Rank Measures and Exponential Lower Bounds for Multilinear Secret Sharing

Shahram Khazaei

Department of Mathematical Sciences, Sharif University of Technology, Tehran, Iran
shahram.khazaei@sharif.ir

September 2026

Abstract

Multilinear secret sharing can amortize share size over a vector secret, so scalar linear lower bounds need not survive normalization by the secret dimension. We prove that the Razborov–Gál rank measure does survive this amortization: every multi-target monotone span program satisfies a rank inequality in which the target dimension appears multiplicatively. Combined with the characteristic-separated witnesses of Pitassi and Robere, this gives an explicit family of access structures whose average and maximum multilinear information ratios are exponential over every finite field, settling the worst-case order at $2^{\Theta(n)}$. We also study arbitrary sharing algorithms with affine-linear reconstruction under non-perfect security notions, including the pairwise statistical security of Beimel–Othman–Peter and the partial privacy of Jafari–Khazaei. Combining these affine-reconstruction results with the degree-reduction theorem of Beimel–Othman–Peter gives corresponding fixed-degree lower bounds under both security notions, in particular for subexponential secret dimension.

1 Introduction

A secret-sharing scheme distributes a secret among a set of n participants. The authorized sets can recover the secret, whereas the unauthorized sets should obtain no information about it. The family of authorized sets is called an access structure. If the secret consists of d_0 field elements and participant i receives d_i field elements, then the information ratio of participant i is $\rho_i = d_i/d_0$. Let

$$\mathrm{cv}(\Pi) = (\rho_1, \ldots, \rho_n)$$

denote the *contribution vector*, or *convec*, of Π, and

$$\rho_{\max} = \max_i \rho_i, \qquad \rho_{\mathrm{av}} = \frac{1}{n}\sum_i \rho_i, \qquad \rho_{\mathrm{tot}} = \sum_i \rho_i = \frac{d}{d_0},$$

where $d = \sum_i d_i$. The first two are the standard efficiency measures.

In a scalar linear scheme the secret is one field element and every share is linear in the secret and fresh randomness. A multilinear scheme shares a vector of d_0 field elements in the same way. Allowing d_0 to grow is a form of amortization, and it can strictly reduce the normalized share size for a fixed access structure [2, 7]. A lower bound proved for $d_0 = 1$ therefore need not survive normalization by d_0.

For $d_0 = 1$ the relevant linear-algebraic model is the monotone span program (MSP) [22]. Exponential lower bounds are known: Robere, Pitassi, Rossman, and Cook proved them over the reals [28], and Pitassi and Robere over every field [26]. Both use the Razborov–Gál rank measure [27, 17, 18]. For multilinear schemes, the best previously known explicit lower bound was quasipolynomial, $n^{\Omega(\log n)}$, obtained by Beimel, Ben-Efraim, Padró, and Tyomkin (BBPT) by extending a unique-intersection rank argument to multi-target monotone span programs (MTMSPs) [8]. Beimel's recent monograph records this state of affairs and asks whether the worst-case multilinear information ratio is exponential [7, Question 10.4]. We answer that question.

1.1 Amortized rank theorem and exponential lower bound

Let $f : \{0,1\}^n \to \{0,1\}$ be monotone, let $U = f^{-1}(1)$ and $V = f^{-1}(0)$, and let A be a matrix over a field $\mathbb{F}$ with rows indexed by U and columns by V. Participant i separates a pair $(u, v) \in U \times V$ when $u_i = 1$ and $v_i = 0$; write R_i for the set of such pairs and $A|_{R_i}$ for the matrix obtained by setting all entries outside R_i to zero. The Razborov–Gál rank measure is

$$\mu_{\mathbb{F}}(f) = \max_{A \neq 0} \frac{\operatorname{rank}_{\mathbb{F}} A}{\max_i \operatorname{rank}_{\mathbb{F}}(A|_{R_i})}.$$

The numerator measures a global obstruction, while the denominator measures the largest part of that obstruction visible to one participant.

A multi-target monotone span program (MTMSP) is the matrix model for multilinear secret sharing introduced by BBPT [8]. Its rows are labeled by participants and it has a d_0-dimensional target space. An authorized set must span the whole target, while an unauthorized set must meet the target trivially. Participant i labels d_i rows when its share has d_i field coordinates.

The following is the main result of the paper. For $d_0 = 1$ it is the standard Razborov–Gál lower bound for MSPs [27, 17, 18]; the new content is that the same obstruction scales linearly with an arbitrary target dimension.

Theorem 1.1 (Amortized rank-measure theorem)**.** *Let $\mathbb{F}$ be any field, and let a d_0-target MTMSP over $\mathbb{F}$ compute f. If participant i labels d_i rows, then every nonzero matrix $A \in \mathbb{F}^{U \times V}$ satisfies*

$$d_0 \operatorname{rank} A \leq \sum_{i=1}^{n} d_i \operatorname{rank}(A|_{R_i}).$$

Consequently,

$$\frac{\sum_i d_i}{d_0} \geq \mu_{\mathbb{F}}(f).$$

Combining the theorem with the characteristic-separated rank witnesses of Pitassi and Robere gives the following field-uniform consequence; see Theorems 4.6 and 4.7.

Theorem 1.2 (Exponential multilinear lower bound)**.** *There is an explicit family $\{\Gamma_n\}$ and constants $\alpha > 0$ and n_0 such that, for every finite field $\mathbb{F}$ and every $n \geq n_0$, every perfect multilinear scheme for Γ_n satisfies*

$$\rho_{\mathrm{av}} \geq \frac{2^{\alpha n}}{n} \qquad \text{and} \qquad \rho_{\max} \geq \frac{2^{\alpha n}}{n}.$$

Consequently, the worst-case multilinear information ratio is $2^{\Theta(n)}$.

1.2 Relation to previous rank arguments

The hard access structures in Section 4 are the characteristic-separated families of Pitassi and Robere [26]; the scalar all-field consequence is already recorded in [10, extended version, Theorem 4.10]. The new point is that the same Razborov–Gál obstruction survives arbitrary target dimension.

BBPT's earlier multilinear lower bound uses a unique-intersection witness for which every participant restriction has rank at most one [8, Theorem 4.5]. In that case Theorem 3.1 specializes to

$$d_0 \operatorname{rank} A \le \sum_i d_i \operatorname{rank}(A|_{R_i}) \le \sum_i d_i = d,$$

so $d/d_0 \ge \operatorname{rank} A$, which is exactly the BBPT unique-intersection form. The Razborov–Gál witnesses needed for exponential bounds have higher-rank participant restrictions. The new point in Theorem 3.1 is that these local ranks are retained rather than collapsed to one.

1.3 Beyond multilinear sharing

The paper also studies schemes whose sharing algorithm is arbitrary but whose authorized reconstruction functions are affine-linear. Perfect correctness alone is enough to replace such a scheme by a perfectly correct multilinear one with the same secret space, no larger participant dimensions, and no more information available to any coalition. Thus affine reconstruction, rather than linear sharing, is the structural property needed for the scheme-level extension of the rank method.

This replacement has two different security consequences. Under Jafari–Khazaei partial privacy [21], every unauthorized coalition still leaves a positive fraction of the secret unrecovered, which leads to the partial-target rank bound of Section 5.2. Under pairwise statistical privacy with parameter $\delta < 1$, any nonconstant affine-linear secret information recoverable by an unauthorized coalition would force statistical distance one between the views corresponding to two secrets; the support-disjointness step is the one used in BOP's Claim 10.8 [10, extended version, Claim 10.8]. Hence the multilinear replacement is perfect. Section 5 states these transformations precisely and compares them with BOP's linear-reconstruction transformation [10, extended version, Construction 10.6 and Theorem 10.9] and the Jafari–Khazaei partial-to-perfect transformation [21, Proposition 5.5].

Finally, Beimel, Othman, and Peter show how degree-r reconstruction can be replaced by multilinear sharing after enlarging each participant's share by at most a degree-dependent $O_r(d^r)$ number of field coordinates, where d is the original total share dimension [10, Lemma 4.7, extended version]. Their construction does not increase the information available to an unauthorized coalition. Combining it with the two security branches above gives fixed-degree lower bounds under both pairwise statistical privacy and partial privacy; in particular, subexponential secret dimension still gives exponential normalized lower bounds. See Section 6.

1.4 Related work

Secret sharing was introduced independently by Shamir [30] and Blakley [12]. General access structures were realized by Ito, Saito, and Nishizeki [19] and related to monotone formulas by Benaloh and Leichter [11]; see [6, 7] for surveys. For more than three decades the best universal secret-sharing constructions had information ratio $2^{n-o(n)}$. Liu and Vaikuntanathan improved this to $2^{0.994n+o(n)}$ for unrestricted schemes and $2^{0.999n+o(n)}$ for linear schemes [23]. Subsequent work reduced the unrestricted upper bound to $2^{0.892n+o(n)}$ [3], $2^{0.637n+o(n)}$ [4], and $2^{(\log_2(3/2))n+o(n)}$ [5].

Nir recently obtained $2^{0.496n+o(n)}$ for unrestricted schemes and, for linear schemes, the explicit upper bound $2^{\lceil n/2\rceil-1}+1$ [24]. Since scalar linear schemes are a special case of multilinear schemes, the latter is also a $2^{n/2+O(1)}$ upper bound for multilinear schemes. On the lower-bound side for unrestricted perfect schemes, Williamson gives $\Omega(n)$ maximum and $\Omega(n^2)$ total normalized share size, improving the previous general bounds of Csirmaz by a logarithmic factor [14, 31].

For linear secret sharing, monotone span programs provide the standard scalar model [22]. Exponential rank lower bounds were obtained over the reals in [28] and over every field in [26]. BBPT introduced multi-target monotone span programs and proved the first explicit superpolynomial lower bound for multilinear schemes [8]. Applebaum and Arkis exhibited access structures for which increasing the secret dimension strictly improves normalized complexity [2]. Thus a scalar lower bound must be shown to scale with the target dimension before it can be used against multilinear schemes.

Polynomial reconstruction was studied in [25, 10, 9]. BOP proved both the linear-reconstruction perfecting theorem and the fixed-degree reduction used here [10]. The partial-secret-sharing results of Jafari and Khazaei that are used here are described next to Table 1 and in Section 5.2.

1.5 Organization

Section 2 gives the secret-sharing, reconstruction, MTMSP, and rank-measure models. Section 3 proves the amortized rank theorem and states the partial-target form used later; Section 4 derives the field-uniform exponential lower bound. Section 5 treats affine reconstruction, while Section 6 gives the bounded-degree consequence. Section 7 discusses the remaining questions. Appendices A–C provide the proof-complexity background used in Section 4, the reconstruction-dual proof used in Section 5.2, and a concrete non-multilinear example for pairwise statistical distance below one.

2 Preliminaries

We first fix the secret-sharing and security terminology, then recall the standard multilinear sharing-map model and the usual matrix presentation of multi-target monotone span programs. Random variables are used for the information-theoretic security notions; matrices are used for the multilinear and rank arguments.

2.1 Access structures and general secret-sharing schemes

Let $P=[n]=\{1,\dots,n\}$ be the participant set. A family $\Gamma\subseteq 2^P$ is monotone if $A\in\Gamma$ and $A\subseteq B$ imply $B\in\Gamma$. An access structure is a nonempty monotone family with $\emptyset\notin\Gamma$. Sets in Γ are authorized; the other sets are unauthorized. The associated monotone Boolean function is

$$f_\Gamma(x)=1 \quad\Longleftrightarrow\quad \{i : x_i=1\}\in\Gamma.$$

The dual access structure is

$$\Gamma^*=\{A\subseteq P : P\setminus A\notin\Gamma\},$$

and the minimal authorized sets are

$$\min\Gamma=\{A\in\Gamma :\ \text{no proper subset of } A \text{ belongs to } \Gamma\}.$$

A secret-sharing scheme Π consists of a finite secret domain $\mathcal{S}_0$, finite share domains $\mathcal{S}_1, \dots, \mathcal{S}_n$, and a randomized sharing map

$$\mathrm{Share} : \mathcal{S}_0 \longrightarrow \mathcal{S}_1 \times \cdots \times \mathcal{S}_n.$$

Take X_0 uniform on $\mathcal{S}_0$, run the sharing algorithm conditionally on X_0, and denote the resulting share random variables by $X_1, \dots, X_n$. For $A \subseteq P$, put $X_A = (X_i)_{i \in A}$; equivalently, $\mathrm{Share}_A(s)$ denotes the conditional distribution of X_A given $X_0 = s$.

Correctness is defined relative to an access structure. The scheme is *perfectly correct* for Γ if, for every $A \in \Gamma$, there is a deterministic reconstruction function Rec_A such that

$$\Pr[\mathrm{Rec}_A(X_A) = X_0 \mid X_0 = s] = 1$$

for every $s \in \mathcal{S}_0$.

For distributions X, Y on the same finite set, their statistical distance (SD) is

$$\mathrm{SD}(X, Y) = \frac{1}{2} \sum_z |\Pr[X = z] - \Pr[Y = z]|.$$

Definition 2.1 (Pairwise statistical privacy)**.** A perfectly correct scheme has *pairwise statistical privacy parameter* δ for Γ if

$$\mathrm{SD}\big(\mathrm{Share}_B(s), \mathrm{Share}_B(s')\big) \le \delta$$

for every unauthorized B and every pair of secrets s, s'. Perfect privacy is the case $\delta = 0$.

Following Farràs, Hansen, Kaced, and Padró [16], define the *access function* of Π by

$$\Phi_\Pi(A) = \frac{I(X_0; X_A)}{H(X_0)} \qquad (A \subseteq P).$$

Following Jafari and Khazaei [21], we use the following notion of partial privacy.

Definition 2.2 (Partial privacy)**.** A scheme Π is *partially private* for Γ if it has perfect correctness and

$$\eta = \eta(\Pi; \Gamma) := 1 - \max_{B \notin \Gamma} \Phi_\Pi(B) > 0.$$

Equivalently,

$$\eta = \min_{B \notin \Gamma} \frac{H(X_0 \mid X_B)}{H(X_0)}.$$

We call η the *partial-privacy advantage*. We use η rather than the symbol δ of [21] to distinguish this information gap from the pairwise statistical-privacy parameter in Definition 2.1.

The information ratio of participant i is

$$\rho_i(\Pi) = \frac{\log |\mathcal{S}_i|}{\log |\mathcal{S}_0|}.$$

We use alphabet sizes in this definition.[1] The *contribution vector*, or *convec*, of Π is

$$\mathrm{cv}(\Pi) = (\rho_1(\Pi), \dots, \rho_n(\Pi)),$$

and

$$\rho_{\max} = \max_i \rho_i, \qquad \rho_{\mathrm{av}} = \frac{1}{n} \sum_{i=1}^{n} \rho_i, \qquad \rho_{\mathrm{tot}} = \sum_{i=1}^{n} \rho_i.$$

The choice of logarithm base is immaterial. For schemes whose alphabets are powers of $\mathbb{F}$, this is the field-coordinate convec used throughout the lower-bound statements below.

[1] Another common convention measures a share by entropy and uses $H(X_i)/H(X_0)$; in particular, the convec in [20, 21] is entropy-based. We use alphabet sizes because the rank lower bounds in this paper count field coordinates.

2.2 Classes of sharing and reconstruction

For schemes over $\mathbb{F}$ whose secret and share alphabets are

$$\mathcal{S}_0 = \mathbb{F}^{d_0}, \qquad \mathcal{S}_i = \mathbb{F}^{d_i} \quad (i \in [n]), \qquad d = \sum_{i=1}^{n} d_i,$$

the convec is

$$\mathrm{cv}(\Pi) = \left(\frac{d_1}{d_0}, \ldots, \frac{d_n}{d_0} \right).$$

We use the following classes throughout the paper.

- A *scalar-linear scheme* has $d_0 = 1$ and its sharing map is $\mathbb{F}$-linear in the secret and independent uniform random field elements.
- A *multilinear scheme* is defined in the same way with arbitrary secret dimension $d_0 \geq 1$.
- An *affine-reconstructable scheme* may have an arbitrary sharing algorithm, but its reconstruction functions are affine-linear in the share coordinates supplied to them.
- A *degree-r reconstructable scheme* may have an arbitrary sharing algorithm, but its reconstruction functions are polynomials of total degree at most r in the share coordinates supplied to them.

When the affine reconstruction functions have zero constant term, we say that the scheme is *linearly reconstructable*.

2.3 Multi-target monotone span programs

Definition 2.3 (Multi-target monotone span program; BBPT [8])**.** A d_0-target monotone span program over $\mathbb{F}$ consists of a matrix

$$M \in \mathbb{F}^{d \times b},$$

a labeling $\pi : [d] \to [n]$ of its rows by participants, and d_0 linearly independent target vectors spanning a subspace $T_0 \leq \mathbb{F}^b$. We assume $1 \leq d_0 < b$; by Convention 2.4 below this is no loss. For $A \subseteq P$, let M_A be the rows labeled by participants in A. The program

- accepts A if $T_0 \subseteq \mathrm{rowspan}(M_A)$;
- rejects A if $\mathrm{rowspan}(M_A) \cap T_0 = \{0\}$.

It computes Γ if it accepts exactly the authorized sets and rejects exactly the unauthorized sets.

Convention 2.4 (Zero columns)**.** The strict inequality $d_0 < b$ follows [8]. Whenever it is not already satisfied, we append one zero column to the program matrix and to every target vector. This changes neither the row counts nor acceptance or rejection. We use this convention without further comment.

After an invertible change of coordinates in $\mathbb{F}^b$, we may and shall take

$$E = (I_{d_0}\ 0) \in \mathbb{F}^{d_0 \times b}, \qquad T_0 = \mathrm{rowspan}(E).$$

If participant i labels d_i rows, then $d = \sum_i d_i$. The normalized total size is d/d_0.

The correspondence between perfect multilinear schemes and multi-target span programs is standard; see BBPT [8] and Beimel [7, Claim 4.15]. We state the direction used for the lower bound and include the short proof because it makes explicit how perfect privacy gives the strong-rejection condition used throughout Section 3.

Lemma 2.5 (From multilinear schemes to MTMSPs)**.** *Let Π be a perfect multilinear secret-sharing scheme over $\mathbb{F}$ realizing Γ. Suppose the secret lies in $\mathbb{F}^{d_0}$ and participant i receives d_i field elements. Then there is a d_0-target MTMSP computing Γ in which participant i labels exactly d_i rows.*

Proof. Write the full share vector as

$$Mx, \qquad x = (s, r) \in \mathbb{F}^b,$$

where the first d_0 coordinates of x are the secret coordinates, and let

$$E = (I_{d_0}\ 0), \qquad T_0 = \mathrm{rowspan}(E).$$

Apply Convention 2.4 if necessary. The rows of M are already partitioned by participants, with exactly d_i rows belonging to participant i.

Let A be authorized. Perfect reconstruction implies that two vectors x, x' with $M_A x = M_A x'$ must have the same secret, and hence

$$\ker M_A \subseteq \ker E.$$

Taking annihilators gives $T_0 = \mathrm{rowspan}(E) \subseteq \mathrm{rowspan}(M_A)$, so A is accepted.

Now let B be unauthorized and write

$$M_B = (S_B\ R_B),$$

according to secret and random coordinates. Since r is uniform, the view of B on secret s is uniform on the affine coset

$$S_B s + \mathrm{colspan}(R_B).$$

Perfect privacy for all secrets is therefore equivalent to

$$\mathrm{colspan}(S_B) \subseteq \mathrm{colspan}(R_B).$$

Suppose $y \in \mathrm{rowspan}(M_B) \cap T_0$. Then for some row vector a,

$$y = aM_B = (u\ 0).$$

Thus $aR_B = 0$. Since every column of S_B lies in the column span of R_B, also $aS_B = 0$, and therefore $u = 0$. Hence $y = 0$ and

$$\mathrm{rowspan}(M_B) \cap T_0 = \{0\}.$$

Thus every unauthorized set is strongly rejected. □

For a monotone function f, let $s^{\mathbb{F}}_{d_0}(f)$ be the minimum number of rows of a d_0-target MTMSP computing f. Such programs exist for every d_0: take d_0 independent copies of a one-target program. We call

$$\mathrm{amSP}_{\mathbb{F}}(f) = \inf_{d_0 \geq 1} \frac{s^{\mathbb{F}}_{d_0}(f)}{d_0}$$

the *amortized monotone span-program complexity* of f. For $d_0 = 1$ this reduces to ordinary MSP complexity. The reverse MTMSP-to-multilinear construction is standard [8] and will be invoked only where it is needed in Theorem 5.6.

2.4 The Razborov–Gál rank measure

Let $f : \{0,1\}^n \to \{0,1\}$ be monotone, and write

$$U = f^{-1}(1), \qquad V = f^{-1}(0).$$

For $i \in [n]$ define

$$R_i = \{(u,v) \in U \times V : u_i = 1,\ v_i = 0\}.$$

Monotonicity implies that the rectangles R_i cover $U \times V$: if $u \in U$ and $v \in V$, then $u \not\leq v$, so $u_i = 1, v_i = 0$ for some i.

If $A \in \mathbb{F}^{U \times V}$, let $A|_{R_i}$ be the matrix obtained by retaining $A[u,v]$ on R_i and putting zero elsewhere. Thus $A|_{R_i}$ has the same row and column index sets as A: it is a mask, not a submatrix formed by deleting rows or columns.

Definition 2.6 (Razborov–Gál rank measure)**.** For $A \neq 0$, define

$$\mu_{\mathbb{F}}(f, A) = \frac{\operatorname{rank}_{\mathbb{F}} A}{\max_{i \in [n]} \operatorname{rank}_{\mathbb{F}}(A|_{R_i})}, \qquad \mu_{\mathbb{F}}(f) = \max_{A \neq 0} \mu_{\mathbb{F}}(f, A).$$

For an access structure Γ, write $\mu_{\mathbb{F}}(\Gamma) = \mu_{\mathbb{F}}(f_\Gamma)$.

The witness A is an auxiliary matrix used to lower-bound span-program size; it is not part of the MSP or MTMSP. A pair (u,v) belongs to R_i exactly when participant i separates the positive input u from the negative input v, that is, when $i \in u \setminus v$. Informally, $\operatorname{rank} A$ measures the independent global patterns carried by the witness, while $\operatorname{rank}(A|_{R_i})$ measures the part that remains when one keeps only distinctions assigned to participant i.

The denominator in Definition 2.6 is nonzero whenever $A \neq 0$, because the R_i cover every entry. This measure is the standard Razborov–Gál lower bound for MSPs; see [27, 17, 18].

3 Rank lower bounds for multi-target span programs

The scalar Razborov–Gál argument must be strengthened so that its global rank obstruction grows in proportion to the target dimension. The proof achieves this by assigning each participant a low-rank block matrix: for every positive–negative pair these blocks sum to the identity on the target, while participant i contributes only on pairs separated by i.

Theorem 3.1 (Rank measure lower-bounds normalized MTMSP size)**.** *Let $\mathbb{F}$ be any field, let f be monotone, and let a d_0-target MTMSP over $\mathbb{F}$ compute f. If participant i labels d_i rows and $d = \sum_i d_i$, then for every nonzero $A \in \mathbb{F}^{U \times V}$,*

$$d_0 \operatorname{rank} A \leq \sum_{i=1}^{n} d_i \operatorname{rank}(A|_{R_i}).$$

Consequently,

$$\frac{d}{d_0} \geq \frac{\operatorname{rank} A}{\max_i \operatorname{rank}(A|_{R_i})}.$$

Maximizing over A and then over the target dimension gives

$$\operatorname{amSP}_{\mathbb{F}}(f) \geq \mu_{\mathbb{F}}(f).$$

The proof uses two elementary linear-algebraic certificates associated with accepted and rejected inputs. For a positive input u, a reconstruction certificate L_u records coefficients that express all d_0 target vectors simultaneously as combinations of rows available to u. For a negative input v, a rejection certificate W_v annihilates every row available to v while acting as the identity on the target. The product $L_{u,i}M_iW_v$ therefore isolates the contribution of participant i to the distinction between u and v.

Lemma 3.2 (Simultaneous reconstruction)**.** *For every $u \in U$ there is $L_u \in \mathbb{F}^{d_0 \times d}$ such that*

$$L_u M = E,$$

and every column of L_u corresponding to a row unavailable on u is zero.

Proof. Since u is accepted, every row of E belongs to $\mathrm{rowspan}(M_u)$. Choose a reconstruction coefficient vector for each of the d_0 rows of E, pad it by zero coefficients outside the rows available on u, and stack the d_0 vectors. □

Lemma 3.3 (Simultaneous rejection)**.** *For every $v \in V$ there is $W_v \in \mathbb{F}^{b \times d_0}$ satisfying*

$$M_v W_v = 0, \qquad E W_v = I_{d_0}.$$

Proof. Put $S = \mathrm{rowspan}(M_v)$ and let

$$S^\perp = \{w \in \mathbb{F}^b : M_v w = 0\}.$$

Consider the linear map $E : S^\perp \to \mathbb{F}^{d_0}$. We claim that it is onto. If $y \in \mathbb{F}^{d_0}$ annihilates $E(S^\perp)$, then

$$E^T y \in (S^\perp)^\perp = S.$$

On the other hand, $E^T y \in T_0$. Since v is rejected, $S \cap T_0 = \{0\}$, and hence $E^T y = 0$. As E has full row rank, $y = 0$. Thus the image $E(S^\perp)$ has trivial annihilator and equals $\mathbb{F}^{d_0}$.

For every standard basis vector $e_s \in \mathbb{F}^{d_0}$, choose $w_s \in S^\perp$ with $E w_s = e_s$. Then $W_v = [w_1 \ \cdots \ w_{d_0}]$ has the required properties. □

Let $M_i \in \mathbb{F}^{d_i \times b}$ denote the row block labeled by participant i. Let $L_{u,i} \in \mathbb{F}^{d_0 \times d_i}$ be the corresponding block of columns in L_u. For $u \in U$, $v \in V$, define

$$B_i[u, v] = L_{u,i} M_i W_v \in \mathbb{F}^{d_0 \times d_0}.$$

In this paper, we refer to the resulting block matrix B_i as the tile of participant i. We regard it as an ordinary matrix whose rows are indexed by $(u, a) \in U \times [d_0]$, whose columns are indexed by $(v, b) \in V \times [d_0]$, and whose $((u, a), (v, b))$-entry is the (a, b)-entry of the block $B_i[u, v]$.

Lemma 3.4 (Participant tiles)**.** *For every participant i, its tile B_i satisfies:*

(i) $B_i[u, v] = 0$ *whenever* $(u, v) \notin R_i$*;*

(ii) $\sum_{i=1}^{n} B_i[u, v] = I_{d_0}$ *for every* $(u, v) \in U \times V$*;*

(iii) $\operatorname{rank} B_i \le d_i$*.*

Proof. If $u_i = 0$, then the i-block of the reconstruction matrix vanishes: $L_{u,i} = 0$. If $v_i = 1$, every row of M_i belongs to M_v, so $M_i W_v = 0$ by Lemma 3.3. This proves (i).

For (ii), use the block decomposition of $L_u M$:

$$\sum_i B_i[u, v] = \sum_i L_{u,i} M_i W_v = L_u M W_v = E W_v = I_{d_0}.$$

For (iii), vertically stack the matrices $L_{u,i}$ over $u \in U$ and horizontally concatenate the matrices $M_i W_v$ over $v \in V$:

$$P_i = \begin{pmatrix} L_{u_1,i} \\ L_{u_2,i} \\ \vdots \end{pmatrix} \in \mathbb{F}^{d_0|U| \times d_i}, \qquad Q_i = \begin{pmatrix} M_i W_{v_1} & M_i W_{v_2} & \cdots \end{pmatrix} \in \mathbb{F}^{d_i \times d_0|V|}.$$

With the preceding row and column indexing, $B_i = P_i Q_i$. Therefore $\operatorname{rank} B_i \le d_i$. This also covers $d_i = 0$: the factorization is then through the zero-dimensional space and $B_i = 0$. □

We use the standard inequality

$$\operatorname{rank}(X \circ Y) \le \operatorname{rank}(X) \operatorname{rank}(Y),$$

for the Hadamard product $X \circ Y$. It holds over every field: write each matrix as a sum of rank-one matrices and distribute the entrywise product.

Proof of Theorem 3.1. Let $A_i = A|_{R_i}$ and let J_{d_0} be the all-ones $d_0 \times d_0$ matrix. Define the matrix C_i blockwise by

$$C_i[u, v] = A[u, v] B_i[u, v].$$

Thus C_i is the tile of participant i weighted by the rank witness A. By Lemma 3.4(i),

$$C_i = B_i \circ (A_i \otimes J_{d_0}).$$

Since the nonzero all-ones matrix J_{d_0} has rank one over every field, including when the characteristic divides d_0,

$$\operatorname{rank}(A_i \otimes J_{d_0}) = \operatorname{rank} A_i.$$

Therefore

$$\operatorname{rank} C_i \le \operatorname{rank} B_i \operatorname{rank}(A_i \otimes J_{d_0}) \le d_i \operatorname{rank} A_i.$$

On the other hand, Lemma 3.4(ii) gives

$$\sum_i C_i = A \otimes I_{d_0}.$$

Hence

$$d_0 \operatorname{rank} A = \operatorname{rank}(A \otimes I_{d_0}) \le \sum_i \operatorname{rank} C_i \le \sum_i d_i \operatorname{rank} A_i \le d \max_i \operatorname{rank} A_i.$$

The first three terms prove the first inequality in the theorem. The last inequality and division by $d_0 \max_i \operatorname{rank} A_i$ give the normalized-size bound. Maximizing over $A \neq 0$, then taking the infimum over the target dimension, proves $\mathrm{amSP}_{\mathbb{F}}(f) \ge \mu_{\mathbb{F}}(f)$. □

Theorem 3.1 assumes that every positive input spans the full d_0-dimensional target. In Section 5, affine reconstruction will naturally produce a linear program in which every positive input is guaranteed only a λ-dimensional intersection with the target. We therefore state the corresponding rank inequality here. It also follows by combining Jafari–Khazaei's partially-correct-to-perfect transformation [21, Proposition 5.4] with Theorem 3.1; the direct proof below avoids that amplification step and keeps the rank mechanism visible.

Theorem 3.5 (Partial-target form of the rank bound)**.** *Let $\mathbb{F}$ be any field and let f be monotone with positive inputs U and negative inputs V. Let M be a labeled matrix with a d_0-dimensional target space T_0, participant i labeling d_i rows. Suppose that for some integer $1 \le \lambda \le d_0$,*

$$\dim\big(T_0 \cap \operatorname{rowspan}(M_u)\big) \ge \lambda \qquad (u \in U),$$

while every negative input is strongly rejected:

$$T_0 \cap \operatorname{rowspan}(M_v) = \{0\} \qquad (v \in V).$$

Then every nonzero $A \in \mathbb{F}^{U \times V}$ satisfies

$$\lambda \operatorname{rank} A \le \sum_{i=1}^{n} d_i \operatorname{rank}(A|_{R_i}).$$

Consequently,

$$\frac{d}{\lambda} \ge \mu_{\mathbb{F}}(f), \qquad d = \sum_i d_i.$$

Proof. Choose coordinates so that $T_0 = \operatorname{rowspan}(E)$ with

$$E = (I_{d_0}\ 0) \in \mathbb{F}^{d_0 \times b}.$$

For each $u \in U$, choose a λ-dimensional subspace of $T_0 \cap \operatorname{rowspan}(M_u)$. It has a basis of the form $P_u E$, where $P_u \in \mathbb{F}^{\lambda \times d_0}$ has row rank λ. Choose $L_u \in \mathbb{F}^{\lambda \times d}$, supported only on rows available to u, such that

$$L_u M = P_u E.$$

For every $v \in V$, the rejection argument of Lemma 3.3 gives $W_v \in \mathbb{F}^{b \times d_0}$ with

$$M_v W_v = 0, \qquad E W_v = I_{d_0}.$$

Writing $L_{u,i}$ and M_i for the participant-i blocks, define

$$B_i[u, v] = L_{u,i} M_i W_v \in \mathbb{F}^{\lambda \times d_0}.$$

As in Lemma 3.4, $B_i[u, v] = 0$ outside R_i, the corresponding ordinary matrix B_i has rank at most d_i, and

$$\sum_i B_i[u, v] = P_u.$$

Let $A_i = A|_{R_i}$ and let J_{λ,d_0} be the all-ones $\lambda \times d_0$ matrix. Weight the blocks by A:

$$C_i[u, v] = A[u, v] B_i[u, v].$$

Then

$$C_i = B_i \circ (A_i \otimes J_{\lambda,d_0}),$$

so the Hadamard-product rank inequality gives

$$\operatorname{rank} C_i \le d_i \operatorname{rank} A_i.$$

The sum $D = \sum_i C_i$ has blocks

$$D[u,v] = A[u,v]P_u.$$

We claim

$$\operatorname{rank} D \ge \lambda \operatorname{rank} A.$$

Let $r = \operatorname{rank} A$ and choose r linearly independent rows $a_{u_1}, \dots, a_{u_r}$ of A. The row space contributed by the block row indexed by u_j is

$$a_{u_j} \otimes \operatorname{rowspan}(P_{u_j}),$$

which has dimension λ. These r tensor-product subspaces form a direct sum: if $\sum_j a_{u_j} \otimes z_j = 0$, linear independence of the a_{u_j} forces every $z_j = 0$. Hence the selected block rows already have rank λr, proving the claim.

Therefore

$$\lambda \operatorname{rank} A \le \operatorname{rank} D \le \sum_i \operatorname{rank} C_i \le \sum_i d_i \operatorname{rank}(A|_{R_i}).$$

Bounding the restricted ranks by their maximum and dividing by λ gives $d/\lambda \ge \mu_{\mathbb{F}}(f)$ after maximizing over A. □

We also need the following elementary consequence of the definition in Section 5. For a monotone Boolean function f, the dual function is

$$f^*(x) = 1 - f(\mathbf{1} - x).$$

Complementing an input exchanges positive and negative inputs and preserves the participant-separation relation after transposition.

Lemma 3.6 (Dual invariance of the rank measure)**.** *For every access structure Γ on P and every field $\mathbb{F}$,*

$$\mu_{\mathbb{F}}(\Gamma^*) = \mu_{\mathbb{F}}(\Gamma).$$

Proof. Let $U = \Gamma$ and $V = 2^P \setminus \Gamma$. The positive sets of Γ^* are $P \setminus V$ and its negative sets are $P \setminus U$. Given a witness $A \in \mathbb{F}^{U \times V}$, define the witness A^* for Γ^* by

$$A^*[P \setminus v,\ P \setminus u] = A[u,v].$$

Thus A^* is obtained from A^T by relabeling rows and columns, so $\operatorname{rank} A^* = \operatorname{rank} A$. Moreover,

$$i \in (P \setminus v) \setminus (P \setminus u) \quad \Longleftrightarrow \quad i \in u \setminus v,$$

so for every participant i, the restricted matrix of A^* is likewise a relabeled transpose of $A|_{R_i}$. Hence all participant-restricted ranks are preserved. Taking maxima gives $\mu_{\mathbb{F}}(\Gamma^*) \ge \mu_{\mathbb{F}}(\Gamma)$, and applying the same argument to $(\Gamma^*)^* = \Gamma$ gives equality. □

4 Exponential lower bounds for multilinear secret sharing

The exponential rank witnesses used in this section come from algebraic proofs that certain Boolean formulas are unsatisfiable. Two algebraic proof systems are involved: Polynomial Calculus and Nullstellensatz. For the formulas used by Pitassi and Robere, a linear lower bound on the degree of such algebraic refutations can be converted, by their lifting theorem, into an exponential lower bound for the Razborov–Gál rank measure. The construction is characteristic-sensitive, so we use two formula families, one suited to odd characteristic and one to characteristic two, and then combine them into a single access-structure family valid over every finite field. The section proceeds in that order.

Before invoking the proof-complexity results, we record two elementary monotonicity properties of the rank measure. The first passes from a partially defined monotone function to a total monotone extension. The second says that fixing inputs to zero cannot increase the rank measure. Both will be used below.

Lemma 4.1 (Rank-measure monotonicity)**.** *(i) Let $h : \{0,1\}^N \to \{0,1,*\}$ be a partially defined monotone function and let f be a total monotone extension of h. Then*

$$\mu_{\mathbb{F}}(f) \geq \mu_{\mathbb{F}}(h),$$

where the rank measure of h is defined using its specified positive and negative inputs.

(ii) Let f be a monotone function on $[n]$, let $D \subseteq [n]$, and let f' be the nonconstant restriction obtained by fixing every input in D to zero. Then

$$\mu_{\mathbb{F}}(f) \geq \mu_{\mathbb{F}}(f').$$

Proof. For (i), let $U_0 = h^{-1}(1)$ and $V_0 = h^{-1}(0)$, and let Z be a rank witness on $U_0 \times V_0$. Since f extends h, its positive and negative input sets contain U_0 and V_0, respectively. Extend Z to the full positive–negative matrix of f by putting zero in every new row and column. This preserves $\operatorname{rank} Z$, and for every input coordinate i it also preserves the rank after masking by the rectangle $u_i = 1, v_i = 0$. Thus the same matrix witnesses the same rank-measure ratio for f.

For (ii), identify an input x of f' with the input $\iota(x)$ of f obtained by inserting zeros in the coordinates of D. Let U', V' be the positive and negative inputs of f', and let A' be a rank witness on $U' \times V'$. Define A on the positive–negative matrix of f by

$$A[\iota(x), \iota(y)] = A'[x, y] \qquad (x \in U',\ y \in V'),$$

and put zero in every other entry. Then $\operatorname{rank} A = \operatorname{rank} A'$. For $i \notin D$, the mask $A|_{R_i}$ is an embedded copy of $A'|_{R'_i}$ and has the same rank. For $i \in D$, every row index supporting a nonzero entry of A has $u_i = 0$, so $A|_{R_i} = 0$. Hence the denominator in the rank-measure ratio is unchanged, and every witness for f' yields the same ratio for f. Taking the maximum over A' gives $\mu_{\mathbb{F}}(f) \geq \mu_{\mathbb{F}}(f')$. □

4.1 Two algebraic proof-complexity inputs

Recall that a CNF formula is a conjunction of clauses, each clause being a disjunction of literals. Nullstellensatz and Polynomial Calculus are algebraic proof systems that certify that such a formula has no Boolean satisfying assignment; their degree measures the minimum polynomial degree needed for the certificate or derivation. Appendix A recalls the formal definitions and polynomial encoding.

For an unsatisfiable CNF formula C and a field $\mathbb{F}$, write $\mathrm{NS}_{\mathbb{F}}(C)$ for the minimum degree of a Nullstellensatz refutation of C over $\mathbb{F}$, and $\mathrm{PC}_{\mathbb{F}}(C)$ for the minimum degree of a Polynomial Calculus refutation.

Proposition 4.2 (Uniform proof-complexity input; Buss–Grigoriev–Impagliazzo–Pitassi)**.** *For every fixed prime p there is an explicit family of unsatisfiable constant-width CNF formulas $\{C_{p,t}\}$ with $N_{p,t} = \Theta(t)$ variables and $O(N_{p,t})$ clauses, and a constant $\varepsilon_p > 0$, such that for every field $\mathbb{F}$ with $\operatorname{char}\mathbb{F} \neq p$,*

$$\mathrm{PC}_{\mathbb{F}}(C_{p,t}) \geq \varepsilon_p N_{p,t}.$$

Consequently, by the standard simulation $\mathrm{PC} \leq \mathrm{NS}$*,*

$$\mathrm{NS}_{\mathbb{F}}(C_{p,t}) \geq \varepsilon_p N_{p,t}$$

for the same fields.

For each fixed prime p, we use the mod-p counting-principle formulas of Buss, Grigoriev, Impagliazzo, and Pitassi. The underlying linear Polynomial Calculus degree lower bound is stated as Theorem 2 in their preliminary version [13, Theorem 2 in the preliminary version]; the constant-width CNF formulation used here is recorded by Pitassi and Robere [26, Theorem 7.3]. Alekhnovich and Razborov obtain a closely related uniform lower bound for Boolean mod-p Tseitin formulas [1].

A lifting gadget is a fixed small Boolean function used to replace the inputs of the original formula so that algebraic degree hardness becomes rank hardness for the resulting monotone function. We write g_λ for the Pitassi–Robere gadget of their Definition 4.1. Let J denote the all-ones matrix with the same dimensions as g. Pitassi and Robere call a gadget g *good* over $\mathbb{F}$ if, for every pair of same-size matrices A, B over $\mathbb{F}$,

$$\operatorname{rank}(J \otimes A + g \otimes B) = \operatorname{rank}(A) + \operatorname{rank}(g)\operatorname{rank}(B).$$

This exact rank-additivity property prevents cancellation when polynomial degree is converted into matrix rank.

Proposition 4.3 (Rank lifting; Pitassi–Robere)**.** *Let C be an unsatisfiable constant-width CNF on N variables with $O(N)$ clauses, let $\mathbb{F}$ be a field, and suppose*

$$\mathrm{NS}_{\mathbb{F}}(C) \geq \varepsilon N$$

for some $\varepsilon > 0$. Choose an integer $\lambda > 2^{1/\varepsilon}$ and let g_λ denote the Pitassi–Robere gadget. Then g_λ is good and

$$\operatorname{rank}_{\mathbb{F}}(g_\lambda) = \lambda$$

for every field $\mathbb{F}$, and the associated partially defined monotone function f_{C,g_λ} on $O(N)$ Boolean inputs satisfies

$$\mu_{\mathbb{F}}(f_{C,g_\lambda}) \geq 2^{\Omega(N)}.$$

The implicit constant depends only on ε, the clause width, and λ.

Pitassi and Robere's lifting theorem turns a linear Nullstellensatz-degree lower bound into an exponential lower bound for the rank measure, provided the gadget is good and has sufficiently large rank. Their gadget g_λ has rank exactly λ over every field and satisfies the required goodness condition [26, Theorems 3.4–3.5 and Lemma 4.2]. Thus, once the proof-complexity lower bound is fixed, choosing λ large enough gives the desired exponential rank lower bound uniformly over the relevant fields. A later theorem of de Rezende, Meir, Nordström, Pitassi, Robere, and Vinyals extends this lifting principle to a much broader class of high-rank gadgets [15].

4.2 The characteristic-separated families

To obtain a lower bound that is uniform over every finite field, we use two characteristic-separated families. The choice $p = 2$ gives the required lower bound over every field of odd characteristic, while $p = 3$ gives it over fields of characteristic two. In each case the proof-complexity lower bound and the lifting gadget can be chosen uniformly over the relevant fields. Proposition 4.2 provides the former, and Pitassi–Robere Lemma 4.2 provides the latter.

Theorem 4.4 (Characteristic-separated rank witnesses; Pitassi–Robere, restated)**.** *There are explicit families of total monotone functions $\{f_{2,n}\}$ and $\{f_{3,n}\}$, constants $\alpha_2, \alpha_3 > 0$, and thresholds n_2, n_3 with the following properties:*

(i) for every finite field $\mathbb{F}$ of odd characteristic and every $n \geq n_2$,

$$\mu_{\mathbb{F}}(f_{2,n}) \geq 2^{\alpha_2 n};$$

(ii) for every finite field $\mathbb{F}$ of characteristic two and every $n \geq n_3$,

$$\mu_{\mathbb{F}}(f_{3,n}) \geq 2^{\alpha_3 n}.$$

The constants and thresholds are independent of the particular finite field.

Proof. Fix $p \in \{2, 3\}$. Let $\{C_{p,t}\}$ and $\varepsilon_p > 0$ be as in Proposition 4.2; write $N_{p,t}$ for the number of variables. Fix once and for all an integer

$$\lambda_p > 2^{1/\varepsilon_p}.$$

By Pitassi–Robere Lemma 4.2, the same gadget g_{λ_p} is good and has rank λ_p over every field. Hence Proposition 4.3 gives, for every field $\mathbb{F}$ with $\operatorname{char} \mathbb{F} \neq p$, a partially defined monotone function $\varphi_{p,t}$ on $L_{p,t} = O(N_{p,t}) = O(t)$ inputs satisfying

$$\mu_{\mathbb{F}}(\varphi_{p,t}) \geq 2^{c_p N_{p,t}} = 2^{\Omega(t)},$$

where $c_p > 0$ is independent of $\mathbb{F}$.

Choose the canonical monotone total extension

$$\widetilde{f}_{p,t}(x) = 1 \quad \Longleftrightarrow \quad \text{some specified positive input of } \varphi_{p,t} \text{ is coordinatewise at most } x.$$

Because $\varphi_{p,t}$ is partially monotone, this agrees with all specified negative inputs. By Lemma 4.1(i),

$$\mu_{\mathbb{F}}(\widetilde{f}_{p,t}) \geq \mu_{\mathbb{F}}(\varphi_{p,t}) \geq 2^{\Omega(t)}.$$

Let $C_p > 0$ satisfy $L_{p,t} \leq C_p t$ for all t. For each sufficiently large desired input length n, choose $t = \lfloor n/C_p \rfloor$ and let $f_{p,n}$ be the function obtained by padding $\widetilde{f}_{p,t}$ with ignored inputs until it has exactly n inputs. Restricting the dummy coordinates to zero recovers $\widetilde{f}_{p,t}$, so Lemma 4.1(ii) gives

$$\mu_{\mathbb{F}}(f_{p,n}) \geq \mu_{\mathbb{F}}(\widetilde{f}_{p,t}).$$

Since $t = \Omega(n)$, there are constants $\alpha_p > 0$ and n_p such that

$$\mu_{\mathbb{F}}(f_{p,n}) \geq 2^{\alpha_p n}$$

for every $n \geq n_p$ in the corresponding characteristic range.

Taking $p = 2$ gives (i), since every odd characteristic differs from 2. Taking $p = 3$ gives (ii), since characteristic two differs from 3. □

Corollary 4.5. *Under the respective hypotheses of Theorem 4.4,*

$$\mathrm{amSP}_{\mathbb{F}}(f_{p,n}) \geq 2^{\alpha_p n} \qquad (p \in \{2, 3\}).$$

Proof. Apply Theorems 3.1 and 4.4. □

4.3 One family over every finite field

To obtain one family that works over all finite fields, let $f_{2,k}$ and $f_{3,k}$ use disjoint sets of k variables and define

$$F_{2k}(x, y) = f_{2,k}(x) \vee f_{3,k}(y).$$

Each $f_{p,k}$ is nonconstant and monotone, and hence $f_{p,k}(0^k) = 0$. We apply the zero-restriction part of Lemma 4.1 directly to the rank measure; the amortized span-program bound will then follow from Theorem 3.1.

Theorem 4.6 (One family over every finite field)**.** *There is an explicit monotone family $\{F_n\}$, a constant $\alpha > 0$, and a threshold n_0 such that, for every finite field $\mathbb{F}$ and every $n \geq n_0$,*

$$\mu_{\mathbb{F}}(F_n) \geq 2^{\alpha n}.$$

Consequently,

$$\mathrm{amSP}_{\mathbb{F}}(F_n) \geq 2^{\alpha n}.$$

Proof. First let $n = 2k$. If $\operatorname{char}\mathbb{F} = 2$, restricting the first block of F_{2k} to zero gives $f_{3,k}$; hence Lemma 4.1(ii) and Theorem 4.4(ii) give

$$\mu_{\mathbb{F}}(F_{2k}) \geq \mu_{\mathbb{F}}(f_{3,k}) \geq 2^{\alpha_3 k}.$$

If $\operatorname{char}\mathbb{F} \neq 2$, restricting the second block to zero gives $f_{2,k}$, and similarly

$$\mu_{\mathbb{F}}(F_{2k}) \geq \mu_{\mathbb{F}}(f_{2,k}) \geq 2^{\alpha_2 k}.$$

Thus

$$\mu_{\mathbb{F}}(F_{2k}) \geq 2^{\min(\alpha_2,\alpha_3)k}.$$

For odd $n = 2k + 1$, let F_n ignore its last input. Restricting that input to zero recovers F_{2k}, so Lemma 4.1(ii) gives the same lower bound. Taking

$$\alpha = \frac{1}{4}\min\{\alpha_2, \alpha_3\} \qquad \text{and} \qquad n_0 = 2\max\{n_2, n_3\} + 1$$

gives constants valid simultaneously for every finite field. The amortized span-program bound follows from Theorem 3.1. □

Theorem 4.7 (Exponential multilinear information ratio)**.** *Let Γ_n be the access structure represented by F_n in Theorem 4.6. With the constants $\alpha > 0$ and n_0 from that theorem, for every finite field $\mathbb{F}$ and every $n \geq n_0$, every perfect multilinear scheme realizing Γ_n satisfies*

$$\rho_{\mathrm{av}} \geq \frac{2^{\alpha n}}{n} \qquad \textit{and} \qquad \rho_{\max} \geq \frac{2^{\alpha n}}{n}.$$

Proof. By Lemma 2.5, the scheme gives a d_0-target MTMSP in which participant i labels exactly d_i rows. By Theorem 4.6,

$$\frac{d}{d_0} = \sum_{i=1}^{n} \frac{d_i}{d_0} \geq 2^{\alpha n}.$$

Therefore

$$\rho_{\mathrm{av}} = \frac{d}{n d_0} \geq \frac{2^{\alpha n}}{n} = 2^{\Omega(n)}.$$

Since $\rho_{\max} \geq \rho_{\mathrm{av}}$, the same lower bound holds for the maximum ratio. □

Scalar attribution. The all-finite-fields exponential lower bound for scalar linear schemes is already recorded in BOP [10, extended version, Theorem 4.10], with attribution there to Pitassi–Robere and Robere's thesis [29]. The new point here is Theorem 3.1: the same rank obstruction persists for arbitrary target dimension.

Corollary 4.8 (Worst-case asymptotic order)**.** *For an access structure Γ, let $\rho_{\max}^{\mathrm{ML}}(\Gamma)$ be the infimum of $\rho_{\max}$ over all perfect multilinear schemes for Γ, over all finite fields and all secret dimensions. Define*

$$R_{\mathrm{multi}}(n) = \max_{\Gamma \text{ on } n \text{ participants}} \rho_{\max}^{\mathrm{ML}}(\Gamma).$$

Then

$$\frac{2^{\alpha n}}{n} \le R_{\mathrm{multi}}(n) \le \binom{n}{\lfloor n/2 \rfloor} \le 2^n,$$

so $R_{\mathrm{multi}}(n) = 2^{\Theta(n)}$.

Proof. The lower bound is Theorem 4.7. For the upper bound, use the standard minimal-authorized-set construction underlying the classical general-access-structure schemes of Ito–Saito–Nishizeki and Benaloh–Leichter [19, 11]. For each $A \in \min \Gamma$, share one copy of the same scalar secret by an additive $|A|$-out-of-$|A|$ linear scheme among the participants of A. Participant i receives one field element for each minimal authorized set containing i. Every authorized set contains a minimal authorized set and can reconstruct from its corresponding copy, while an unauthorized set misses at least one additive share in every copy and therefore learns nothing. Hence

$$\rho_{\max}^{\mathrm{ML}}(\Gamma) \le \max_i |\{A \in \min \Gamma : i \in A\}| \le |\min \Gamma|.$$

Since $\min \Gamma$ is an antichain, Sperner's theorem gives

$$|\min \Gamma| \le \binom{n}{\lfloor n/2 \rfloor}.$$

□

5 Affine-reconstructable schemes

Recall the two scheme-level quantities

$$\Phi_\Pi(A) = \frac{I(X_0; X_A)}{H(X_0)} \qquad \text{and} \qquad \mathrm{cv}(\Pi) = (\rho_1(\Pi), \ldots, \rho_n(\Pi)).$$

The first measures the normalized secret information available to a coalition, while the second is the contribution vector of the scheme. Affine reconstruction allows us to replace arbitrary sharing by multilinear sharing while controlling both quantities. Before giving the construction, Table 1 places this transformation beside the two earlier results used later in the paper.

Result	Hypotheses on the original scheme Π	Guarantee for the resulting scheme $\widehat{\Pi}$
BOP [10, Thm. 10.9]	arbitrary sharing; linear reconstruction; perfect correctness; stated with pairwise parameter $\delta \le 1/2$	perfect multilinear; same secret space; $\rho_{\mathrm{tot}}(\widehat{\Pi}) = \rho_{\mathrm{tot}}(\Pi)$
JK [21, Prop. 5.5]	linear scheme; perfect correctness; partial privacy with advantage η	perfect linear scheme over a finite extension; $\mathrm{cv}(\widehat{\Pi}) = \mathrm{cv}(\Pi)/\eta$
Theorem 5.1	affine-reconstructable; perfect correctness	perfectly correct multilinear scheme with the same secret space, $\mathrm{cv}(\widehat{\Pi}) \le \mathrm{cv}(\Pi)$, and $\Phi_{\widehat{\Pi}} \le \Phi_{\Pi}$
Theorem 5.6	affine-reconstructable; perfect correctness; pairwise parameter $\delta < 1$	perfect multilinear scheme with the same secret space and $\mathrm{cv}(\widehat{\Pi}) \le \mathrm{cv}(\Pi)$

Table 1: Related transformations, stated in terms of the hypotheses on the original scheme and the guarantee for the resulting scheme.

5.1 Affine linearization

Affine reconstruction alone is enough to replace the sharing algorithm by a multilinear one without increasing any participant dimension or the access function of any coalition. The theorem below assumes no privacy condition; partial privacy and pairwise statistical privacy enter only afterward. Its construction is closely related to the reconstruction-dual argument of BOP [10, extended version, Construction 10.6], while the pointwise access-function comparison is the feature needed for the partial-privacy branch.

Theorem 5.1 (Affine linearization)**.** *Let Π be an affine-reconstructable scheme over $\mathbb{F}$ that is perfectly correct for Γ, with secret space $\mathbb{F}^{d_0}$ and participant-i share space $\mathbb{F}^{d_i}$. Then there is a multilinear scheme $\widehat{\Pi}$ over $\mathbb{F}$, perfectly correct for Γ, with the same secret space $\mathbb{F}^{d_0}$ such that*

$$\mathrm{cv}(\widehat{\Pi}) \le \mathrm{cv}(\Pi)$$

componentwise and

$$\Phi_{\widehat{\Pi}}(A) \le \Phi_{\Pi}(A) \qquad (A \subseteq P).$$

In particular, participant i uses at most d_i field coordinates in $\widehat{\Pi}$.

Proof. Let Ω be the set of all tuples

$$\omega = (s, x_1, \dots, x_n)$$

that can be produced by the sharing algorithm. Work in

$$\mathcal{V} = \mathbb{F}^{\Omega}/\langle \mathbf{1} \rangle,$$

the vector space of $\mathbb{F}$-valued functions on Ω modulo constant functions. Let

$$\bar{s}_j = [\omega \mapsto s_j(\omega)] \quad (j \in [d_0]),$$

and

$$\bar{x}_{i,t} = [\omega \mapsto x_{i,t}(\omega)] \quad (i \in [n],\ t \in [d_i]).$$

Put

$$S = \mathrm{span}\{\bar{s}_1, \dots, \bar{s}_{d_0}\}, \qquad U_i = \mathrm{span}\{\bar{x}_{i,t} : t \in [d_i]\}, \qquad U_A = \sum_{i \in A} U_i.$$

Because the secret ranges over all of $\mathbb{F}^{d_0}$, the classes $\bar{s}_1, \ldots, \bar{s}_{d_0}$ are linearly independent, so

$$\dim S = d_0, \qquad \dim U_i \le d_i.$$

Exact affine reconstruction implies $S \subseteq U_A$ for every authorized A: modulo constants, each affine reconstruction identity becomes a linear relation expressing the corresponding $\bar{s}_j$ in terms of the share-coordinate classes available to A.

Choose bases of S and the U_i, and put

$$W = S + \sum_{i=1}^{n} U_i \subseteq \mathcal{V}.$$

Only the restriction of a functional to W can affect the secret and shares, so let Z be a uniformly random linear functional on W. Define $\widehat{X}_0$ to be the coordinates of $Z|_S$ and $\widehat{X}_i$ the coordinates of $Z|_{U_i}$. After choosing a linear section of the restriction map $W^* \to S^*$, this joint random variable is generated by a fixed linear map of $\widehat{X}_0$ and independent uniform randomness; hence it defines a multilinear scheme $\widehat{\Pi}$. Since $S \subseteq U_A$ for authorized A, coalition A determines $Z|_S$, so $\widehat{\Pi}$ is perfectly correct for Γ. Also participant i receives $\dim U_i \le d_i$ field elements, proving the convec inequality.

It remains to compare access functions. For a coalition A, let $\mathcal{L}_A$ be the vector space of linear forms $\beta \in (\mathbb{F}^{d_0})^*$ for which $\beta \cdot s$ is an affine-linear function of A's shares on every possible sharing. The map

$$\beta \longmapsto [\beta \cdot s]$$

identifies $\mathcal{L}_A$ with $S \cap U_A$: one direction follows by reducing an affine recovery identity modulo constants, and the converse follows because membership in U_A says that the difference from a linear combination of A's share coordinates is constant on Ω.

For the multilinear scheme $\widehat{\Pi}$, the standard dimension formula for linear random variables gives

$$I(\widehat{X}_0; \widehat{X}_A) = \dim(S \cap U_A) \log |\mathbb{F}|,$$

and hence

$$\Phi_{\widehat{\Pi}}(A) = \frac{\dim \mathcal{L}_A}{d_0}.$$

If $\ell = \dim \mathcal{L}_A$, choose ℓ independent forms in $\mathcal{L}_A$. Their joint value on the uniform secret $X_0 \in \mathbb{F}^{d_0}$ has entropy $\ell \log |\mathbb{F}|$ and is a deterministic function of X_A. Therefore

$$I(X_0; X_A) \ge \ell \log |\mathbb{F}|,$$

which proves $\Phi_{\widehat{\Pi}}(A) \le \Phi_{\Pi}(A)$. □

Recall from Section 2.1 that partial privacy with advantage $\eta > 0$ means, in addition to perfect correctness,

$$\Phi_{\Pi}(B) \le 1 - \eta \qquad (B \notin \Gamma).$$

Corollary 5.2 (Partial privacy is preserved)**.** *If Π in Theorem 5.1 is partially private for Γ with advantage $\eta > 0$, then the multilinear scheme $\widehat{\Pi}$ constructed there is partially private for Γ with advantage at least η.*

Proof. For every unauthorized B,

$$\Phi_{\widehat{\Pi}}(B) \le \Phi_{\Pi}(B) \le 1 - \eta.$$

Authorized sets are perfectly correct in $\widehat{\Pi}$ by Theorem 5.1. □

5.2 Partial targets and direct lower-bound transfer

Partial privacy does not in general make the multilinear scheme obtained above perfect. Nevertheless, it guarantees that every unauthorized coalition leaves a positive fraction of the secret unrecovered. We now show that this remaining fraction is sufficient for the partial-target rank bound, yielding a lower bound proportional to the privacy parameter. The proof of the following theorem is given in Appendix B.

Theorem 5.3 (Partial dual linearization)**.** *Let Π be an affine-reconstructable scheme over $\mathbb{F}$, perfectly correct for Γ, with uniform secret $\mathbb{F}^{d_0}$ and participant-i share space $\mathbb{F}^{d_i}$. If Π is partially private with advantage $\eta > 0$, then there is a labeled matrix for f_{Γ^*} with a d_0-dimensional target such that every negative input strongly rejects the target, every positive input meets the target in dimension at least some integer λ, participant i labels exactly d_i rows, and*

$$\lambda \geq \eta d_0.$$

Proof. See Appendix B. □

Corollary 5.4 (Direct rank transfer under partial privacy)**.** *Let Π satisfy the hypotheses of Theorem 5.3, let $d = \sum_i d_i$, and let λ be as in that theorem. Then*

$$\frac{d}{\lambda} \geq \mu_{\mathbb{F}}(\Gamma).$$

Consequently,

$$\rho_{\mathrm{tot}}(\Pi) = \frac{d}{d_0} \geq \mu_{\mathbb{F}}(\Gamma) \max\left\{\eta, \frac{1}{d_0}\right\}.$$

The corresponding family-wise bounds for average and maximum ratios follow from

$$\rho_{\mathrm{av}} = \frac{1}{n}\rho_{\mathrm{tot}}, \qquad \rho_{\max} \geq \rho_{\mathrm{av}}.$$

Proof. Apply Theorem 3.5 to the program of Theorem 5.3 on Γ^* and use dual invariance, $\mu_{\mathbb{F}}(\Gamma^*) = \mu_{\mathbb{F}}(\Gamma)$, from Lemma 3.6. Thus $d/\lambda \geq \mu_{\mathbb{F}}(\Gamma)$. Since $\lambda/d_0 \geq \eta$ and $\lambda \geq 1$,

$$\frac{d}{d_0} = \frac{d}{\lambda}\frac{\lambda}{d_0} \geq \mu_{\mathbb{F}}(\Gamma) \max\left\{\eta, \frac{1}{d_0}\right\}.$$

□

The inequality $d/\lambda \geq \mu_{\mathbb{F}}(\Gamma)$ can alternatively be obtained by applying Jafari–Khazaei's partial-to-perfect transformation [21, Proposition 5.4] to the partial linear scheme of Theorem 5.3 and then applying Theorem 3.1. The direct argument above avoids that amplification and keeps the partial-target rank mechanism explicit.

Corollary 5.5 (Exponential lower bounds under partial privacy)**.** *For the explicit family $\{\Gamma_n\}$ of Theorem 4.7, every affine-reconstructable partially private scheme over a finite field with uniform secret $\mathbb{F}^{d_0}$ and advantage $\eta > 0$ satisfies*

$$\rho_{\mathrm{tot}} \geq 2^{\alpha n} \max\left\{\eta, \frac{1}{d_0}\right\},$$

and

$$\rho_{\mathrm{av}}, \rho_{\max} \geq \frac{2^{\alpha n}}{n} \max\left\{\eta, \frac{1}{d_0}\right\}.$$

In particular, all three ratios remain $2^{\Omega(n)}$ whenever either $\eta = 2^{-o(n)}$ or $d_0 = 2^{o(n)}$.

Proof. Use Corollary 5.4 and $\mu_{\mathbb{F}}(\Gamma_n) \geq 2^{\alpha n}$ from Theorem 4.6. The average and maximum bounds follow from $\rho_{\text{av}} = \rho_{\text{tot}}/n$ and $\rho_{\max} \geq \rho_{\text{av}}$. □

5.3 Pairwise privacy below one

Recall from Section 2.1 that pairwise statistical privacy with parameter δ bounds the statistical distance between the views of every unauthorized coalition on any two secrets. The condition $\delta < 1$ is strictly weaker than perfect privacy; Appendix C gives a 2-out-of-2 example. For an affine-reconstructable scheme, however, this strict inequality is enough to make the multilinear scheme from Theorem 5.1 perfect.

Theorem 5.6 (Perfect replacement below pairwise distance one)**.** *Let Π be an affine-reconstructable scheme over $\mathbb{F}$, perfectly correct for Γ, with secret space $\mathbb{F}^{d_0}$ and pairwise statistical privacy parameter $\delta < 1$. Then the multilinear scheme $\widehat{\Pi}$ from Theorem 5.1 is perfect for Γ. It has the same secret space $\mathbb{F}^{d_0}$ and*

$$\text{cv}(\widehat{\Pi}) \leq \text{cv}(\Pi)$$

componentwise. By padding with deterministic zero coordinates, the original field-coordinate convec can be matched exactly if desired.

Proof. Suppose some unauthorized B had $\Phi_{\widehat{\Pi}}(B) > 0$. By the construction in the proof of Theorem 5.1, this means that B can recover affinely a nonconstant linear form $\beta \cdot X_0$ of the secret. Choose secrets $s, s' \in \mathbb{F}^{d_0}$ with $\beta \cdot s \neq \beta \cdot s'$. The supports of $\text{Share}_B(s)$ and $\text{Share}_B(s')$ lie in two disjoint affine level sets of the recovery function, so their statistical distance is one, contradicting $\delta < 1$. Hence

$$\Phi_{\widehat{\Pi}}(B) = 0 \qquad (B \notin \Gamma).$$

Since $\widehat{\Pi}$ is perfectly correct for every authorized set, it is a perfect multilinear scheme for Γ. □

The support-disjointness step in this proof is the same mechanism used in BOP's Claim 10.8 [10, extended version, Claim 10.8]; here it is combined with the affine-linearization theorem above in the form needed for the subsequent rank transfer.

The conclusion is existential: the original scheme need be neither multilinear nor perfect; Appendix C gives an explicit example with both features. What Theorem 5.6 guarantees is a perfect multilinear replacement with the same secret space and participant dimensions no larger than those of the original scheme.

Corollary 5.7 (Rank lower bound for affine-reconstructable schemes)**.** *Under the hypotheses of Theorem 5.6, if the shares contain $d = \sum_i d_i$ field elements in total, then*

$$\rho_{\text{tot}} = \frac{d}{d_0} \geq \mu_{\mathbb{F}}(\Gamma).$$

Proof. Apply Theorem 5.6, then Lemma 2.5 and Theorem 3.1. □

Corollary 5.8 (Exponential lower bound with pairwise statistical privacy)**.** *For the explicit family $\{\Gamma_n\}$ of Theorem 4.7, over every finite field every affine-reconstructable scheme satisfying the pairwise hypothesis of Theorem 5.6 has*

$$\rho_{\text{tot}} \geq 2^{\Omega(n)}, \qquad \rho_{\text{av}} \geq 2^{\Omega(n)}, \qquad \rho_{\max} \geq 2^{\Omega(n)}.$$

Proof. Apply Theorem 5.6 and then Theorems 4.6 and 4.7. □

6 Degree-bounded reconstruction

Affine reconstruction is the degree-one case of polynomial reconstruction. BOP's degree-reduction theorem moves the nonlinear monomials from reconstruction into the sharing procedure, with polynomial overhead in the original total share dimension. This allows the lower bounds of Section 5 to extend to every fixed reconstruction degree. The pairwise-statistical and partial-privacy assumptions lead to parallel consequences.

6.1 Degree reduction

Here degree refers to polynomials in the share coordinates available to the reconstructing set. With d original share coordinates, the number of monomials of total degree at most r is at most

$$\binom{d+r}{r} = O_r(d^r).$$

BOP share the values of the nonlinear monomials during the sharing procedure so that reconstruction becomes linear. For an unauthorized set, the extra values it receives are either independent masks or functions of its original view; consequently the transformed view contains no more information about the secret than the original one [10, Lemma 4.7, extended version].

Lemma 6.1 (Degree reduction; Beimel, Othman, and Peter, extended version, Lemma 4.7)**.** *Let* Π *be a degree-*r *reconstructable scheme with a* d_0*-dimensional secret, and let* d *be the total number of field elements in all shares. Then* Π *can be transformed into a multilinear scheme with the same secret dimension and the same correctness and pairwise statistical-privacy parameters, in which each participant receives at most* $K_r d^r$ *field elements, where* K_r *depends only on* r*. Moreover, for every unauthorized coalition, the transformed view can be generated from its original view by secret-independent randomized postprocessing.*

The last property is the form needed for partial privacy: by data processing, the mutual information between an unauthorized transformed view and the secret cannot increase.

6.2 Pairwise statistical privacy

Under pairwise statistical distance strictly below one, the multilinear scheme produced by the degree reduction can be perfected with no further coordinate overhead by Theorem 5.6.

Theorem 6.2 (Degree-r reconstruction under pairwise statistical privacy)**.** *There are constants* $\alpha > 0$ *and* n_0*, and for every fixed* $r \geq 2$ *a constant* $K_r > 0$*, such that the explicit family* $\{\Gamma_n\}$ *of Theorem 4.7 has the following property. Let* $n \geq n_0$ *and let* Π *be a perfectly correct degree-*r *reconstructable scheme over a finite field, with secret space* $\mathbb{F}^{d_0}$*, pairwise statistical privacy parameter* $\delta < 1$*, and* d *total share field elements. Then*

$$\rho_{\mathrm{tot}}(\Pi) = \frac{d}{d_0} \geq K_r^{-1/r} n^{-1/r} 2^{\alpha n/r} d_0^{-(1-1/r)}.$$

Consequently, for every fixed r*, if* $d_0 = 2^{o(n)}$*, then*

$$\rho_{\mathrm{tot}}(\Pi) \geq 2^{\Omega(n)}, \qquad \rho_{\mathrm{av}}(\Pi) \geq 2^{\Omega(n)}, \qquad \rho_{\max}(\Pi) \geq 2^{\Omega(n)}.$$

Proof. Apply Lemma 6.1. The resulting multilinear scheme has the same secret dimension d_0, pairwise statistical privacy parameter below one, and participant dimensions at most $K_r d^r$. By Theorem 5.6, it has a perfect multilinear replacement with no larger participant dimensions. By Theorem 4.7, that replacement has maximum ratio at least $2^{\alpha n}/n$, so some participant receives at least

$$\frac{d_0 2^{\alpha n}}{n}$$

field elements. Therefore

$$K_r d^r \geq \frac{d_0 2^{\alpha n}}{n},$$

and hence

$$d \geq K_r^{-1/r} n^{-1/r} d_0^{1/r} 2^{\alpha n/r}.$$

Dividing by d_0 gives the displayed bound. If $d_0 = 2^{o(n)}$ and r is fixed, its logarithm is

$$\frac{\alpha}{r} n - o(n) - \frac{1}{r} \log_2 n - O_r(1) = \Omega(n).$$

The average ratio is ρ_{tot}/n and the maximum is at least the average; division by n is absorbed by the exponential lower bound. □

6.3 Partial privacy

The same degree reduction preserves JK partial privacy. Indeed, if Π' denotes the transformed scheme, then for every unauthorized coalition B the data-processing inequality gives

$$I(X_0; X'_B) \leq I(X_0; X_B),$$

and hence

$$\Phi_{\Pi'}(B) \leq \Phi_{\Pi}(B).$$

Thus the partial-privacy advantage cannot decrease. Combining this observation with Corollary 5.5 gives the corresponding fixed-degree lower bound.

Theorem 6.3 (Degree-r reconstruction under partial privacy)**.** *With the constants and family of Theorem 6.2, let Π be a perfectly correct degree-r reconstructable scheme over a finite field, with secret space $\mathbb{F}^{d_0}$, total share dimension d, and partial-privacy advantage $\eta > 0$. Then*

$$\rho_{\mathrm{tot}}(\Pi) \geq K_r^{-1/r} n^{-1/r} 2^{\alpha n/r} \max\left\{\eta^{1/r} d_0^{-(1-1/r)}, d_0^{-1}\right\}.$$

Consequently, for every fixed r, if $d_0 = 2^{o(n)}$, then

$$\rho_{\mathrm{tot}}(\Pi),\ \rho_{\mathrm{av}}(\Pi),\ \rho_{\max}(\Pi) = 2^{\Omega(n)}.$$

Proof. Apply Lemma 6.1. Let Π' be the resulting multilinear scheme. Its secret dimension is d_0, every participant receives at most $K_r d^r$ field elements, and the preceding data-processing argument gives

$$\eta(\Pi'; \Gamma_n) \geq \eta(\Pi; \Gamma_n) = \eta.$$

By Corollary 5.5,

$$\rho_{\max}(\Pi') \geq \frac{2^{\alpha n}}{n} \max\left\{\eta, \frac{1}{d_0}\right\}.$$

Hence some participant in Π' receives at least

$$\frac{2^{\alpha n}}{n} \max\{\eta d_0, 1\}$$

field elements. Comparing this with the upper bound $K_r d^r$ gives

$$K_r d^r \geq \frac{2^{\alpha n}}{n} \max\{\eta d_0, 1\}.$$

Taking rth roots and dividing by d_0 yields the stated bound. If $d_0 = 2^{o(n)}$, the second term inside the maximum is $2^{-o(n)}$, so the total ratio is exponential; the same follows for the average and maximum ratios. □

7 Discussion and open problems

The main theorem shows that every Razborov–Gál rank obstruction survives target amortization. Together with the Pitassi–Robere witnesses, this settles the worst-case multilinear information ratio at the scale $2^{\Theta(n)}$. Three questions remain.

The exponent. Our lower bound gives $2^{\alpha n}$ for some $\alpha > 0$. On the upper side, Nir's explicit linear construction [24] already gives

$$R_{\text{multi}}(n) \leq 2^{\lceil n/2 \rceil - 1} + 1 = 2^{n/2 + O(1)},$$

because scalar linear schemes are a special case of multilinear schemes. Determining the optimal multilinear exponent remains open.

Arbitrarily long secrets with degree-bounded reconstruction. For fixed $r > 1$, Theorem 6.2 loses the factor $d_0^{1-1/r}$ and therefore gives an exponential normalized lower bound for subexponential secret dimension, more generally whenever $\log_2 d_0$ stays below the threshold implicit in the theorem. The partial-privacy bound of Theorem 6.3 has the analogous dependence on d_0 and the privacy advantage. Removing these losses would settle the remaining long-secret case of the BOP question. A direct degree-r obstruction stable under amortization may be needed.

Leaky forbidden sets. The partial-target rank theorem weakens acceptance but still requires strong rejection: a negative set must have zero intersection with the target. Jafari–Khazaei provide an indirect route around nonzero forbidden intersections for linear partial schemes by perfecting the scheme first [21, Proposition 5.4]. The reconstruction dual of Theorem 5.3 instead avoids leaky negative targets by exploiting affine-reconstruction structure. It remains open whether the rank method itself can tolerate nonzero forbidden target intersections without first perfecting the scheme or using such reconstruction-specific structure.

Acknowledgements

The author thanks Mohammad Hossein Kalantari for many useful discussions. Generative AI tools were used for limited mathematical discussion and for language and editorial polishing. The author is solely responsible for the mathematical content.

References

[1] M. V. Alekhnovich and A. A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003.

[2] B. Applebaum and B. Arkis. On the power of amortization in secret sharing: d-uniform secret sharing and CDS with constant information rate. *ACM Transactions on Computation Theory*, 12(4), Article 24, 2020. Preliminary version in TCC 2018.

[3] B. Applebaum, A. Beimel, O. Farràs, O. Nir, and N. Peter. Secret-sharing schemes for general and uniform access structures. In *Advances in Cryptology—EUROCRYPT 2019*, LNCS 11478, pages 441–471, 2019.

[4] B. Applebaum, A. Beimel, O. Nir, and N. Peter. Better secret-sharing via robust conditional disclosure of secrets. In *Proceedings of STOC 2020*, pages 280–293, 2020.

[5] B. Applebaum and O. Nir. Upslices, downslices, and secret-sharing with complexity of 1.5^n. In *Advances in Cryptology—CRYPTO 2021*, LNCS 12827, pages 627–655, 2021.

[6] A. Beimel. Secret-sharing schemes: A survey. In Y. Chee et al., editors, *Coding and Cryptology*, LNCS 6639, pages 11–46. Springer, 2011.

[7] A. Beimel. Secret-sharing schemes for general access structures: An introduction. IACR Cryptology ePrint Archive, Report 2025/518, Version 1.1, 15 August 2026. `https://eprint.iacr.org/2025/518`.

[8] A. Beimel, A. Ben-Efraim, C. Padró, and I. Tyomkin. Multi-linear secret-sharing schemes. In *Theory of Cryptography—TCC 2014*, LNCS 8349, pages 394–418, 2014. `https://doi.org/10.1007/978-3-642-54242-8_17`.

[9] A. Beimel, O. Farràs, and O. Lasri. Improved polynomial secret-sharing schemes. In *Theory of Cryptography—TCC 2023*, 2023. Extended version: IACR ePrint 2023/1158.

[10] A. Beimel, H. Othman, and N. Peter. Quadratic secret sharing and conditional disclosure of secrets. *IEEE Transactions on Information Theory*, 69(11):7295–7316, 2023. `https://doi.org/10.1109/TIT.2023.3296588`. Extended version: IACR ePrint 2021/285, revised 23 February 2022, `https://eprint.iacr.org/2021/285`.

[11] J. C. Benaloh and J. Leichter. Generalized secret sharing and monotone functions. In *Advances in Cryptology—CRYPTO '88*, LNCS 403, pages 27–35. Springer, 1990. `https://doi.org/10.1007/0-387-34799-2_3`.

[12] G. R. Blakley. Safeguarding cryptographic keys. In *Proceedings of the National Computer Conference*, volume 48, pages 313–317, 1979.

[13] S. R. Buss, D. Grigoriev, R. Impagliazzo, and T. Pitassi. Linear gaps between degrees for the polynomial calculus modulo distinct primes. *Journal of Computer and System Sciences*, 62(2):267–289, 2001. `https://doi.org/10.1006/jcss.2000.1726`.

[14] L. Csirmaz. The size of a share must be large. *Journal of Cryptology*, 10:223–231, 1997.

[15] S. F. de Rezende, O. Meir, J. Nordström, T. Pitassi, R. Robere, and M. Vinyals. Lifting with simple gadgets and applications to circuit and proof complexity. In *Proceedings of FOCS 2020*, pages 24–30, 2020. https://doi.org/10.1109/FOCS46700.2020.00011.

[16] O. Farràs, T. B. Hansen, T. Kaced, and C. Padró. On the information ratio of non-perfect secret sharing schemes. *Algorithmica*, 79(4):987–1013, 2017. https://doi.org/10.1007/s00453-016-0217-9.

[17] A. Gál. A characterization of span program size and improved lower bounds for monotone span programs. *Computational Complexity*, 10(4):277–296, 2001.

[18] A. Gál and P. Pudlák. A note on monotone complexity and the rank of matrices. *Information Processing Letters*, 87(6):321–326, 2003.

[19] M. Ito, A. Saito, and T. Nishizeki. Secret sharing scheme realizing general access structure. *Electronics and Communications in Japan (Part III)*, 72(9):56–64, 1989. https://doi.org/10.1002/ecjc.4430720906.

[20] A. Jafari and S. Khazaei. On Abelian and homomorphic secret sharing schemes. *Journal of Cryptology*, 34(4), Article 43, 2021. https://doi.org/10.1007/s00145-021-09410-2.

[21] A. Jafari and S. Khazaei. Partial secret sharing schemes. *IEEE Transactions on Information Theory*, 69(8):5364–5385, 2023. https://doi.org/10.1109/TIT.2023.3265093.

[22] M. Karchmer and A. Wigderson. On span programs. In *Proceedings of the 8th Annual Structure in Complexity Theory Conference*, pages 102–111, 1993.

[23] T. Liu and V. Vaikuntanathan. Breaking the circuit-size barrier in secret sharing. In *Proceedings of STOC 2018*, pages 699–708, 2018. https://doi.org/10.1145/3188745.3188936.

[24] O. Nir. Resolving the complexity of linear secret sharing. IACR Cryptology ePrint Archive, Report 2026/1667, 12 August 2026. https://eprint.iacr.org/2026/1667.

[25] A. Paskin-Cherniavsky and S. Radune. On polynomial secret sharing schemes. In *1st Conference on Information-Theoretic Cryptography (ITC 2020)*, LIPIcs 163, Article 12, 2020.

[26] T. Pitassi and R. Robere. Lifting Nullstellensatz to monotone span programs over any field. In *Proceedings of STOC 2018*, pages 1207–1219, 2018. https://doi.org/10.1145/3188745.3188914.

[27] A. A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.

[28] R. Robere, T. Pitassi, B. Rossman, and S. A. Cook. Exponential lower bounds for monotone span programs. In *Proceedings of FOCS 2016*, pages 406–415, 2016.

[29] R. Robere. *Unified Lower Bounds for Monotone Computation.* PhD thesis, University of Toronto, 2018. https://hdl.handle.net/1807/92007.

[30] A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[31] C. Williamson. Secret sharing at the Shannon ceiling. arXiv:2608.17047, 17 August 2026. https://arxiv.org/abs/2608.17047.

A Proof-complexity background for Section 4

This appendix recalls the algebraic proof-complexity notions needed in Section 4. Let C be a CNF formula on Boolean variables $x_1, \dots, x_N$, that is, variables taking values in $\{0,1\}$. We work over a field $\mathbb{F}$ together with the Boolean axioms

$$x_i^2 - x_i = 0 \qquad (i \in [N]).$$

For a literal ℓ, write $\ell(x) = x_i$ for the positive literal x_i and $\ell(x) = 1 - x_i$ for the negative literal $\neg x_i$. A clause $\ell_1 \vee \cdots \vee \ell_k$ is represented by

$$p_C(x) = \prod_{j=1}^{k} (1 - \ell_j(x)),$$

so that every Boolean assignment satisfying the clause makes p_C vanish.

A Nullstellensatz refutation of an unsatisfiable CNF is an identity

$$1 = \sum_C g_C(x) p_C(x) + \sum_{i=1}^{N} h_i(x)(x_i^2 - x_i),$$

with $g_C, h_i \in \mathbb{F}[x_1, \dots, x_N]$. Its degree is the maximum total degree of the displayed summands, and $\mathrm{NS}_{\mathbb{F}}(C)$ is the minimum such degree.

Polynomial Calculus starts from the clause polynomials and Boolean axioms and derives new polynomials by $\mathbb{F}$-linear combinations and multiplication by variables, ending with the constant polynomial 1. The degree of a derivation is the maximum degree appearing in it; $\mathrm{PC}_{\mathbb{F}}(C)$ is the minimum such degree. The standard simulation gives

$$\mathrm{PC}_{\mathbb{F}}(C) \le \mathrm{NS}_{\mathbb{F}}(C),$$

so a lower bound on Polynomial Calculus degree also lower-bounds Nullstellensatz degree.

Pitassi and Robere's lifting theorem associates to such a CNF C and a fixed constant-size gadget g a partially defined monotone function $f_{C,g}$. Recall that "good" is the exact rank-additivity condition stated before Proposition 4.3. In the form used here, if C has N variables, constant width, $O(N)$ clauses, and $\mathrm{NS}_{\mathbb{F}}(C) \ge \varepsilon N$, and if g is a good gadget whose matrix rank exceeds $2^{1/\varepsilon}$, then

$$\mu_{\mathbb{F}}(f_{C,g}) \ge 2^{\Omega(N)}.$$

The constants depend only on ε, the clause width, and the fixed gadget. Thus the role of Propositions 4.2 and 4.3 in the main text is exactly

$$\text{PC degree lower bound} \Longrightarrow \text{NS degree lower bound} \Longrightarrow \text{rank-measure lower bound}.$$

No other proof-complexity machinery is used in the paper.

B Reconstruction dual for partial affine reconstruction

Partial privacy gives an information gap for unauthorized coalitions, while the rank theorem of Theorem 3.5 needs a geometric lower bound on the target intersection of positive inputs and zero target intersection for negative inputs. This appendix proves Theorem 5.3 by converting affine reconstruction vectors into exactly that dual geometry. The orthogonality calculation follows the same pattern as duality arguments for linear secret sharing [16, Theorem 6.10] and for abelian schemes [20, Proposition 7.1].

Proof of Theorem 5.3. Let participant i receive d_i field elements and let $d = \sum_i d_i$. Index the share coordinates by

$$C = C_1 \sqcup \cdots \sqcup C_n,$$

where C_i indexes participant i's coordinates, and denote a full share vector by $\sigma(s, r) \in \mathbb{F}^d$.

The dual matrix. For every minimal authorized set $A \in \min \Gamma$ and every secret coordinate $j \in [d_0]$, perfect correctness and affine reconstructability give a vector $\lambda^{A,j} \in \mathbb{F}^d$, supported on the coordinates held by A, and a constant $b_{A,j} \in \mathbb{F}$ such that

$$\langle \lambda^{A,j}, \sigma(s, r) \rangle + b_{A,j} = s_j$$

for every valid sharing. Form the matrix

$$\Lambda \in \mathbb{F}^{d \times (d_0 | \min \Gamma|)}$$

whose column indexed by (A, j) is $\lambda^{A,j}$, and label each row indexed by C_i with participant i. For $j \in [d_0]$, let τ_j be the vector whose (A, j')-coordinate is $\mathbf{1}[j = j']$, and put

$$T_0 = \operatorname{span}\{\tau_1, \ldots, \tau_{d_0}\}.$$

Thus T_0 has dimension d_0.

For $X \subseteq P$, write

$$R_X = \operatorname{rowspan}(\Lambda_X), \qquad K_X = \ker(\Lambda_X).$$

Negative inputs. If $X \notin \Gamma^*$, then $P \setminus X \in \Gamma$. Choose a minimal authorized set $A \subseteq P \setminus X$. Suppose that

$$\sum_j \alpha_j \tau_j \in R_X.$$

Every reconstruction vector $\lambda^{A,j'}$ is supported inside A, which is disjoint from X. Hence every row of Λ_X is zero in the coordinate (A, j'). Examining this coordinate gives $\alpha_{j'} = 0$ for every j'. Therefore

$$R_X \cap T_0 = \{0\}.$$

Every negative input of f_{Γ^*} thus strongly rejects the target.

Positive inputs. Let $X \in \Gamma^*$ and put $B = P \setminus X \notin \Gamma$. Set

$$t_X = \dim(R_X \cap T_0).$$

Define

$$Q : \mathbb{F}^{d_0 | \min \Gamma|} \longrightarrow \mathbb{F}^{d_0}, \qquad Q(w)_j = \tau_j w = \sum_{A \in \min \Gamma} w_{A,j}.$$

With respect to the standard bilinear form,

$$K_X = R_X^{\perp}, \qquad \ker Q = T_0^{\perp}.$$

Consequently,

$$\begin{aligned}\dim Q(K_X) &= \dim K_X - \dim(K_X \cap T_0^{\perp}) \\ &= \dim R_X^{\perp} - \dim(R_X + T_0)^{\perp} \\ &= \dim(R_X + T_0) - \dim R_X \\ &= d_0 - \dim(R_X \cap T_0) = d_0 - t_X.\end{aligned}$$

Now take $w = (w_{A,j})_{A,j} \in K_X$ and put

$$\mu = \Lambda w = \sum_{A,j} w_{A,j} \lambda^{A,j} \in \mathbb{F}^d.$$

Since $\Lambda_X w = 0$, the vector μ is supported on the share coordinates of B. On every valid sharing,

$$\langle \mu, \sigma(s,r) \rangle = \sum_{A,j} w_{A,j}(s_j - b_{A,j}) = Q(w) \cdot s - \gamma_w, \qquad \gamma_w = \sum_{A,j} w_{A,j} b_{A,j}.$$

Thus coalition B can recover affinely every secret linear form whose coefficient vector lies in $Q(K_X)$. A basis of $Q(K_X)$ gives $d_0 - t_X$ independent linear forms of the uniform secret. Their joint entropy is $(d_0 - t_X)\log|\mathbb{F}|$, and their joint value is a deterministic function of X_B. Hence

$$I(X_0; X_B) \geq (d_0 - t_X)\log|\mathbb{F}|,$$

so

$$\Phi_\Pi(B) \geq 1 - \frac{t_X}{d_0}.$$

Partial privacy with advantage η gives $\Phi_\Pi(B) \leq 1 - \eta$, and therefore

$$t_X \geq \eta d_0 \qquad (X \in \Gamma^*).$$

Taking

$$\lambda = \min_{X \in \Gamma^*} t_X$$

gives $\lambda \geq \eta d_0$. Participant i labels exactly the d_i rows indexed by C_i, completing the proof. □

C A non-multilinear scheme below pairwise distance one

The conclusion of Theorem 5.6 is a replacement statement: the original scheme need be neither multilinear nor perfect. The following elementary example shows that both distinctions can occur simultaneously.

Consider the 2-out-of-2 access structure over $\mathbb{F}_3$. Let the secret S and a random variable R be independent and uniform on $\mathbb{F}_3$, and define

$$X_1 = R^2, \qquad X_2 = S + R^2.$$

Each participant receives one field element. The authorized set $\{1,2\}$ reconstructs the secret linearly:

$$S = X_2 - X_1.$$

Thus the scheme is affine-reconstructable, indeed linearly reconstructable.

The scheme is not perfect. Participant 1 sees R^2, whose distribution is independent of the secret. For participant 2, conditioned on $S = s$,

$$\Pr[X_2 = s] = \frac{1}{3}, \qquad \Pr[X_2 = s + 1] = \frac{2}{3},$$

with all arithmetic in $\mathbb{F}_3$. Hence for any two distinct secrets s, s',

$$\mathrm{SD}\big(X_2 \mid S = s,\ X_2 \mid S = s'\big) = \frac{2}{3}.$$

Therefore the pairwise statistical privacy parameter is $\delta = 2/3 < 1$, but the scheme is not perfectly private.

The scheme is also genuinely non-multilinear. For a fixed secret s, the joint share distribution is supported on

$$\{(0, s), (1, s + 1)\},$$

with probabilities $1/3$ and $2/3$, respectively. By contrast, in a multilinear scheme over $\mathbb{F}_3$, after fixing the secret the share vector is an affine-linear image of independent uniform field elements. Its distribution is therefore uniform on an affine subspace of the share space, whose cardinality is a power of 3. The two-point, nonuniform distribution above cannot have this form.

The perfect multilinear replacement promised by Theorem 5.6 can be seen explicitly in this example. Take a uniform $T \in \mathbb{F}_3$ and set

$$\widehat{X}_1 = T, \qquad \widehat{X}_2 = S + T.$$

This is a perfect scalar-linear 2-out-of-2 scheme with the same one-dimensional secret and one field coordinate for each participant.