errors and erasures in random network coding. IEEE Transactions on Information Theory **54**(8), 3579–3591 (2008). https://doi.org/10.1109/TIT.2008.926449, https://arxiv.org/abs/cs/0703061
12. Kovačević, M., Tan, V.Y.F.: Codes in the space of multisets—coding for permutation channels with impairments. IEEE Transactions on Information Theory **64**(7), 5156–5169 (2018). https://doi.org/10.1109/TIT.2017.2789292, https://arxiv.org/abs/1612.08837
13. Krawczyk, H.: Secret sharing made short. In: Advances in Cryptology – CRYPTO '93. Lecture Notes in Computer Science, vol. 773, pp. 136–146. Springer, Berlin, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_12
14. Paskin-Cherniavsky, A., Olimid, R.F.: On cryptographic anonymity and unpredictability in secret sharing. Information Processing Letters **161**, 105965 (2020). https://doi.org/10.1016/j.ipl.2020.105965

15. Phillips, S.J., Phillips, N.C.: Strongly ideal secret sharing schemes. Journal of Cryptology **5**(3), 185–191 (1992). https://doi.org/10.1007/BF02451114
16. Rota, G.C.: On the foundations of combinatorial theory I. Theory of Möbius functions. Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete **2**, 340–368 (1964). https://doi.org/10.1007/BF00531932
17. Shamir, A.: How to share a secret. Communications of the ACM **22**(11), 612–613 (1979). https://doi.org/10.1145/359168.359176
18. Stinson, D.R., Vanstone, S.A.: A combinatorial approach to threshold schemes. SIAM Journal on Discrete Mathematics **1**(2), 230–236 (1988). https://doi.org/10.1137/0401024
19. Zhang, Y., Sun, Y., Zhang, X., Ge, G.: Random Reed–Solomon codes correcting permutations, insertions, and deletions over polynomial-size alphabets. arXiv preprint arXiv:2606.22344 (2026), https://arxiv.org/abs/2606.22344