

Resolving the Worst-Case Complexity of Linear Secret Sharing

Oded Nir*

September 18, 2026

Abstract

A secret-sharing scheme allows a dealer to distribute a secret among n parties such that only predefined “authorized” sets of parties can reconstruct the secret, and all other “unauthorized” sets learn nothing about it. Families of authorized sets are called access structures, and a scheme is called linear if its sharing map is linear in the secret and the dealer’s randomness. We show that every n -party access structure can be realized by a linear secret-sharing scheme over every finite field with at most $2^{\lceil n/2 \rceil - 1} + 1$ field elements per share. Taking the field to be $\mathbb{F}_2$ yields the same bound in bits for one-bit secrets. A known counting lower bound for monotone span programs shows that for every fixed finite field, almost all access structures require share size $2^{n/2 - o(n)}$ in every linear scheme, which makes our upper bound tight. Our scheme is considerably simpler than previous schemes that obtained share size $2^{cn + o(n)}$ with $1/2 \leq c < 1$.

We also present a variant of this linear construction that is tailored for monotone k -DNF access structures (also known as k -upslices). Then, by combining it with existing non-linear schemes, we derive a scheme for all access structures with share size $2^{0.494n + o(n)}$. This improves the previous upper bound of $2^{0.585n + o(n)}$ by Applebaum and Nir (CRYPTO 2021), and establishes a separation between the worst-case non-linear and linear exponents.

Additionally, we prove that almost all monotone functions require monotone span programs of size $2^{n/2 - o(n)}$ simultaneously over all fields, finite or infinite, improving the previous bound of $2^{n/3 - o(n)}$ for this setting. In particular, almost all access structures require information ratio $2^{n/2 - o(n)}$ in linear schemes over every finite field, including fields whose size depend on n .

Some of the constructions and proofs were discovered in conversations prompted by the author with GPT-5.6 Sol and Claude Fable 5.

1 Introduction

Secret-sharing schemes allow a dealer to distribute a secret among n parties so that every *authorized* coalition can reconstruct the secret while every other coalition learns nothing about it. The family of authorized coalitions is a monotone collection, identified with a monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and called an access structure¹. Secret-sharing schemes were first introduced by Shamir [35] and Blakley [19] for threshold functions, and Ito, Saito, and Nishizeki [28] were the first to realize general monotone functions.

A scheme is called *linear* over a field $\mathbb{F}$ if the secret is an element in $\mathbb{F}$, the dealer’s randomness is a vector of elements in $\mathbb{F}$, and every share is a vector of $\mathbb{F}$ -linear combinations of the secret

*Ben-Gurion University, Beer Sheva, Israel. Email: odednir123@gmail.com. ORCID: 0009-0008-7860-8151

¹Throughout the paper, we treat access structures as either families of authorized sets or their corresponding functions, according to the view which best fits the context.

and the randomness². Linear schemes are equivalent to the computational model of monotone span programs [30], and linearity is essential in many applications, including secure multiparty computation [20] and attribute-based encryption [36, 7], since parties can locally add shares of two secrets and obtain shares of their sum.

Upper bounds on the share size. The most important complexity measure for secret sharing is the size of the shares. Formally, given an access structure F we let $\text{SSize}(F)$ be the minimal total share size over all schemes that realize F with one-bit secrets. Most schemes discussed in this work have shares with size that is exponential in the number of parties n , and so we compare asymptotic share sizes through their exponential rates. We say that a construction whose worst-case share size is $2^{cn+o(n)}$ has exponent c , and following [4], we define the *secret-sharing exponent* of a monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ (for one-bit secrets) as $\mathbf{S}(f) := n^{-1} \cdot \log_2 \text{SSize}(f)$, and define the (worst-case) *secret-sharing exponent* of a family of functions $\mathcal{F}$ to be $\limsup_{n \rightarrow \infty} \max_{f \in \mathcal{F}} \mathbf{S}(f)$. We then denote by $\mathbf{S}$ the exponent of the family of all monotone functions over $\{0, 1\}^n$, and by $\mathbf{S}_\ell$ the *linear exponent*, which is defined analogously to $\mathbf{S}$, except that $\text{SSize}(f)$ is replaced with $\text{LSSize}(f)$, the minimal complexity of a linear scheme that realizes f .

For general access structures, a sequence of works has progressively reduced these exponents. Liu and Vaikuntanathan [31] were the first to break the barrier of $2^{n-o(n)}$ and prove that both $\mathbf{S}$ and $\mathbf{S}_\ell$ are smaller than 1. Their construction was based on conditional disclosure of secrets protocols (CDS) [26] composed according to a monotone formula. Applebaum et al. [3] improved their results by replacing one of the ingredients in their construction with a recursive scheme, and Applebaum et al. [4] offered a further improvement by introducing and designing robust conditional disclosure of secrets protocols (RCDS). The schemes of Applebaum and Nir [6] were based on a decomposition of functions into upslices (equivalently, monotone k-DNFs) and downslices (equivalently, monotone k-CNFs), and Alon, Beimel, and Lasri [1] optimized one part of the linear schemes of [4, 6] to obtain another improvement. The table below lists the exponents of these schemes, along with the exponents of the schemes presented in this paper:

Work	Non-linear exponent	Linear exponent
Ito-Saito-Nishizeki (1987) [28]	1	1
Liu-Vaikuntanathan (2018) [31]	0.994	0.999
Applebaum et al. (2019) [3]	0.892	0.942
Applebaum et al. (2020) [4]	0.637	0.762
Applebaum-Nir (2021) [6]	$\log_2(3/2) \approx 0.585$	0.7576
Alon-Beimel-Lasri (2025) [1]	—	0.7563
Theorem 1.1 (this work)	0.5	0.5
Theorem 1.5 (this work)	0.494	—

Table 1: Exponents c in share-size bounds of the form $2^{cn+o(n)}$. A dash indicates that the work does not provide a new non-linear/linear construction.

Lower bounds on the share size. On the lower bound front, for unrestricted schemes Csirmaz constructed n -party access structures for which the total share size is at least $\Omega(n^2/\log n)$ times

²Equivalently, a scheme is linear if every authorized coalition can reconstruct the secret as an $\mathbb{F}$ -linear combination of its shares [9].

the size of the secret [22], and recently this bound was improved to $\Omega(n^2)$ by Williamson [37].

For linear schemes, the strongest lower bounds are based on their equivalence to monotone span programs over finite fields. A *monotone span program* (MSP) consists of a matrix over a field, with each row assigned to one of the parties, and a target vector e_1 . A coalition is authorized exactly when the target vector is a linear combination of the rows assigned to its members. The size of the program is its number of rows, which corresponds to the total number of field elements shared in the equivalent linear secret-sharing scheme. In contrast to secret-sharing schemes, the definition of MSPs naturally extends to *infinite* fields.

Counting monotone span programs gives dramatically stronger bounds than the ones for unrestricted secret-sharing schemes. We define the *share length* of a party in a linear scheme as the number of field elements in its share (the maximal share-length is also known as the *information ratio* of the scheme), and list several known lower bounds for linear schemes. First, for every fixed finite field $\mathbb{F}$, almost all n -party access structures require total share length $2^{n/2-o(n)}$ [8]. Moreover, a field-uniform version exists for bit-complexity: almost all n -party access structures require share size $2^{n/2-o(n)}$ in every finite field [12, 11]. For every field $\mathbb{F}$ (whose size may depend on n), almost all n -party access structures require monotone span programs of size $2^{n/3-o(n)}$ over $\mathbb{F}$ [34], and the same lower bound was later shown to hold for almost all access structures simultaneously over every finite field [12].

In summary, prior work left open, among other questions, the following gaps. For linear schemes, there was an exponential gap between the upper and lower bounds on share length over small fixed fields, and an even larger gap over arbitrarily large fields. It was also unknown whether unrestricted schemes could achieve a smaller worst-case exponent than linear ones. We resolve these questions in the results presented next.

1.1 Our Results

1.1.1 Linear Secret Sharing

Our first result is the following:

Theorem 1.1 (Main theorem). *For every finite field $\mathbb{F}$, every monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be realized by a linear secret-sharing scheme over $\mathbb{F}$ with maximal share length of $2^{\lceil n/2 \rceil - 1} + 1$ field elements. When the field is $\mathbb{F}_2$ the maximal share size is $2^{\lceil n/2 \rceil - 1} + 1$ bits.*

This theorem has two direct consequences. First, it implies that every monotone function over n variables can be computed by a monotone span program over every finite field with size at most $n(2^{\lceil n/2 \rceil - 1} + 1)$, as the size of the program corresponds to the total number of field elements given to the parties. Second, if we combine it with the counting lower bound of [8], we get that the worst-case linear exponent $\mathbf{S}_\ell$ is exactly $1/2$. We also note that our scheme works over any finite abelian group, since the construction uses only addition, subtraction, and the sampling of uniform elements, which makes it a *black-box secret-sharing scheme*, as defined in [23, 21].

Next, by introducing some modifications to our main construction we obtain more than the bound for general functions: we can prove an upper bound on the exponent of every *layer*. Following [6], we define (k, n) -upslices as functions over n bits that can be represented by logical formulas in monotone disjunctive normal form (DNF) in which each term consists of exactly k variables. These functions are useful in the context of secret sharing since every monotone function can be decomposed into a disjunction of its upslices, i.e., $f = \bigvee_{a \in [n]} f_a$ where f_a is

the (a, n) -upslice function that agrees with f on its a -weight inputs (hereafter referred to as the a -upslice of f). Hence, due to standard closure properties of secret sharing over OR gates (Fact A.6), the secret-sharing complexity of worst-case monotone functions is at most n times larger than the secret-sharing complexity of upslices. We prove the following theorem which resolves the complexity of “high” upslices, and which will help us build the non-linear schemes with exponent below $1/2$.

Theorem 1.2 (Upslice bounds). *For every finite field $\mathbb{F}$, and for all integers $n \geq 1$, $1 \leq k \leq n$, $0 \leq q \leq n - 1$, and $\max\{0, k - (n - q)\} \leq t \leq \min\{k, q\}$, every (k, n) -upslice can be realized by a linear secret-sharing scheme over $\mathbb{F}$ with total share length at most*

$$\frac{1}{p_{k,q,t}} \left[\binom{n-q}{k-t} + \sum_{j=t}^q \binom{q}{j} \right] O(n^2 \log(n+1)),$$

where $p_{k,q,t} = \frac{\binom{q}{t} \binom{n-q}{k-t}}{\binom{n}{k}}$. Moreover, when $t = \lceil kq/n \rceil$, the same bound holds without the factor $1/p_{k,q,t}$.

Tight bounds for constant-density high upslices. For $\alpha \in (0, 1)$, let $U_\ell(\alpha)$ denote the worst-case linear exponent of $\lceil \alpha n \rceil$ -upslices, and let H_2 be the binary entropy function, with $H_2(0) = H_2(1) = 0$. For every constant $\alpha \in [1/2, 1)$, a simple choice of parameters $k = \lceil \alpha n \rceil$, $q = \lfloor n/2 \rfloor$, and $t = \lceil kq/n \rceil$ in Theorem 1.2 implies that $U_\ell(\alpha) \leq \frac{H_2(\alpha)}{2}$. With these parameters, both terms in the share-length bound are at most $2^{H_2(\alpha)n/2+o(n)}$ and the $1/p_{k,q,t}$ factor is removed.

A matching lower bound can be derived from known results. More concretely, Beimel et al. [15] define the rank of an access structure as the maximal size of its minterms, and prove in [15, Theorem 6.8] that for every family L of access structures with rank at most r , almost all access structures require maximal share size of $\Omega\left(\sqrt{\log |L|/rn}\right)$ from any linear secret-sharing scheme for one-bit secrets. For every k , the family of k -upslices has rank at most k and size $2^{\binom{n}{k}}$, and therefore by approximating the constant-density binomial coefficient we get that almost every upslice of density α requires linear share size $\Omega\left(\sqrt{\binom{n}{\alpha n}/\alpha n^2}\right) = 2^{H_2(\alpha)n/2-o(n)}$. This implies the following corollary:

Corollary 1.3 (High upslice exponents). *For every constant $\alpha \in [1/2, 1)$ it holds that $U_\ell(\alpha) = \frac{H_2(\alpha)}{2}$.*

Upper bounds for constant-density low upslices. For upslices of constant density $0 < \alpha < 1/2$, we choose $q = \left\lfloor \frac{H_2(\alpha)}{1+H_2(\alpha)} n \right\rfloor$ and $t = \lceil kq/n \rceil$, which implies the following corollary, proved as part of Corollary 6.4³.

Corollary 1.4 (Low upslice exponents). *For every constant $\alpha \in (0, 1/2)$ it holds that*

$$U_\ell(\alpha) \leq \frac{H_2(\alpha)}{1 + H_2(\alpha)} < \frac{1}{2}.$$

³These exponents improve on those of the previously best-known upslice schemes of [14] (based on [6]), even when compared with their non-linear schemes for upslices with density below $1/2$.

By examining different assignments for q and t (specifically, choosing t different from $\lceil kq/n \rceil$), we can obtain sharper bounds than the analytical estimate of Corollary 1.4. A simple numerical optimization leverages this to later yield an exponent of 0.494 for unrestricted schemes, as presented in the following Theorem 1.5 (compared to an exponent of 0.496, which can be obtained without this optimization).

1.1.2 Non-linear Secret Sharing

Theorem 1.1 already improves the best unrestricted exponent from $\log_2(3/2) \approx 0.585$ [6] to $1/2$. One could have conjectured that this is also the worst-case complexity of general schemes, but we show that this is not the true bound. To go below $1/2$, we combine our linear scheme for upslices with a few building blocks and ideas developed in previous works, including the now-classical framework of Liu and Vaikuntanathan [31], duality for upslices [6], and a non-linear scheme of [4]. We prove the following theorem:

Theorem 1.5 (Beyond the linear exponent). *Every monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be realized by a secret-sharing scheme for one-bit secrets with total share size $2^{0.494n+o(n)}$. Hence $\mathbf{S} \leq 0.494 < \frac{1}{2} = \mathbf{S}_\ell$, and the non-linear secret-sharing exponent is strictly smaller than the linear one.*

Previously, a separation between the unrestricted and linear exponents was known only under unproven hypotheses [6, 14], or for restricted families of functions such as slices [31, 15].

We also prove that a similar separation exists between linear schemes and *quadratic* schemes. In a quadratic secret-sharing scheme, both sharing and reconstruction are computed by polynomials of degree at most two over the same field [16]. We obtain the following separation:

Theorem 1.6 (Quadratic schemes beyond the linear exponent). *Every monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be realized by a secret-sharing scheme with quadratic sharing and reconstruction with total share size at most $2^{0.4994n+o(n)}$. Hence the quadratic secret-sharing exponent is at most $0.4994 < \frac{1}{2} = \mathbf{S}_\ell$, and is strictly smaller than the linear one.*

In contrast to the linear case, this quadratic upper bound is not known to be tight: the best-known lower bound for schemes with reconstruction degree d is $2^{n/(d+1)-o(n)}$, which for $d = 2$ gives only $2^{n/3-o(n)}$ [16].

Remark 1.7 (Schemes with reconstruction of degree d). The same argument as in Theorem 1.6 extends beyond quadratic reconstruction. Using the degree- d reconstruction RCDS protocols of Beimel, Farràs, and Lasri [13, Corollary 7.2] in our construction gives a sequence of degree- d -reconstruction schemes whose exponents approach our unrestricted exponent of Theorem 1.5 as $d \rightarrow \infty$.

1.1.3 A Monotone Span-Program Lower Bound over All Fields

We prove the following lower bound on the size of monotone span programs over all fields.

Theorem 1.8 (Monotone span programs over all fields). *Almost all n -party access structures require monotone span-program size at least $2^{n/2-o(n)}$ over every field.*

By the known equivalence of monotone span programs over finite fields and linear secret sharing, this proves that almost all n -party access structures require linear secret-sharing schemes

with total share length $2^{n/2-o(n)}$ over every finite field. Together with our upper bounds, this resolves the worst-case complexity of linear secret sharing uniformly over all finite fields.

This lower bound allows us to obtain yet another separation. As observed in [33], some constructions that are based on *conditional disclosure of secrets* (CDS) protocols [26] with exponent c can be converted to multilinear schemes for sufficiently long secrets with information ratio $2^{cn+o(n)}$, by using multilinear CDS protocols with constant information rate [2]. A multilinear scheme is a secret-sharing scheme in which the secret is a vector of d field elements and the shares are linear functions of the secret elements and the dealer's randomness over the same field (a linear scheme is then a multilinear one with $d = 1$). This observation of [33] also holds for our non-linear construction, since its non-linear parts are based on a formula with one layer of CDS gates (see [5] for a description of this formula), and the rest of its steps are linear. Together with Theorem 1.8, this yields the following separation.

Corollary 1.9 (Separation between linear and multilinear schemes for long secrets). *There exists a multilinear secret-sharing scheme for all access structures and for sufficiently long secrets with information ratio $2^{0.494n+o(n)}$, while almost all access structures require information ratio $2^{n/2-o(n)}$ in every linear scheme, even for long secrets (i.e., for large fields).*

1.2 Our techniques

The linear scheme for general access structures. Our scheme is based on a relatively simple concept. In a nutshell, the n secret-sharing parties are partitioned into left and right parts L and R . Then, the dealer distributes to the left parties many sets of additive shares of maskings of the secret, where each masking is done with a different set of random pads, and the right parties receive different sets of the pads. The same approach was used in the 2-server RCDS protocol of [17] (which in turn was inspired by the work of [25], and also appeared in [4, Lemma B.1]), and our scheme can be seen as a natural development of it that uses different sets of pads for the right-side parties.

We now provide more details. Fix an arbitrary finite field $\mathbb{F}$. Let F be an access structure over n parties, let $s \in \mathbb{F}$ be the secret, and fix a partition of $[n]$ into two disjoint parts L and R , that we will call the left side or the *masking side*, and the right side or the *disclosure side*. For every $T \subseteq R$, the dealer samples an independent uniform pad $\omega_T \in \mathbb{F}$, and for every $U \subseteq L$ it computes a masked value of the secret $m_U = s + \sum_{T: F(U \cup T)=0} \omega_T$. That is, for a left-side set U , m_U is the secret masked by pads ω_T for right-side sets T such that $U \cup T$ is unauthorized in F . The dealer then shares these masked values and the pads according to the following steps.

- For every nonempty $U \subseteq L$, share m_U additively among the parties of U . Give a copy of $m_\emptyset$ to every party of R .
- For every $T \subseteq R$, give ω_T to every party in $R \setminus T$.

We first give a sketch of the relatively straightforward arguments for correctness and privacy. Let B be a coalition, and write $C = B \cap L$ and $D = B \cap R$. If B is authorized, its parties can reconstruct m_C from the additive scheme (or receive $m_\emptyset$ directly), and additionally, they hold every pad that masks the secret in m_C : indeed, if a pad ω_T masks m_C , then by definition the coalition $C \cup T$ is unauthorized. Consequently, $D \not\subseteq T$, since otherwise $B = C \cup D$ would be contained in the unauthorized coalition $C \cup T$ and would also be unauthorized by monotonicity. Therefore, every party in $D \setminus T$ holds the pad ω_T . Subtracting these pads from m_C recovers s .

If B is unauthorized, the pad ω_D (for $D = B \cap R$) masks every value m_U (shared to a set $U \subseteq C$) that the parties of B can reconstruct. This holds since $U \cup D \subseteq B$, which implies that $U \cup D$ is unauthorized by monotonicity. In addition, ω_D is not disclosed to the disclosure-side parties of B , since they are exactly the parties in D . Therefore, a single hidden pad masks the secret in the entire view of the coalition, which suffices for proving privacy.

For the share size, each party of L receives one additive share for every subset $U \subseteq L$ that contains it, and hence holds $2^{|L|-1}$ field elements. Each party of R receives $m_\emptyset$ and one pad for every T that does not contain it, and hence holds $2^{|R|-1} + 1$ field elements. Choosing $|R| = \lceil n/2 \rceil$ gives the desired share length.

The linear scheme for upslices. Our linear scheme for upslices is based on a realization of a more structured family of functions we call *partitioned k -upslices*. In a partitioned k -upslice f_Π , parametrized by a partition Π of the inputs to left and right parts, every minterm has weight k , with exactly $k - t$ 1's on the left side and t on the right. Then, a general upslice can be realized by composing partitioned upslices over a family of partitions with a suitable covering property. We describe the scheme for partitioned upslices and then their composition.

Given a partition Π into L and R of sizes $n - q$ and q respectively, every minterm of a partitioned upslice f_Π is of the form $U \cup V$ with $|U| = k - t$, $|V| = t$ and $U \subseteq L$, $V \subseteq R$. We think of the parties of L and R as the masking-side parties and the disclosure-side parties in the general construction above.

Now, in contrast to the linear scheme for general functions, the dealer computes the masked values only for sets $U \subseteq L$ of size $k - t$, and samples independent uniform pads $\omega_T \in \mathbb{F}$ only for sets in $\mathcal{T} = \{T \subseteq R \text{ with } |T| \geq t\}$ (which will later allow for savings in the share size). To ensure that subsets of R of size smaller than t receive no information, we sample one extra independent uniform pad $\eta \in \mathbb{F}$ that masks every m_U and is shared among the parties of R using a t -out-of- q threshold scheme (Definition A.4). That is, for every $U \in \binom{L}{k-t}$, the dealer computes

$$m_U = s + \eta + \sum_{T \in \mathcal{T}: f_\Pi(U \cup V) = 0 \text{ for every } V \in \binom{T}{t}} \omega_T.$$

The dealer then distributes the following shares:

- For every $U \in \binom{L}{k-t}$, share m_U additively among the parties of U .
- Share η among the parties of R using a linear t -out-of- q threshold scheme over $\mathbb{F}$. For every $T \subseteq R$ with $|T| \geq t$, give ω_T to every party in $R \setminus T$.

Correctness and privacy follow from similar arguments to those in the general scheme, with one additional distinction in the privacy argument. For an unauthorized coalition B , write $D = B \cap R$. If $|D| \geq t$, the undisclosed pad ω_D masks the secret in every m_U that the coalition can reconstruct. If $|D| < t$, the threshold sharing keeps η hidden, and this pad instead masks the secret in all such values.

The share length will now depend on k and decrease as k moves away from $n/2$. There are $\binom{n-q}{k-t}$ masked values m_U , and each is additively shared among $k - t \leq n$ parties, summing to at most $n \binom{n-q}{k-t}$ field elements in total. For every set T of size j , the field-element pad ω_T is given to the $q - j \leq n$ parties of $R \setminus T$, contributing at most $n \binom{q}{j}$ field elements in total for all sets T of size j . Finally, by Definition A.4, the threshold sharing of η over $\mathbb{F}$ uses $\lceil \log_{|\mathbb{F}|}(q + 1) \rceil$ field

elements per party of R , for a total share length of $O(n \log n)$. Thus, the total share length is at most

$$\left[\binom{n-q}{k-t} + \sum_{j=t}^q \binom{q}{j} \right] O(n \log n). \quad (1)$$

It remains to realize general upslices from partitioned ones. This can be done with a standard construction based on a family of partitions $\Pi_1, \dots, \Pi_z$ that have a covering property: every k -subset $A \subseteq [n]$ has profile $(k-t, t)$ in at least one partition. Given a (k, n) -upslice f_k , for every $i \in [z]$ let f_{Π_i} be the partitioned upslice whose minterms are precisely the minterms of f_k having profile $(k-t, t)$ in Π_i . By the covering property, every minterm of f_k is a minterm of at least one f_{Π_i} , while every minterm of each f_{Π_i} is a minterm of f_k . Therefore, $f_k = \bigvee_{i=1}^z f_{\Pi_i}$, which results in share length (1) times z by the standard closure of secret sharing under OR.

We can choose the partitions at random: for a uniformly random q -set R and $L = [n] \setminus R$, a fixed k -set has the desired profile $(k-t, t)$ with probability $p_{k,q,t}$, as defined in Theorem 1.2. By the union bound, choosing $z = O(n/p_{k,q,t})$ independent partitions gives a positive probability of covering every k -set simultaneously, so such a cover exists. Alternatively, for the typical profile $t = \lceil kq/n \rceil$, which yields the optimal exponent for high upslices with $k \geq n/2$, we can use an explicit family of only n partitions: arrange the parties on a cycle and let R range over all cyclic windows of q consecutive parties. These window partitions cover every k -set with the typical profile, by the classical sliding-window argument, which also appeared in discrete necklace splitting [27].

Going below the linear exponent. To prove Theorems 1.5 and 1.6, we follow the framework of Liu and Vaikuntanathan [31], and express any monotone function f in terms of three monotone functions, f_{bot} , f_{top} , f_{mid} , that agree with f on inputs of a low, high, and middle Hamming weight, respectively. In more detail, let $\delta \in (0, 1/2)$ be a parameter that sets the boundaries between these ranges, and then the low range is below $a = (1/2 - \delta)n$, the high is above $b = (1/2 + \delta)n$, and the middle is within $[a, b]$. The decomposition holds for every such δ , whose value we choose later to optimize the share size.

We formally define the functions above as follows. $f_{\text{bot}}(x) = 1$ exactly when x contains a minterm of f of weight less than a . Thus it agrees with f on inputs of weight less than a . Dually, we define $f_{\text{top}}(x) = 0$ exactly when x is contained in a maxterm of f of weight greater than b . Then it agrees with f on inputs of weight greater than b . Finally, f_{mid} agrees with f on inputs of weight in $[a, b]$, outputs 0 below this interval, and outputs 1 above it. We call this last function the δ -midslice of f . Then, f can be expressed in the following way:

$$f = f_{\text{top}} \wedge (f_{\text{mid}} \vee f_{\text{bot}}).$$

We include the case-by-case verification of this identity (as done in [31, Lemma 3.3]) in Section 6. Thus, it suffices to realize these three access structures and compose their schemes using OR and AND.

Next, we realize f_{bot} with parameter δ using our upslice construction, relying on the fact that f_{bot} can be described as an OR of k -upslices with $k \leq \lceil (1/2 - \delta)n \rceil - 1$. We then realize f_{top} with the same share size as f_{bot} using secret-sharing duality. The dual of a monotone function $g : \{0, 1\}^n \rightarrow \{0, 1\}$ is $g^*(x) = 1 - g(\bar{x})$, where $\bar{x}$ denotes the bitwise complement of x . It is known that if g has a linear secret-sharing scheme, then g^* has a linear scheme over the same field with the same total share size (e.g., [24]). Then, the minterms of f_{top}^* are the complements

of the maxterms of f_{top} , and therefore have size less than $n - b = a$. Hence, the dual of every f_{top} is an instance of f_{bot} , and both function families can be linearly realized with the same share sizes⁴.

For the middle part, we use the non-linear midslice construction of Applebaum et al. [4]. We note that by the counting lower bound of [15], non-linearity is necessary to achieve a δ -midslice exponent below $1/2$ for every $\delta > 0$. Choosing the width $\delta = 0.08026$ along with other optimized parameters in the upslice scheme gives the exponent 0.494. Replacing only the middle component with a quadratic midslice construction based on the work of [16] and choosing $\delta = 0.02664$ (again, with other optimized parameters) gives the exponent 0.4994. Since the remaining components and the outer compositions are linear, this replacement preserves quadratic sharing and reconstruction.

The lower bound over all fields. Our lower bound for monotone span programs over all fields relies on an upper bound of Nelson [32] on the number of possible zero patterns of polynomials. A *zero pattern* of a family of N polynomials is a list of N zeros and nonzeros obtained by evaluating all the polynomials at the same point over some field. Using zero patterns for span-program lower bounds goes back to Babai et al. [8]. They showed that for every field $\mathbb{F}$ of characteristic zero, almost all monotone functions require MSPs of size $2^{n/3-o(n)}$. Rónyai et al. [34] extended this result to every field. Beimel and Farràs [12] proved the same lower bound as [34] with a different quantifier order by using the newer zero-pattern result of Nelson. They showed that almost all access structures require share length $2^{n/3-o(n)}$ in every linear secret-sharing scheme, simultaneously over all finite fields.

We improve the bound in the stronger sense of [12]. Our proof strategy generally follows that of previous work: We encode each monotone span program of size d as an assignment to a list of polynomials, such that MSPs that compute different functions will be encoded to a different zero pattern of this list. Then, the upper bound on the zero patterns will bound the possible number of functions that can be computed by an MSP of size d .

The source of our improvement is a more efficient encoding of span programs as polynomials. Specifically, in previous works, an MSP matrix was encoded by the linear dependencies among all subsets of its d rows. In contrast, we use the fact that once the row labeling is fixed, the shares of each coalition depend on the rows belonging to its parties, and authorization depends only on whether adjoining the target vector increases the rank of these rows. We construct a single family of polynomials whose zero patterns determine these two ranks for every coalition, and hence determine the access structure.

1.3 Open Problems

Improved exponent via optimal linear schemes for low upslices. By Corollary 1.4 and the lower bound argument above, it holds that the exponent $U_\ell(\alpha)$ of upslices with density α below $1/2$ satisfies

$$\frac{H_2(\alpha)}{2} \leq U_\ell(\alpha) \leq \frac{H_2(\alpha)}{1 + H_2(\alpha)}.$$

⁴The same duality argument implies that the duals of k -upslices, which are the $(n - k)$ -downslice functions, studied in [6], can be realized by linear schemes with the same share sizes promised for upslices in Theorem 1.2 and its corollaries.

Hence, it remains open whether every $(\alpha n, n)$ -upslice with $\alpha < \frac{1}{2}$ can be realized by a linear scheme with total share size $2^{H_2(\alpha)n/2+o(n)}$. Such optimal schemes for low upslices can be plugged directly into the proof of Theorem 1.5, and lead to an improved unrestricted exponent smaller than 0.491 instead of 0.494, with δ that is roughly 0.079.

Monotone span programs vs. formulas over slices. As mentioned before, linear secret-sharing schemes are equivalent to the computational model of monotone span programs (MSPs). In contrast, all the known nontrivial upper bounds for non-linear secret sharing for general functions arise from constructions in a different computational model, named formulas over slices (FOS), formally introduced in [5]. Prior to this work, the FOS constructions for general functions were stronger than the MSP ones, peaking at FOS of size $1.5^{n+o(n)}$ [6]. Following our results, MSPs are now in the lead, and a combined FOS-MSP model bypasses the MSP lower bound. This brings forward the following questions, some of them already discussed in [5]: 1) Is the FOS model stronger than MSPs for general functions? More concretely, does every function have an FOS of size $2^{(0.5-\epsilon)n+o(n)}$ for some constant $\epsilon > 0$? 2) Can a computational model of formulas with both MSP and slice gates help circumvent the existing FOS lower bound of $2^{\Omega(n/\log^2 n)}$ of [5]? Alternatively, can we prove significant lower bounds in this stronger model?

1.4 Organization

In Section 2, we fix notation. In Section 3, we present the linear scheme for general access structures and prove Theorem 1.1. In Section 4, we present the scheme for a single partitioned upslice and prove its correctness, privacy, and share-size bound. In Section 5, we describe partition families with a covering property and combine them with the results of the previous section to prove upslice bound of Theorem 1.2. In Section 6, we prove the non-linear and quadratic bounds of Theorems 1.5 and 1.6. In Section 7, we prove the lower bound of Theorem 1.8.

In Appendix A, we provide the formal definitions of secret-sharing schemes, basic secret-sharing tools, and function families used in the paper, and in Appendix B we provide proofs of some technical lemmas used in the paper.

1.5 Statement on AI Use

The upslice construction, the lower bound result and their proofs were developed through author-prompted conversations with GPT-5.6 Sol and Claude Fable 5. We reached the rest of the upper bounds later based on the ideas of the upslice construction.

We asked GPT-5.6 Sol to devise a research plan to improve the share size of linear secret sharing for general access structures. This first led to a linear scheme with exponent 0.73 based on schemes for upslices, implemented with two-server RCDS protocols similar to those of [4]. We then asked both models to improve the result by suggesting specific techniques. This led GPT-5.6 Sol to output a linear scheme with exponent $2/3$, based on an optimization of the previous RCDS-based upslice scheme. Asking for further improvements led to a linear scheme with exponent $1/2$ based on a K -server RCDS protocol with $K = \Theta(n^{1/3})$. We then identified several ways to simplify the scheme, and feeding these ideas back to GPT-5.6 Sol led to more simplifications and to the scheme presented in the paper. The lower bound result was obtained with a one-shot prompt.

We prepared the entire manuscript, using model-generated drafts as a starting point in some cases. We manually verified all of the proofs.

2 Preliminaries

Standard background on secret sharing and formal definitions of slices, multislices, midslices and upslices are deferred to Appendix A. The formal definition of monotone span programs is deferred to Section 7 where it is used.

2.1 Notation

All logarithms are base two, except $\ln$, which denotes the natural logarithm. We identify a coalition with a subset of $[n]$. For a finite set R and an integer r , $\binom{R}{r}$ denotes the family of all r -subsets of R . The binary entropy function is $H_2(\rho) = -\rho \log \rho - (1 - \rho) \log(1 - \rho)$, with $H_2(0) = H_2(1) = 0$. We denote by $\kappa(\rho)$ the function that equals 1 for $\rho \in [0, 1/2]$ and $H_2(\rho)$ for $\rho \in [1/2, 1]$.

Notation on Monotone Boolean Functions. The *weight* of an input $x \in \{0, 1\}^n$, denoted $\text{wt}(x)$, is the number of bits in x that are one, i.e., $\text{wt}(x) = |\{i : x_i = 1\}|$. For two strings $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in \{0, 1\}^n$, we say that $x \leq y$ if $x_i \leq y_i$ for every $1 \leq i \leq n$. A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is *monotone* if $x \leq y$ implies $f(x) \leq f(y)$.

A *minterm* of a monotone function f is a minimal input $x \in \{0, 1\}^n$ such that $f(x) = 1$, i.e., for every $y \neq x$, if $y \leq x$ then $f(y) = 0$. A *maxterm* of a monotone function f is a maximal input $x \in \{0, 1\}^n$ such that $f(x) = 0$, i.e., for every $y \neq x$, if $y \geq x$ then $f(y) = 1$.

3 Linear Secret Sharing for General Access Structures

In this section, we describe the linear scheme for general access structures and prove Theorem 1.1. Let $\mathbb{F}$ be a finite field, and let f be an access structure over $[n]$. Let $L \sqcup R$ be a partition of $[n]$ into nonempty sets with $|L| = \ell$ and $|R| = q$. For every $U \subseteq L$ and $T \subseteq R$, define the public coefficient $\gamma_U(T) \in \{0, 1\} \subseteq \mathbb{F}$ by

$$\gamma_U(T) = 1 \iff f(U \cup T) = 0. \quad (2)$$

We next describe the scheme. Its correctness, privacy, and share size are established in Theorem 3.2, which implies Theorem 1.1.

Construction 3.1 (Linear scheme for general access structures).

Let $s \in \mathbb{F}$ be the secret. The dealer samples independent uniform field elements $(\omega_T)_{T \subseteq R}$, and distributes the following shares over $\mathbb{F}$.

Masking side. For every $U \subseteq L$, compute

$$m_U = s + \sum_{T \subseteq R} \gamma_U(T) \omega_T. \quad (3)$$

1. For every nonempty $U \subseteq L$, share m_U among the parties of U using additive secret sharing.

Disclosure side. Do the following over the parties of R :

1. Give a copy of $m_\emptyset$ to every party of R .
2. For every $T \subseteq R$, give ω_T to each party in $R \setminus T$.

Theorem 3.2 (The general linear scheme). *For every finite field $\mathbb{F}$, every access structure f over $[n]$, and every partition $[n] = L \sqcup R$ into nonempty sets with $|L| = \ell$ and $|R| = q$, Construction 3.1 is a linear secret-sharing scheme over $\mathbb{F}$ for f , in which every party of L holds $2^{\ell-1}$ elements of $\mathbb{F}$ and every party of R holds $2^{q-1} + 1$ elements of $\mathbb{F}$. Consequently, taking $q = \lceil n/2 \rceil$, every n -party access structure can be realized over every finite field $\mathbb{F}$ by a linear secret-sharing scheme with a maximal share length $2^{\lceil n/2 \rceil - 1} + 1$ field elements.*

Proof. We first notice that the scheme is linear over $\mathbb{F}$, as every distributed value is an $\mathbb{F}$ -linear combination of the secret and the dealer's random field elements. For the share size, a party of L holds one field element for each of the $2^{\ell-1}$ nonempty sets $U \subseteq L$ containing it, and a party of R holds a copy of $m_\emptyset$, together with the field element ω_T for each of the 2^{q-1} sets $T \subseteq R$ not containing it.

Correctness. Let B be authorized, and set $C = B \cap L$ and $D = B \cap R$. We first claim that the parties of B can compute m_C . If $C \neq \emptyset$, the parties of C hold all the additive shares of m_C and reconstruct it. If $C = \emptyset$, then $B = D$ is nonempty (as it is an authorized coalition), and each of its parties holds a copy of $m_\emptyset = m_C$.

Next, if D is empty, then $C = B$ is authorized. By monotonicity, $f(C \cup T) = 1$ for every $T \subseteq R$, so all the coefficients $\gamma_C(T)$ vanish and $m_C = s$. Otherwise, consider any $T \subseteq R$ with $\gamma_C(T) = 1$, i.e., $f(C \cup T) = 0$ by (2), so that ω_T masks the secret. It cannot hold that $D \subseteq T$, since monotonicity would then give $f(B) = f(C \cup D) \leq f(C \cup T) = 0$, contradicting that B is authorized. Hence, $D \not\subseteq T$, and a party in $D \setminus T \subseteq B$ has ω_T . Subtracting in $\mathbb{F}$ every masking pad ω_T from m_C yields s .

Privacy. Let B be unauthorized, and set $C = B \cap L$ and $D = B \cap R$. The proof idea is to show that ω_D masks the secret in every “visible” linear combination in which the secret appears, and that ω_D is not released to the parties of B and is independent of all the other random elements. Lemma 3.3 below then shows that this suffices for privacy.

Formally, for every nonempty $U \not\subseteq C$, the parties of B hold only a strict subset of the additive shares of m_U . By the privacy of additive sharing, and the fact that these sharings are done with mutually independent fresh randomness, the joint view of all these incomplete sharings is distributed as $N\rho_1$ for some matrix N and a random vector ρ_1 that is independent of the secret, the pads, and the randomness used in all remaining shares.

Let ρ_2 collect all dealer randomness occurring in the remaining view of B , except for the pad ω_D . This pad is uniform and independent of (ρ_1, ρ_2) .

We claim that, in every linear combination in the remaining view, the coefficient of s equals the coefficient of ω_D . Consider first the complete additive sharing associated with a nonempty $U \subseteq C$. Since $U \cup D \subseteq B$, monotonicity and $f(B) = 0$ imply that $f(U \cup D) = 0$. Hence, $\gamma_U(D) = 1$ by (2), so m_U contains s and ω_D with the same coefficient. It follows that every additive share of m_U also has equal coefficients on s and ω_D : the fresh random shares have coefficient zero on both, while the remaining share has coefficient one on both. If $D \neq \emptyset$, the same argument with $U = \emptyset$ shows that every visible copy of $m_\emptyset$ has coefficient one on both s and ω_D . Finally, consider the pads individually disclosed to B by item (2). None of these is ω_D , since that pad is given only to parties in $R \setminus D$, whereas the disclosure-side parties of B are exactly those in D . Thus, every individually disclosed pad has coefficient zero on both s and ω_D .

Therefore, for some vector a and matrix M , the entire view of B is distributed as

$$\text{View}_B(s) \stackrel{d}{=} (N\rho_1, a(s + \omega_D) + M\rho_2).$$

This satisfies the conditions of Lemma 3.3, and hence the view of B is independent of the secret, which concludes the proof. $\square$

Lemma 3.3 (Common-pad privacy). *Let Π be a linear share-distribution procedure over a finite field $\mathbb{F}$ for a secret $s \in \mathbb{F}$, and fix a coalition B . Denote the view of B when the secret is s by $\text{View}_B(s)$. Suppose that there exist a uniform random pad $\omega \in \mathbb{F}$, random vectors $\rho_1 \in \mathbb{F}^{r_1}$ and $\rho_2 \in \mathbb{F}^{r_2}$ such that ω is independent of the pair (ρ_1, ρ_2) , matrices $N \in \mathbb{F}^{m_0 \times r_1}$ and $M \in \mathbb{F}^{m_1 \times r_2}$, and a vector $a \in \mathbb{F}^{m_1}$, and for every $s \in \mathbb{F}$ it holds that*

$$\text{View}_B(s) \stackrel{d}{=} (N\rho_1, a(s + \omega) + M\rho_2).$$

Then the distribution of $\text{View}_B(s)$ is independent of s .

Proof. Fix two secrets $s_0, s_1 \in \mathbb{F}$ and define $\omega' = \omega + s_0 - s_1$. Since translation by $s_0 - s_1$ is a bijection of $\mathbb{F}$, the variable ω' is uniform over $\mathbb{F}$. Moreover, because ω is independent of (ρ_1, ρ_2) , the triple $(\omega', \rho_1, \rho_2)$ has the same distribution as (ω, ρ_1, ρ_2) . Using the same values of ρ_1 and ρ_2 , and $\omega' = \omega + s_0 - s_1$, we have the pointwise identity

$$\begin{aligned} (N\rho_1, a(s_1 + \omega') + M\rho_2) &= (N\rho_1, a(s_1 + \omega + s_0 - s_1) + M\rho_2) \\ &= (N\rho_1, a(s_0 + \omega) + M\rho_2). \end{aligned}$$

The left-hand side has the distribution of $\text{View}_B(s_1)$, and the right-hand side has the distribution of $\text{View}_B(s_0)$. Hence these two views are identically distributed. Since s_0 and s_1 were arbitrary, the view of B is independent of the secret. $\square$

4 A Linear Scheme for One Partitioned Upslice

In this section, we show how to build secret-sharing schemes over any finite field for partitioned upslices. In Section 5, we will show how to combine these schemes with a suitable set of partitions in order to realize k -upslice functions, thus proving Theorem 1.2. Throughout this section, fix a finite field $\mathbb{F}$, integers $1 \leq k \leq n$, a partition $\Pi : [n] = L \sqcup R$ with $|R| = q$, and an integer t with $1 \leq t \leq \min(k, q)$ and $1 \leq k - t \leq n - q$. We again call L the *masking side* and R the *disclosure side*. A *partitioned upslice* f_Π with profile $(k - t, t)$ is a k -upslice in which every minterm M satisfies

$$|M \cap L| = k - t, \quad |M \cap R| = t.$$

Thus, for $U \in \binom{L}{k-t}$ and $V \in \binom{R}{t}$, we have $f_\Pi(U \cup V) = 1$ if and only if $U \cup V$ is a minterm of f_Π . We call a set $T \subseteq R$ with $|T| \geq t$ a *trace*. We let $\mathcal{T}$ denote the family of traces, and for every $U \in \binom{L}{k-t}$ and $T \in \mathcal{T}$ define the public coefficient $\gamma_U(T) \in \{0, 1\} \subseteq \mathbb{F}$ by

$$\gamma_U(T) = 1 \iff f_\Pi(U \cup T) = 0. \quad (4)$$

By monotonicity and the definition of f_Π , this is equivalent to $\gamma_U(T) = 1$ iff $f_\Pi(U \cup V) = 0$ for every $V \in \binom{T}{t}$.

We next describe the scheme. Its correctness, privacy, and share size are established in Theorem 4.2.

Construction 4.1 (Linear secret sharing for a partitioned upslice).

Let $s \in \mathbb{F}$ be the secret. The dealer samples independent uniform field elements η and $(\omega_T)_{T \in \mathcal{T}}$, and distributes the following shares over $\mathbb{F}$.

Masking side. For every $U \in \binom{L}{k-t}$, compute

$$m_U = s + \eta + \sum_{T \in \mathcal{T}} \gamma_U(T) \omega_T. \quad (5)$$

1. Share m_U among the parties of U using additive secret sharing.

Disclosure side. Do the following over the parties of R :

1. Share η using the $\mathbb{F}$ -linear t -out-of- q Shamir threshold scheme of Definition A.4.
2. For every $T \in \mathcal{T}$, give ω_T to each party in $R \setminus T$.

Theorem 4.2 (A linear scheme for partitioned upslices). *For every finite field $\mathbb{F}$, Construction 4.1 is a linear secret-sharing scheme over $\mathbb{F}$ realizing f_Π . Setting $m = \lceil \log_{|\mathbb{F}|}(q+1) \rceil$, its total share length is at most*

$$(k-t) \binom{n-q}{k-t} + qm + q \sum_{j=t}^q \binom{q}{j}$$

elements of $\mathbb{F}$.

Proof. The proof is very similar to that of Theorem 3.2. We first notice that the scheme is linear over $\mathbb{F}$, as every distributed value is an $\mathbb{F}$ -linear combination of the secret and the dealer's randomness, and the threshold scheme of Definition A.4 is $\mathbb{F}$ -linear. For the share length, a party in L holds one field element for each of the $\binom{n-q-1}{k-t-1}$ sets U containing it, for a total of $(n-q) \binom{n-q-1}{k-t-1} = (k-t) \binom{n-q}{k-t}$ field elements over all parties in L . The parties in R hold together qm field elements from the threshold scheme. In addition, every trace pad is given to at most q parties, and $|\mathcal{T}| = \sum_{j=t}^q \binom{q}{j}$. This proves the claimed field-element bound.

Correctness. Let B be authorized in f_Π . Hence, B contains some minterm M of f_Π . Let $U = M \cap L$ and $V = M \cap R$, so that $|V| = t$. The parties first reconstruct m_U from the additive shares of the masking side. Next, since $|V| = t$, the parties on the disclosure side can reconstruct η . For every trace T with $\gamma_U(T) = 1$ that masks the secret in m_U it cannot hold that $V \subseteq T$, since that would imply that $f_\Pi(U \cup V) = 0$ by (4). Hence, $V \not\subseteq T$, and some party in $V \setminus T$ receives ω_T by disclosure-side item (2). Subtracting in $\mathbb{F}$ the element η and every pad occurring in (5) from m_U yields s .

Privacy. Let B be unauthorized in f_Π , and denote $C = B \cap L$ and $D = B \cap R$. For every $U \in \binom{L}{k-t}$ with $U \not\subseteq C$, the parties of B hold only a strict subset of the additive shares of m_U . By the privacy of additive sharing (Definition A.5), the joint view of all these incomplete sharings can be treated as a part of the first, secret-independent block of Lemma 3.3. It remains to identify a common uniform field element that appears as a pad in every relevant additive sharing on the masking side, and to verify that this pad is independent of the rest of the view. This pad will be η when $|D| < t$ and ω_D when $|D| \geq t$, and in both cases Lemma 3.3 will serve as the formal framework.

Case $|D| < t$. The privacy of the threshold scheme implies that the vector Z_D of threshold shares given to the parties of D is independent of η . Since the threshold sharing is independent of all trace pads and additive-sharing randomness, we can include Z_D with the incomplete

additive sharings in the first block of Lemma 3.3. Let ρ_1 represent this first block, and let ρ_2 consist of all the trace pads and the fresh randomness used in the complete additive sharings. Then, η is independent of (ρ_1, ρ_2) . Outside the first block, the shares of B consist of complete additive shares and separately disclosed trace pads. In every complete additive sharing of $m_U = s + \eta + \sum_{T \in \mathcal{T}} \gamma_U(T) \omega_T$, each share that consists only of fresh randomness has coefficient zero on both s and η , while the last remaining share has coefficient one on both. Every separately disclosed trace pad has coefficient zero on both. Thus, every coordinate outside the first block has equal coefficients on s and η . Therefore, for suitable matrices N, M and a vector a , the view of B can be represented as

$$\text{View}_B(s) \stackrel{d}{=} (N\rho_1, a(s + \eta) + M\rho_2).$$

Applying Lemma 3.3 with the pad η proves privacy in this case.

Case $|D| \geq t$. Take the pad ω_D (note that $D \in \mathcal{T}$ since $|D| \geq t$, so ω_D is defined). It is not released to B , since disclosure-side item (2) gives ω_D only to parties in $R \setminus D$. Moreover, ω_D is uniform and independent of all the remaining dealer randomness. Let (ρ_1, ρ_2) represent the incomplete additive sharings and all remaining randomness in the view of B , excluding ω_D .

Fix $U \in \binom{C}{k-t}$. For every $V \in \binom{D}{t}$, the set $U \cup V$ is contained in the unauthorized coalition B , so $f_\Pi(U \cup V) = 0$. Hence, $\gamma_U(D) = 1$ by (4). It follows that, in every coordinate of the complete additive sharing of m_U , the coefficient of s equals the coefficient of ω_D . Since ω_D is not disclosed separately, the coefficients of ω_D and s in each pad-disclosure share are 0, and the entire view has a representation

$$\text{View}_B(s) \stackrel{d}{=} (N\rho_1, a(s + \omega_D) + M\rho_2)$$

for suitable matrices N, M and a vector a . Applying Lemma 3.3 with the pad ω_D proves privacy in this case and completes the proof. $\square$

5 Linear Secret Sharing for One Upslice

In this section, we show how to decompose an upslice into partitioned upslices and realize it by a linear secret-sharing, thus proving Theorem 1.2. We describe partition families in Section 5.1 and the decomposition in Section 5.2.

5.1 Covering Partition Families

We describe two families of partitions. The family of Lemma 5.1 is the more general one that works for every choice of the parameters q and t , and works by choosing partitions at random until all the relevant sets are covered. The family of Lemma 5.2 has the desired covering property for every choice of q , but only for a balanced choice of $t = \lceil kq/n \rceil$. This choice is considered balanced since it splits the k elements of each covered minterm between the two parts so that both parts have density k/n up to rounding. For this choice of t , this family is smaller and can be computed more efficiently, but it does not give the best bound for upslices with density below $1/2$. We prove the two lemmas.

Lemma 5.1 (Cover by arbitrary-profile partitions). *Let $n \geq 1$, $1 \leq k \leq n$, $0 \leq q \leq n - 1$, and $\max\{0, k - (n - q)\} \leq t \leq \min\{k, q\}$. Let $C(n, k; q, t)$ be the smallest size of a family $\mathcal{R} \subseteq \binom{[n]}{q}$*

such that every $A \in \binom{[n]}{k}$ satisfies $|A \cap R| = t$ for at least one $R \in \mathcal{R}$. Then

$$\frac{1}{p_{k,q,t}} \leq C(n, k; q, t) \leq \left\lceil \frac{\ln \binom{n}{k} + 1}{p_{k,q,t}} \right\rceil,$$

where $p_{k,q,t} = \binom{q}{t} \binom{n-q}{k-t} / \binom{n}{k} = \binom{k}{t} \binom{n-k}{q-t} / \binom{n}{q}$.

Proof. For a fixed k -set A , a uniformly random q -set R satisfies $|A \cap R| = t$ with probability $p_{k,q,t}$: there are $\binom{k}{t}$ ways to choose the elements of R in A , and $\binom{n-k}{q-t}$ ways to choose the remaining elements. Choose independently

$$z = \left\lceil \frac{\ln \binom{n}{k} + 1}{p_{k,q,t}} \right\rceil$$

such sets. The probability that some k -set is not covered is at most

$$\binom{n}{k} (1 - p_{k,q,t})^z \leq \binom{n}{k} e^{-p_{k,q,t} z} \leq e^{-1} < 1.$$

Thus a cover of size at most z exists. For the lower bound, each q -set covers exactly $\binom{q}{t} \binom{n-q}{k-t}$ of the $\binom{n}{k}$ possible k -sets. Hence every cover has size at least $\binom{n}{k} / \left(\binom{q}{t} \binom{n-q}{k-t} \right) = 1/p_{k,q,t}$. $\square$

Lemma 5.2 (Cover by window partitions). *Let $1 \leq k \leq n$ and $0 \leq q \leq n-1$, and set $t = \lceil kq/n \rceil$. Identify $[n]$ with $\mathbb{Z}_n$, and for every $i \in \mathbb{Z}_n$ let the i 'th partition be defined by $R_i = \{i, i+1, \dots, i+q-1\}$ (residues modulo n) and $L_i = [n] \setminus R_i$. Then, every $A \in \binom{[n]}{k}$ satisfies $|A \cap R_i| = t$ for at least one $i \in \mathbb{Z}_n$.*

Proof. The case $q = 0$ is immediate. Otherwise, fix $A \in \binom{[n]}{k}$ and let $g(i) = |A \cap R_i|$. Since $R_{i+1} = (R_i \setminus \{i\}) \cup \{i+q\}$, consecutive values of $g(i)$ differ by at most one. Each of the k elements of A belongs to exactly q windows, so $\sum_i g(i) = kq$. Hence, the average of g is kq/n , and $\min_i g(i) \leq \lfloor kq/n \rfloor \leq t \leq \max_i g(i)$. Since g is integer-valued and changes by at most one along the cycle, it attains every integer between its minimum and its maximum. In particular, it holds that $g(i) = |A \cap R_i| = t$ for some i . $\square$

The number of window partitions is therefore n , and they can be computed efficiently and deterministically.

5.2 The Upper Bounds for Upslices

We now combine a covering family with the construction of Section 4. For the edge cases, we use the formula-based scheme of Benaloh and Leichter [18], which is linear over every finite field and has total share length at most the size of the monotone DNF.

Theorem 5.3 (Layerwise upslice theorem). *For every finite field $\mathbb{F}$, let n, k, q, t be as in Lemma 5.1. Then, every (k, n) -upslice has a linear secret-sharing scheme over $\mathbb{F}$ with total share length at most*

$$\frac{1}{p_{k,q,t}} \left[\binom{n-q}{k-t} + \sum_{j=t}^q \binom{q}{j} \right] O(n^2 \log(n+1)). \quad (6)$$

Moreover, when $t = \lceil kq/n \rceil$, the same bound holds without the factor $1/p_{k,q,t}$.

Proof. Recall the notations $z = \left\lceil \frac{\ln \binom{n}{k} + 1}{p_{k,q,t}} \right\rceil$ and $m = \left\lceil \log_{|\mathbb{F}|}(q+1) \right\rceil$. Let f_k be a (k, n) -upslice, and let $\mathcal{R}$ be a cover of size at most z given by Lemma 5.1. For each $R_i \in \mathcal{R}$, set $L_i = [n] \setminus R_i$, and let f_{Π_i} be the partitioned upslice whose minterms are the minterms A of f_k satisfying $|A \cap R_i| = t$. The covering property implies $f_k = \bigvee_{R_i \in \mathcal{R}} f_{\Pi_i}$. Therefore, it suffices to share the same secret independently according to each partitioned upslice.

If $1 \leq t < k$, then minterms have parties in both the left and right sides, and so by Theorem 4.2 the total share length for each f_{Π_i} is at most $(k-t)\binom{n-q}{k-t} + qm + q \sum_{j=t}^q \binom{q}{j}$. If $t = 0$, all the minterms of f_{Π_i} lie in L_i , and its monotone DNF has size at most $k\binom{n-q}{k}$. If $t = k$, they all lie in R_i , and its monotone DNF has size at most $k\binom{q}{k}$. The DNF scheme of [18] gives the corresponding share-length bounds in these two cases.

In all three cases, the total share length for each f_{Π_i} is at most

$$\left[\binom{n-q}{k-t} + \sum_{j=t}^q \binom{q}{j} \right] O(n \log(n+1)).$$

Since $\ln \binom{n}{k} \leq \ln(2^n) \leq n \ln 2$ and $p_{k,q,t} \leq 1$, we have $z = O(n/p_{k,q,t})$. Summing over the size of the cover family gives (6), which also proves Theorem 1.2.

For the moreover part, when $t = \lceil kq/n \rceil$, we can instead use Lemma 5.2 in the same proof, replacing z by the number of window partitions n . $\square$

6 Beyond the Linear Exponent

We prove Theorems 1.5 and 1.6 by combining our bound for the linear exponent of f_{bot} functions and its dual with non-linear and quadratic midslice constructions. We postpone the technical details of the exponent analysis of f_{bot} to Section 6.3, and start with the proofs of the theorems.

6.1 Non-linear Secret Sharing

Proof of Theorem 1.5. Let δ be a constant in $(0, 1/2)$ to be fixed later, and let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a monotone function. Following [31], we realize f by decomposing it into three parts parametrized by δ . It holds that $f = f_{\text{top}} \wedge (f_{\text{mid}} \vee f_{\text{bot}})$, where:

- f_{bot} is the monotone function whose minterms are the minterms of f of size less than $(1/2 - \delta)n$. Equivalently, $f_{\text{bot}}(x) = 1$ if and only if there exists a minterm y of f such that $\text{wt}(y) < (1/2 - \delta)n$ and $y \leq x$.
- f_{top} is the monotone function whose maxterms are the maxterms of f of size greater than $(1/2 + \delta)n$. Equivalently, $f_{\text{top}}(x) = 0$ if and only if there exists a maxterm y of f such that $\text{wt}(y) > (1/2 + \delta)n$ and $x \leq y$.
- f_{mid} agrees with f on all inputs x of Hamming weight $\text{wt}(x) \in [(1/2 - \delta)n, (1/2 + \delta)n]$, outputs 0 on inputs of smaller weight, and outputs 1 on inputs of larger weight. We call such functions δ -midslice functions.

To verify the identity $f = f_{\text{top}} \wedge (f_{\text{mid}} \vee f_{\text{bot}})$, we first write $a = (1/2 - \delta)n$ and $b = (1/2 + \delta)n$, and notice that monotonicity and the definitions imply $f_{\text{bot}} \leq f \leq f_{\text{top}}$. We then start with inputs with $\text{wt}(x) < a$, in which case $f_{\text{mid}}(x) = 0$. If $f(x) = 1$, then x contains a minterm of f of weight at most $\text{wt}(x) < a$, so $f_{\text{bot}}(x) = 1$. Additionally, the inequality $f \leq f_{\text{top}}$ guarantees that $f_{\text{top}}(x)$ also equals 1, and so does the entire expression. If $f(x) = 0$, then x contains no minterm

of f , so $f_{\text{bot}}(x) = 0$, and $f_{\text{mid}} = 0$ on every input in this range. Thus, the inner OR equals $f(x)$, and so does the entire expression.

The claim for x where $\text{wt}(x) > b$ is exactly analogous. Next, if $a \leq \text{wt}(x) \leq b$, then $f_{\text{mid}}(x) = f(x)$. When $f(x) = 0$, the inequality $f_{\text{bot}} \leq f$ gives $f_{\text{bot}}(x) = 0$, so the inner OR, and hence the entire expression, is 0. When $f(x) = 1$, the inner OR is 1, and $f \leq f_{\text{top}}$ gives $f_{\text{top}}(x) = 1$, so the entire expression is 1. Next, by the closure properties of secret sharing over $\wedge$ and $\vee$, the share size of f is at most the sum of share sizes of $f_{\text{top}}, f_{\text{mid}}, f_{\text{bot}}$, so it remains to realize these three access structures.

Realizing $f_{\text{bot}}, f_{\text{mid}}, f_{\text{top}}$. Set $\alpha = 1/2 - \delta$ and $k_0 = \lceil \alpha n \rceil - 1$, and choose $0 < \rho < \alpha$ with $\tau_{R,\alpha} < 1/2$ as in Lemma 6.5. Every minterm of f_{bot} has size at most $k_0 \leq \lfloor \alpha n \rfloor$, so that lemma realizes f_{bot} over $\mathbb{F}_2$ with total share size $2^{B_\rho(\alpha)n + O(\log n)}$.

We next realize f_{top} by applying the construction used for f_{bot} to its dual $f_{\text{top}}^*(x) = 1 - f_{\text{top}}(\bar{x})$. The minterms of f_{top}^* are the complements of the maxterms of f_{top} , so their sizes are less than αn and hence at most k_0 . Thus, Lemma 6.5 gives a linear scheme for f_{top}^* over $\mathbb{F}_2$ with total share size $2^{B_\rho(\alpha)n + O(\log n)}$, just as for f_{bot} . Since duality preserves linearity, the field, and total share size (e.g., [24]), there exists a scheme that realizes f_{top} with the same bound.

Finally, to realize f_{mid} , we use the following construction for δ -midslices from [4].

Theorem 6.1 (Midslice schemes [4, Lem. 5.10]). *For every constant $\delta \in (0, 1/2)$, every δ -midslice access structure over n parties can be realized by a secret-sharing scheme for one-bit secrets with total share size $2^{M(\delta)n + o(n)}$, where*

$$M(\delta) = \begin{cases} \left(\frac{1}{2} + \delta\right) H_2\left(\frac{1/2 - \delta}{1/2 + \delta}\right), & 0 < \delta \leq \frac{1}{6}, \\ \frac{1}{2} + \delta, & \frac{1}{6} < \delta < \frac{1}{2}. \end{cases}$$

Applying Theorem 6.1 to f_{mid} gives total share size $2^{M(\delta)n + o(n)}$. The overall exponent is therefore at most

$$\max\{B_\rho(1/2 - \delta), M(\delta)\}.$$

Taking $\delta = 0.08026$ and $\rho = 0.39368$ gives

$$B_\rho(1/2 - \delta) < 0.493727, \quad M(\delta) < 0.493700.$$

Both are smaller than 0.494, proving the theorem. $\square$

6.2 Quadratic Secret Sharing

Proof of Theorem 1.6. Following the steps of the proof of Theorem 1.5, it suffices to replace the non-linear δ -midslice upper bound of Theorem 6.1 with a quadratic one. Indeed, the schemes used there for f_{bot} and f_{top} are linear and the standard compositions over OR and AND preserve linearity, and so all steps other than the δ -midslice construction have sharing and reconstruction degree at most one.

We use the following quadratic midslice bound, obtained by combining the quadratic RCDS protocol of [16, Thm. 6.5] with the RCDS-to-midslice compiler of [4, Lemmas 5.5 and 5.8]; see Section B.2 for its derivation.

Theorem 6.2 (Quadratic midslice schemes). *For every constant $\delta \in (0, 1/2)$, every δ -midslice function over n parties can be realized by a secret-sharing scheme for one-bit secrets with quadratic sharing and reconstruction and total share size $2^{\mathcal{Q}(\delta)n+o(n)}$, where*

$$\mathcal{Q}(\delta) = \begin{cases} \frac{1}{3} + \frac{1+2\delta}{3} H_2\left(\frac{1-2\delta}{1+2\delta}\right), & 0 < \delta \leq \frac{1}{6}, \\ \frac{2}{3} + \frac{2\delta}{3}, & \frac{1}{6} \leq \delta < \frac{1}{2}. \end{cases}$$

Similarly to the proof of Theorem 1.5, the exponent for general functions is at most $\max\{B_\rho(1/2 - \delta), \mathcal{Q}(\delta)\}$, for any balanced choice satisfying the conditions of Lemma 6.5 with $\alpha = 1/2 - \delta$. Taking $\delta = 0.02664$ and $\rho = 0.46447$ gives

$$B_\rho(1/2 - \delta) < 0.499317, \quad \mathcal{Q}(\delta) < 0.499293.$$

Both are smaller than 0.4994, proving the theorem. $\square$

6.3 The Exponents of Upslices and f_{bot} Functions

In this subsection, we prove upper bounds on the share size of access structures whose minterms all lie below a fixed density cutoff smaller than $1/2$, directly realizing the bottom part of the decomposition in Section 6.1. In Section 6.3.1, we prove an upper bound in binomial terms. In Section 6.3.2, we express the exponential rates of this bound and the upslice bounds for constant minterm densities, proving Corollary 1.4 and deriving the exponent bound $B_\rho(\alpha)$ used in Sections 6.1 and 6.2.

6.3.1 Composing Upslices

Lemma 6.3 (Linear scheme for f_{bot} functions). *Let n, k, q, t be as in Lemma 5.1, with $k - t \leq (n - q)/2$ and $t \leq q/2$. Over every finite field $\mathbb{F}$, every n -party access structure whose minterms have size at most k has a linear scheme with total share length at most*

$$\frac{1}{p_{k,q,t}} \left[\binom{n-q}{k-t} + \sum_{h=t}^q \binom{q}{h} \right] O(n^4 \log(n+1)). \quad (7)$$

Proof. We compare the three components of the bound in Theorem 5.3 as the minterm size decreases from k to $j \leq k$. We show that the left-side binomial coefficient does not increase, the right-side binomial tail increases by at most a factor of 2, and the inverse covering probability increases by at most a factor of $n+1$. These losses, together with taking the OR of all upslices below k affect the share size only by a polynomial factor.

The inverse covering probability. Fix $1 \leq j \leq k$, a j -set M , and a k -set $A \supseteq M$. For a uniformly random q -set R , the event $|A \cap R| = t$ has probability $p_{k,q,t}$ and implies that $s = |M \cap R|$ satisfies $0 \leq s \leq t$ and $0 \leq j - s \leq k - t$. Hence

$$\sum_{\substack{0 \leq s \leq t \\ 0 \leq j-s \leq k-t}} p_{j,q,s} \geq p_{k,q,t}.$$

There are at most $n+1$ summands, so some such $s = s_j$ satisfies $p_{j,q,s_j} \geq p_{k,q,t}/(n+1)$. This profile is feasible for Theorem 5.3.

The left-side binomial coefficient. Since $j - s_j \leq k - t \leq (n - q)/2$, we have $\binom{n-q}{j-s_j} \leq \binom{n-q}{k-t}$.

The right-side binomial tail. Since $s_j \leq t \leq q/2$, we have $\sum_{h=s_j}^q \binom{q}{h} \leq 2^q \leq 2 \sum_{h=t}^q \binom{q}{h}$.

Applying Theorem 5.3 with parameters (n, j, q, s_j) therefore gives a linear scheme for f_j , the j -upslice of f , with total share length

$$\frac{1}{p_{k,q,t}} \left[\binom{n-q}{k-t} + \sum_{h=t}^q \binom{q}{h} \right] O(n^3 \log(n+1)).$$

Composing the at most $k \leq n$ schemes using OR proves (7). □

The polynomial factor in (7) can be improved, for example by adapting the linear construction in advance to handle f_{bot} access structures directly. However, we do not pursue such optimizations here, since our focus is on the share size exponential rate, which is used in the schemes for general access structures in Sections 6.1 and 6.2.

6.3.2 Upslices with Constant Density and f_{bot} Functions

In this part, we derive estimates of the upslice upper bounds of Theorem 5.3 for upslices with constant densities, i.e., for $\lceil \alpha n \rceil$ -upslices with constant α . This proves Corollary 1.4.

We fix some notation. The constants $\alpha, \lambda, \tau, \tau_L, \tau_R \in (0, 1)$ will represent the upslice density k/n , the fraction q/n of parties on the disclosure side, the minterm's intersection size with the disclosure side normalized by n , namely t/n , and the minterm densities $(k-t)/(n-q)$ and t/q on the masking and disclosure sides, respectively. We then require that $\tau < \min\{\alpha, \lambda\}$, analogous to $t < k$ and $t < q$, and that $\max\{0, \alpha + \lambda - 1\} < \tau$, analogous to $t > 0$ and $k - t < n - q$. We note the two side densities satisfy

$$\tau_L = \frac{\alpha - \tau}{1 - \lambda}, \quad \tau_R = \frac{\tau}{\lambda}, \quad \alpha = (1 - \lambda)\tau_L + \lambda\tau_R.$$

We also define

$$I(\alpha; \lambda, \tau) = H_2(\alpha) - (1 - \lambda)H_2(\tau_L) - \lambda H_2(\tau_R)$$

to represent the exponential cost of $1/p_{k,q,t}$ and note that by concavity of the binary entropy function and since $\alpha = (1 - \lambda)\tau_L + \lambda\tau_R$ it holds that $I(\alpha; \lambda, \tau) \geq 0$. Define $\kappa(\tau_R) = 1$ for $0 \leq \tau_R \leq 1/2$ and $\kappa(\tau_R) = H_2(\tau_R)$ for $1/2 \leq \tau_R \leq 1$, with the usual convention $H_2(0) = H_2(1) = 0$. We now state the relevant upslice upper bound.

Corollary 6.4 (The exponent of upslice bounds for constant density). *For every $\alpha \in (0, 1)$, the linear exponent of $\lceil \alpha n \rceil$ -upslices satisfies $U_\ell(\alpha) \leq u(\alpha)$, where*

$$u(\alpha) = \inf_{\lambda, \tau} [I(\alpha; \lambda, \tau) + \max\{(1 - \lambda)H_2(\tau_L), \lambda\kappa(\tau_R)\}], \quad (8)$$

and the infimum is over constants λ, τ subject to the constraints above. Moreover,

$$u(\alpha) \leq \frac{H_2(\alpha)}{1 + H_2(\alpha)} \quad \text{for } 0 < \alpha < \frac{1}{2}, \quad u(\alpha) \leq \frac{H_2(\alpha)}{2} \quad \text{for } \frac{1}{2} \leq \alpha < 1.$$

Proof. By Theorem 5.3, every (k, n) -upslice has a linear secret-sharing scheme over any finite field $\mathbb{F}$ with total share length at most

$$\frac{1}{p_{k,q,t}} \left[\binom{n-q}{k-t} + \sum_{j=t}^q \binom{q}{j} \right] O(n^2 \log(n+1)).$$

Take $k = \alpha n + O(1)$, $q = \lambda n + O(1)$, and $t = \tau n + O(1)$. By Claim B.3 that handles rounding and estimation of the exponents, the covering factor has exponent $I(\alpha; \lambda, \tau)$, while the two per-partition terms have exponents $(1 - \lambda)H_2(\tau_L)$ and $\lambda\kappa(\tau_R)$. Taking the infimum proves (8) as an upper bound on $U_\ell(\alpha)$.

For the concrete upper bounds of the corollary, take $\tau_L = \tau_R = \alpha$. This represents a balanced profile for the minters, where they have the same density in the left and right sides. In this case, $I(\alpha; \lambda, \tau) = 0$, since the two subtracted entropy terms sum to $(1 - \lambda)H_2(\alpha) + \lambda H_2(\alpha) = H_2(\alpha)$. Writing $h = H_2(\alpha)$, the remaining exponent is $\max\{(1 - \lambda)h, \lambda\kappa(\alpha)\}$, which is minimized at $\lambda = h/(h + \kappa(\alpha))$. This gives $h/(1 + h)$ below the middle layer and $h/2$ above it. $\square$

Realizing f_{bot} functions with a balanced parameter choice. For feasible constants α, λ, τ with $\tau_L, \tau_R < 1/2$, the bounds for a single upslice in Theorem 5.3 and for bounded minters in Lemma 6.3 differ only by polynomial factors. Hence realizing all minterm layers up to density α has the same exponential upper bound as realizing a single α -upslice with these parameters. In particular, over every fixed finite field $\mathbb{F}$, every access structure whose minters have size at most $\lfloor \alpha n \rfloor$ has a linear scheme with total share length $2^{E(\alpha; \lambda, \tau)n + O(\log n)}$, where

$$E(\alpha; \lambda, \tau) = I(\alpha; \lambda, \tau) + \max\{(1 - \lambda)H_2(\tau_L), \lambda\kappa(\tau_R)\}. \quad (9)$$

For the numerical derivation of the general upper bound in Section 6, we optimize this expression for every α . We first restrict to parameter choices that balance the two terms inside the maximum in $E(\alpha; \lambda, \tau)$, representing the masking and disclosure costs, which leaves a one-dimensional family to optimize over.

Lemma 6.5 (The exponent of f_{bot} functions). *Fix constants $0 < \rho < \alpha < 1/2$ and set $\lambda = \lambda(\rho)$ and $\tau_{R,\alpha} = \tau_{R,\alpha}(\rho)$ as follows:*

$$\lambda = \frac{H_2(\rho)}{1 + H_2(\rho)}, \quad \tau_{R,\alpha} = \frac{\alpha - (1 - \lambda)\rho}{\lambda}. \quad (10)$$

Suppose that $\tau_{R,\alpha} < 1/2$. Over every fixed finite field $\mathbb{F}$, every n -party access structure whose minters have size at most $\lfloor \alpha n \rfloor$ has a linear scheme with total share length $2^{B_\rho(\alpha)n + O(\log n)}$, where

$$B_\rho(\alpha) := H_2(\alpha) - \lambda H_2(\tau_{R,\alpha}).$$

Proof. The assumption $\tau_{R,\alpha} < 1/2$ can always be satisfied by choosing $\rho < \alpha$ sufficiently close to α . Indeed,

$$\tau_{R,\alpha} = \rho + \frac{\alpha - \rho}{\lambda(\rho)} \longrightarrow \alpha < \frac{1}{2} \quad \text{as } \rho \nearrow \alpha.$$

Since the weighted average of the two side densities must equal α , this choice of $\tau_{R,\alpha}$ forces

$$\tau_L = \frac{\alpha - \lambda\tau_{R,\alpha}}{1 - \lambda} = \frac{\alpha - [\alpha - (1 - \lambda)\rho]}{1 - \lambda} = \rho.$$

Moreover, $\tau_{R,\alpha} = \rho + (\alpha - \rho)/\lambda > \rho > 0$. Together with the assumption $\tau_{R,\alpha} < 1/2$ and the identity $\tau_L = \rho \in (0, 1/2)$, this shows that both side densities lie in $(0, 1/2)$, as required to apply (9) in the regime covered by the f_{bot} lemma, Lemma 6.3. The masking-side exponent is therefore $(1 - \lambda)H_2(\rho)$. On the disclosure side, $\tau_{R,\alpha} < 1/2$ implies $\kappa(\tau_{R,\alpha}) = 1$: the binomial tail $\sum_{h=t}^q \binom{q}{h}$ contains at least half of all subsets of R , and hence has exponent $\lambda\kappa(\tau_{R,\alpha}) = \lambda$, independently of the precise density in this range. Our choice of λ makes these two costs equal:

$$(1 - \lambda)H_2(\rho) = \frac{H_2(\rho)}{1 + H_2(\rho)} = \lambda.$$

Consequently, the maximum in (9) equals λ , which gives

$$\begin{aligned} B_\rho(\alpha) &= E(\alpha; \lambda, \lambda\tau_{R,\alpha}) = I(\alpha; \lambda, \lambda\tau_{R,\alpha}) + \lambda \\ &= H_2(\alpha) - (1 - \lambda)H_2(\rho) - \lambda H_2(\tau_{R,\alpha}) + \lambda = H_2(\alpha) - \lambda H_2(\tau_{R,\alpha}). \end{aligned} \tag{11}$$

Thus every access structure whose minterms have size at most $\lfloor \alpha n \rfloor$ has total share length $2^{B_\rho(\alpha)n + O(\log n)}$. $\square$

7 A Lower Bound over All Fields

In this section, we derive the lower bound on the size of monotone span programs that holds simultaneously over all fields, proving Theorem 1.8.

Theorem 7.1 (Monotone span programs over all fields). *For every constant $\varepsilon > 0$, almost all n -party access structures require monotone span programs of size at least*

$$\frac{2^{n/2}}{n^{3/4+\varepsilon}} = 2^{n/2 - o(n)}$$

over every field. This remains true even if the field is chosen separately for each access structure and may be finite or infinite.

Here, “almost all” means that the fraction of access structures satisfying the statement tends to one as $n \rightarrow \infty$.

In Section 7.1, we describe definitions and properties of monotone span programs, in Section 7.2 we state Nelson’s zero-pattern bound, and the rest of the subsections are dedicated to the proof.

7.1 Monotone Span Programs and Field Extensions

We start with standard definitions of monotone span programs:

Definition 7.2 (Monotone Span Program [30]). A monotone span program is a quadruple $\mathcal{M} = (\mathbb{F}, M, \rho, v)$, where $\mathbb{F}$ is a field, M is a $d \times e$ matrix over $\mathbb{F}$, $\rho : [d] \rightarrow [n]$ labels each row of M by a party, and $v \in \mathbb{F}^e \setminus \{0\}$ is the target vector. The size of $\mathcal{M}$ is the number d of rows of M . For every coalition $A \subseteq [n]$, let M_A denote the submatrix consisting of the rows of M whose labels lie in A . We say that $\mathcal{M}$ accepts A if the rows of M_A span v . We say that $\mathcal{M}$ computes an access structure f if it accepts A exactly when $f(A) = 1$. As shown below, we may assume that $M \in \mathbb{F}^{d \times d}$ and $v = e_1 = (1, 0, \dots, 0)$.

Theorem 7.3 (Linear SSS and monotone span programs [30]). *A monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be realized by a linear secret-sharing scheme over a finite field $\mathbb{F}$ with total share length d iff it can be computed by a monotone span program of size d over $\mathbb{F}$.*

We can always restrict the matrix of a span program to a set of linearly independent columns without changing the function that is computed by the program, which implies the following well-known fact:

Fact 7.4. *For every MSP $\mathcal{M} = (\mathbb{F}, M, \rho, v)$ with $M \in \mathbb{F}^{d \times e}$, there exists an MSP $\mathcal{M}' = (\mathbb{F}, M', \rho, e_1)$ computing the same function, where $M' \in \mathbb{F}^{d \times d}$ and $e_1 = (1, 0, \dots, 0) \in \mathbb{F}^d$. Thus, the matrix can be taken to be square and the target vector to be e_1 , without changing the field, row labeling, or size of the program.*

The arguments for the number of columns can be found, e.g., in [9], and the freedom in choosing the target vector is implied by a change of basis. Since $v \neq 0$, if we choose an invertible matrix $T \in \mathbb{F}^{e \times e}$ such that $vT = e_1$, then replacing M by MT preserves all linear dependencies between the target and the rows, while changing the target to e_1 .

We observe that the spanning conditions for authorized and unauthorized coalitions can be specified in terms of ranks of sub-matrices of M . For every coalition A ,

$$f(A) = 1 \iff \text{rank}_{\mathbb{F}}(M_A) = \text{rank}_{\mathbb{F}} \begin{pmatrix} M_A \\ e_1 \end{pmatrix}. \quad (12)$$

Before starting the lower bound proof, we note that for a given MSP over a field $\mathbb{F}$, the rank constraints of (12) that determine the access structure do not change when we view the program as a matrix over an extension field of $\mathbb{F}$.

Lemma 7.5 (Preservation of rank). *Let $E/\mathbb{F}$ be a field extension, and let Y be a matrix over $\mathbb{F}$. Then*

$$\text{rank}_{\mathbb{F}}(Y) = \text{rank}_E(Y).$$

In particular, viewing the coefficients of a monotone span program over $\mathbb{F}$ as elements of E preserves its access structure.

Proof. The rank is the largest order of a nonzero square minor. Each minor belongs to $\mathbb{F}$, and the embedding $\mathbb{F} \hookrightarrow E$ is injective, so the minor is zero over $\mathbb{F}$ if and only if it is zero over E . The last assertion follows by applying this equality to both matrices in (12). $\square$

The reason for allowing an extension is that a nonzero polynomial can vanish at every point of a small finite field. For example, $x^2 - x$ vanishes on all of $\mathbb{F}_2$. If $\mathbb{F}$ is infinite, a nonzero polynomial P of total degree D has a nonzero evaluation over $\mathbb{F}$, by applying the Schwartz–Zippel lemma to any finite subset of $\mathbb{F}$ of size greater than D . If $\mathbb{F}$ is finite, we may instead pass to a finite extension $E/\mathbb{F}$ with $|E| > D$.

7.2 Lower Bound for Zero Patterns over All Fields

Our lower bound is based on the following upper bound on the number of zero patterns of polynomials over all fields. Let $f_1, \dots, f_N \in \mathbb{Z}[x_1, \dots, x_m]$. We say that $S \subseteq [N]$ is a *realizable zero pattern* if there are a field K and a point $u \in K^m$ such that

$$S = \{i \in [N] : f_i(u) \neq 0\},$$

where integer coefficients are interpreted through the natural homomorphism $\mathbb{Z} \rightarrow K$. The field and its characteristic are allowed to vary between patterns. We use the following theorem.

Theorem 7.6 (Nelson [32], Theorem 2.1). *Let $f_1, \dots, f_N \in \mathbb{Z}[x_1, \dots, x_m]$ have degree at most D and coefficients of absolute value at most c , where $c, D \geq 1$ are integers. Set $H = \binom{ND+m}{m}$. If a positive integer B satisfies*

$$B > H(\log(3B) + N \log(c(eN)^D)), \quad (13)$$

then the number of realizable zero patterns is at most B .

We next derive the parameter bound used in our proof.

Corollary 7.7. *Suppose $2 \leq d \leq 2^n$, and consider integer polynomials with*

$$m = O(d^2), \quad D = O(d), \quad N \leq 2^n d^{O(1)}, \quad \log c = O(d \log d).$$

The number of their realizable zero patterns, over all fields, is at most $2^{O(d^2 n)}$.

Proof. Since $\log d \leq n$, we have $\log N = O(n)$. The elementary bound $\binom{a}{b} \leq a^b$ gives

$$\log H \leq m \log(ND + m) = O(d^2 n).$$

Also,

$$\log(c(eN)^D) = \log c + D \log(eN) = O(dn).$$

For $B = \lceil 2^{Cd^2 n} \rceil$, the logarithm of the parenthesized factor on the right-hand side of (13) is

$$\log(\log(3B) + N \log(c(eN)^D)) = O(n + \log(C + 1)).$$

Consequently, for a sufficiently large constant C , the logarithm of the entire right-hand side is strictly less than $Cd^2 n$. This choice of B satisfies (13), and the result follows from Theorem 7.6. $\square$

7.3 One Pair of Matrices for All Coalition Ranks

First, we prove that given t matrices $Y_1, \dots, Y_t \in \mathbb{F}^{(d+1) \times (d+1)}$, one can choose a single pair of matrices L and R over a finite extension (that is, a finite-degree extension) of $\mathbb{F}$ such that each Y_j has rank at least k if and only if the upper-left $k \times k$ submatrix of $LY_j R$ has nonzero determinant. Formally, we define

$$\Delta_k(Y_j; L, R) := \det((LY_j R)_{[k],[k]}), \quad j \in [t], \quad 1 \leq k \leq d+1,$$

where $(LY_j R)_{[k],[k]}$ denotes the submatrix formed by the first k rows and the first k columns of $LY_j R$, and prove the following lemma.

Lemma 7.8 (Simultaneous leading minors). *Let $Y_1, \dots, Y_t \in \mathbb{F}^{(d+1) \times (d+1)}$, where $d \geq 1$ and $\mathbb{F}$ is a field. There is a finite extension $E/\mathbb{F}$ and a pair of matrices $L, R \in E^{(d+1) \times (d+1)}$ such that, for every $j \in [t]$ and every $1 \leq k \leq d+1$,*

$$\Delta_k(Y_j; L, R) \neq 0 \iff k \leq \text{rank}_{\mathbb{F}}(Y_j).$$

The fact that the same pair L, R works for all the matrices is important for the zero-pattern argument below. It allows us to use the entries of a single pair as common variables, keeping their number $O(d^2)$ rather than introducing a separate pair for every matrix.

For a single matrix A over any field K , Kaltofen and Saunders [29] show that for unit-triangular Toeplitz matrices L, R with entries drawn uniformly from $S \subseteq K$, all leading principal minors of LAR of order at most $\text{rank}(A)$ are nonzero with probability at least $1 - \text{rank}(A)(\text{rank}(A) + 1)/|S|$. Choose a finite extension $E/\mathbb{F}$ containing a finite subset S with $|S| > t(d + 1)(d + 2)$; when $\mathbb{F}$ is infinite we may take $E = \mathbb{F}$, and when $\mathbb{F}$ is finite we take a sufficiently large finite extension. Regarding $Y_1, \dots, Y_t$ as matrices over E (rank is unchanged by extension), applying their theorem with entries drawn from S and taking a union bound over the t matrices shows that one pair works for all of them. We include a self-contained proof, not relying on [29], in Section B.3.

7.4 Counting Access Structures via Zero Patterns

We combine the simultaneous leading-minor lemma with the zero-pattern bound to count access structures computed by MSPs with at most d rows, over all fields.

Lemma 7.9 (Counting access structures computed by small MSPs). *For every $n \geq 1$, $2 \leq d \leq 2^n$, at most $2^{O(d^2 n)}$ distinct n -party access structures are computed by monotone span programs with at most d rows over some field.*

Proof. Fix n, d as in the statement. By Fact 7.4 and padding with zero rows and columns, we may take each program matrix to be $d \times d$ and its target to be e_1 , without changing its field or access structure. The added rows can be labeled arbitrarily. We first fix a row labeling $\rho : [d] \rightarrow [n]$ and count the access structures computed by programs with this labeling.

We start with a high-level description of the proof. We will define a family of polynomials of size N with integer coefficients, depending only on n, d, ρ , whose variables are the entries of three matrices X, L, R . For each MSP matrix M over any field $\mathbb{F}$ and a labeling ρ , we will assign $X = M$ and choose values for L, R over a suitable finite extension of $\mathbb{F}$. This assigns all variables, so each polynomial evaluates to a field element. This then results in a zero pattern of polynomials that correspond to M , which will allow us to obtain our lower bound via the upper bound on zero patterns of Corollary 7.7. We now move on to the proof.

Defining the polynomial family. Let $X = (x_{ij})_{i,j \in [d]}$ be a matrix of independent formal variables. For every coalition $A \subseteq [n]$, let $Z_A(X)$ be obtained by keeping the rows of X whose labels lie in A and replacing all other rows by zero. Define the following 2^{n+1} matrices of dimension $(d + 1) \times (d + 1)$:

$$Y_A(X) = \begin{pmatrix} Z_A(X) & 0 \\ 0 & 0 \end{pmatrix}, \quad Y_A^+(X) = \begin{pmatrix} Z_A(X) & 0 \\ e_1 & 0 \end{pmatrix}.$$

The upper-left block has d rows and d columns. The last row is zero in $Y_A(X)$ and contains the target e_1 in $Y_A^+(X)$; the last column is zero in both matrices.

Introduce two additional $(d + 1) \times (d + 1)$ matrices L, R whose entries are new independent formal variables. For every $A \subseteq [n]$ and $1 \leq k \leq d + 1$, define

$$p_{A,k}(X, L, R) = \det((LY_A(X)R)_{[k],[k]}), \\ q_{A,k}(X, L, R) = \det((LY_A^+(X)R)_{[k],[k]}).$$

All polynomials use the same variables: the entries of X, L, R . A realizable zero pattern records which of them are nonzero at one assignment to these variables over some field.

Encoding an MSP matrix as a zero pattern. Let $M \in \mathbb{F}^{d \times d}$ be an MSP matrix with labeling ρ over any field $\mathbb{F}$, and assign $x_{ij} = M_{ij}$ for every i, j in the polynomial family above. For every coalition A , padding with zero rows and columns gives $\text{rank } Y_A(M) = \text{rank } M_A$ and $\text{rank } Y_A^+(M) = \text{rank} \begin{pmatrix} M_A \\ e_1 \end{pmatrix}$. We then apply Lemma 7.8 to the entire family

$$\{Y_A(M), Y_A^+(M) : A \subseteq [n]\},$$

by which there exist a finite extension $E/\mathbb{F}$ and one pair of matrices L, R in E such that, simultaneously for every A and k ,

$$\begin{aligned} p_{A,k}(M, L, R) \neq 0 &\iff k \leq \text{rank}_{\mathbb{F}} M_A, \\ q_{A,k}(M, L, R) \neq 0 &\iff k \leq \text{rank}_{\mathbb{F}} \begin{pmatrix} M_A \\ e_1 \end{pmatrix}. \end{aligned}$$

Here, Lemma 7.5 ensures that the ranks are unchanged by passing to E . Thus, (M, L, R) realizes a zero pattern of the fixed polynomial family that records both ranks for every coalition.

Given a pattern obtained this way, let a_A be the largest k for which $p_{A,k}$ is nonzero, and let b_A be the largest k for which $q_{A,k}$ is nonzero. Use zero if no such index exists. These are the two ranks above, so by (12), the access structure is determined by

$$f(A) = 1 \iff a_A = b_A.$$

Thus, each access structure computed by a program with labeling ρ can be represented by a realizable zero pattern of these polynomials, and distinct access structures must have distinct representing patterns. It remains to count the realizable zero patterns of this polynomial family.

Counting the zero patterns. The number of polynomials and the number of variables are, respectively,

$$N = 2^{n+1}(d+1), \quad m = d^2 + 2(d+1)^2.$$

Every entry of either product matrix has degree at most three, so the determinants have degree at most $D = 3(d+1)$.

To bound the coefficients, let Y be either $Y_A(X)$ or $Y_A^+(X)$. Since each entry of Y is zero, one, or a variable, each entry of LYR is a sum of at most $(d+1)^2$ monomials with coefficient one. Expanding a $k \times k$ determinant therefore gives at most $k!(d+1)^{2k}$ signed monomial terms. After collecting equal monomials, each coefficient has absolute value at most this number. Since $k \leq d+1$, we may take

$$c = (d+1)!(d+1)^{2(d+1)}, \quad \log c = O(d \log d).$$

Consequently, the family satisfies

$$m = O(d^2), \quad D = O(d), \quad N \leq 2^n d^{O(1)}, \quad \log c = O(d \log d).$$

Applying Corollary 7.7 gives at most $2^{O(d^2 n)}$ realizable zero patterns over all fields, and hence at most this many access structures for the fixed labeling over all fields. Summing over the n^d possible row labelings gives at most $n^d 2^{O(d^2 n)} = 2^{O(d^2 n)}$ access structures, proving the lemma. $\square$

7.5 Completing the Lower Bound

We now compare the count in Lemma 7.9 with the number of access structures, proving Theorem 7.1.

Proof of Theorem 7.1. Let $k = \lfloor n/2 \rfloor$. Every family $\mathcal{A} \subseteq \binom{[n]}{k}$ determines an access structure by taking its upward closure. Hence, the total number of access structures is at least $2^{\binom{n}{k}}$, where $\binom{n}{k} = \Theta\left(\frac{2^n}{\sqrt{n}}\right)$. We fix $\varepsilon > 0$ and define $d_n = \left\lfloor \frac{2^{n/2}}{n^{3/4+\varepsilon}} \right\rfloor$. For sufficiently large n , we have $2 \leq d_n \leq 2^n$. By Lemma 7.9, at most $2^{Cd_n^2 n}$ access structures can be computed by a monotone span program of size at most d_n over some field. Since

$$d_n^2 n \leq \frac{2^n}{n^{1/2+2\varepsilon}} = o\left(\frac{2^n}{\sqrt{n}}\right),$$

the fraction of such access structures is at most $2^{Cd_n^2 n - \binom{n}{k}} = 2^{-\Omega(2^n/\sqrt{n})} \rightarrow 0$. Outside this family, every monotone span program over every field has size greater than d_n , and hence at least $2^{n/2}/n^{3/4+\varepsilon}$. This proves the monotone span-program lower bound. The consequence for linear secret-sharing schemes over finite fields follows from Theorem 7.3. $\square$

Acknowledgments

We thank Amos Beimel for useful comments, and Benny Applebaum for suggesting the title of the paper. This work was supported by ISF grant 1614/25.

References

- [1] Bar Alon, Amos Beimel, and Or Lasri. Simplified PIR and CDS protocols and improved linear secret-sharing schemes. In Benny Applebaum and Huijia (Rachel) Lin, editors, *Theory of Cryptography - TCC 2025, Proceedings, Part II*, volume 16269 of *Lecture Notes in Computer Science*, pages 365–398. Springer, 2025.
- [2] Benny Applebaum and Barak Arkis. On the power of amortization in secret sharing: d -uniform secret sharing and CDS with constant information rate. In Amos Beimel and S. Dziembowski, editors, *TCC 2018*, volume 11239 of *LNCS*, pages 317–344, 2018.
- [3] Benny Applebaum, Amos Beimel, Oriol Farràs, Oded Nir, and Naty Peter. Secret-sharing schemes for general and uniform access structures. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 441–471. Springer, 2019.
- [4] Benny Applebaum, Amos Beimel, Oded Nir, and Naty Peter. Better secret sharing via robust conditional disclosure of secrets. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, pages 280–293. ACM, 2020.
- [5] Benny Applebaum, Amos Beimel, Oded Nir, Naty Peter, and Toniann Pitassi. Secret sharing, slice formulas, and monotone real circuits. *ACM Transactions on Computation Theory*, 18(2):1–32, 2026.

- [6] Benny Applebaum and Oded Nir. Upslices, downslices, and secret-sharing with complexity of 1.5^n . In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 627–655. Springer, 2021.
- [7] Nuttapong Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014, Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 557–577. Springer, 2014.
- [8] László Babai, Anna Gál, and Avi Wigderson. Superpolynomial lower bounds for monotone span programs. *Comb.*, 19(3):301–319, 1999.
- [9] Amos Beimel. *Secure Schemes for Secret Sharing and Key Distribution*. PhD thesis, Technion, 1996.
- [10] Amos Beimel. Secret-sharing schemes: A survey. In Yeow Meng Chee, Zhenbo Guo, San Ling, Fengjing Shao, Yuansheng Tang, Huaxiong Wang, and Chaoping Xing, editors, *Coding and Cryptology - Third International Workshop, IWCC 2011*, volume 6639 of *Lecture Notes in Computer Science*, pages 11–46. Springer, 2011.
- [11] Amos Beimel. Secret-sharing schemes for general access structures: An introduction. *IACR Cryptol. ePrint Arch.*, 2025:518, 2025.
- [12] Amos Beimel and Oriol Farràs. The share size of secret-sharing schemes for almost all access structures and graphs. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020*, volume 12552 of *Lecture Notes in Computer Science*, pages 499–529. Springer, 2020.
- [13] Amos Beimel, Oriol Farràs, and Or Lasri. Improved polynomial secret-sharing schemes. In Guy N. Rothblum and Hoeteck Wee, editors, *Theory of Cryptography - 21st International Conference, TCC 2023, Part II*, volume 14370 of *Lecture Notes in Computer Science*, pages 374–405. Springer, 2023.
- [14] Amos Beimel, Oriol Farràs, Or Lasri, and Oded Nir. Secret-sharing schemes for high slices. In Elette Boyle and Mohammad Mahmoody, editors, *Theory of Cryptography - 22nd International Conference, TCC 2024, Part IV*, volume 15367 of *Lecture Notes in Computer Science*, pages 581–613. Springer, 2024.
- [15] Amos Beimel, Oriol Farràs, Yuval Mintz, and Naty Peter. Linear secret-sharing schemes for forbidden graph access structures. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Part II*, volume 10678 of *Lecture Notes in Computer Science*, pages 394–423. Springer, 2017.
- [16] Amos Beimel, Hussien Othman, and Naty Peter. Quadratic secret sharing and conditional disclosure of secrets. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 748–778. Springer, 2021.

- [17] Amos Beimel and Naty Peter. Optimal linear multiparty conditional disclosure of secrets protocols. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology - ASIACRYPT 2018 Proceedings, Part III*, volume 11274 of *Lecture Notes in Computer Science*, pages 332–362. Springer, 2018.
- [18] Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In Shafi Goldwasser, editor, *Advances in Cryptology - CRYPTO '88*, volume 403 of *Lecture Notes in Computer Science*, pages 27–35. Springer, 1988.
- [19] George R. Blakley. Safeguarding cryptographic keys. In Richard E. Merwin, Jacqueline T. Zanca, and Merlin Smith, editors, *Proceedings of the 1979 AFIPS National Computer Conference*, volume 48 of *AFIPS Conference Proceedings*, pages 313–317. AFIPS Press, 1979.
- [20] Ronald Cramer, Ivan Damgård, and Ueli M. Maurer. General secure multi-party computation from any linear secret-sharing scheme. In Bart Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*, pages 316–334. Springer, 2000.
- [21] Ronald Cramer and Serge Fehr. Optimal black-box secret sharing over arbitrary abelian groups. In *Annual International Cryptology Conference*, pages 272–287. Springer, 2002.
- [22] László Csirmaz. The size of a share must be large. *J. Cryptol.*, 10(4):223–231, 1997.
- [23] Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In Gilles Brassard, editor, *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 307–315. Springer, 1989.
- [24] Anna Gál. *Combinatorial methods in Boolean function complexity*. PhD thesis, University of Chicago, 1996.
- [25] Romain Gay, Iordanis Kerenidis, and Hoeteck Wee. Communication complexity of conditional disclosure of secrets and attribute-based encryption. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015, Proceedings, Part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 485–502. Springer, 2015.
- [26] Yael Gertner, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. Protecting data privacy in private information retrieval schemes. *J. Comput. Syst. Sci.*, 60(3):592–629, 2000.
- [27] Charles H Goldberg and Douglas B West. Bisection of circle colorings. *SIAM Journal on Algebraic Discrete Methods*, 6(1):93–106, 1985.
- [28] Mitsuru Ito, Akira Saito, and Takao Nishizeki. Secret sharing schemes realizing general access structure. In *Globecom 87*, pages 99–102. IEEE, 1987. Journal version: Multiple assignment scheme for sharing secret. *J. Cryptol.*, 6(1):15–20, 1993.
- [29] Erich L. Kaltofen and B. David Saunders. On wiedemann’s method of solving sparse linear systems. In Harold F. Mattson, Teo Mora, and T. R. N. Rao, editors, *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, 1991, Proceedings*, volume 539 of *Lecture Notes in Computer Science*, pages 29–38. Springer, 1991.

- [30] Mauricio Karchmer and Avi Wigderson. On span programs. In *Proceedings of the Eighth Annual Structure in Complexity Theory Conference, San Diego, CA, USA, May 18-21, 1993*, pages 102–111. IEEE Computer Society, 1993.
- [31] Tianren Liu and Vinod Vaikuntanathan. Breaking the circuit-size barrier in secret sharing. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 699–708. ACM, 2018.
- [32] Peter Nelson. Almost all matroids are nonrepresentable. *Bulletin of the London Mathematical Society*, 50(2):245–248, January 2018.
- [33] Anat Paskin-Cherniavsky and Artiom Radune. On polynomial secret sharing schemes. In Yael Tauman Kalai, Adam D. Smith, and Daniel Wichs, editors, *1st Conference on Information-Theoretic Cryptography, ITC 2020*, volume 163 of *LIPICs*, pages 12:1–12:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [34] Lajos Rónyai, László Babai, and Murali K Ganapathy. On the number of zero-patterns of a sequence of polynomials. *Journal of the American Mathematical Society*, 14(3):717, 2001.
- [35] Adi Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, 1979.
- [36] Hoeteck Wee. Dual system encryption via predicate encodings. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 616–637. Springer, 2014.
- [37] Christopher Williamson. Secret sharing at the shannon ceiling. *arXiv preprint arXiv:2608.17047*, 2026.

A Omitted Preliminaries

A.1 Secret-sharing schemes

We define a secret-sharing scheme similarly to the definitions given in [14]. For more information about this definition and secret-sharing in general, see [10]. We start by defining an access structure, which is the collection of sets of parties that are authorized to reconstruct the secret. We describe an access structure by a monotone Boolean function.

Definition A.1 (Access Structures). An n -party access structure is a monotone function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ such that $f(0^n) = 0$. Let $P = \{P_1, \dots, P_n\}$ be a set of parties. For an input $x = (x_1, \dots, x_n) \in \{0, 1\}^n$, we define the set of parties that it represents as $I_x = \{P_i : x_i = 1\}$. For every $x \in \{0, 1\}^n$, if $f(x) = 1$, then we say that I_x is authorized; if $f(x) = 0$, then we say that I_x is unauthorized.

Definition A.2 (Secret-Sharing Schemes). A *secret-sharing scheme* Π with domain of secrets S , such that $|S| \geq 2$, is a mapping from $S \times R$, where R is some finite set called the set of random strings, to a tuple of n -sets $S_1 \times S_2 \times \dots \times S_n$, where S_j is called the *domain of shares* of P_j . For an input $x \in \{0, 1\}^n$, we denote $\Pi_x(s; r)$ as the restriction of $\Pi(s; r)$ to its I_x -entries, i.e., $(sh_j)_{j: x_j=1}$.

A secret-sharing scheme Π with domain of secrets S realizes an access structure $f : \{0, 1\}^n \rightarrow \{0, 1\}$ if the following two requirements hold:

Correctness. For any input $x \in \{0, 1\}^n$ such that $f(x) = 1$ there exists a *reconstruction function* $\text{Recon}_x : \prod_{\{i: x_i=1\}} S_i \rightarrow S$ such that $\text{Recon}_x(\Pi_x(s; r)) = s$ for every secret $s \in S$ and every random string $r \in R$.

Privacy. For any input $x \in \{0, 1\}^n$ such that $f(x) = 0$ and every pair of secrets $s, s' \in S$, the distributions $\Pi_x(s; r)$ and $\Pi_x(s'; r)$ are identical, where the distributions are over the choice of r from R with uniform distribution.

Given a secret-sharing scheme Π , define the *size* of the secret as $\log |S|$, the *share size* of party P_j as $\log |S_j|$, the *share size* as $\max_{1 \leq j \leq n} \{\log |S_j|\}$, the *total share size* as $\sum_{j=1}^n \log |S_j|$, and the *information ratio* as $\frac{\max_{1 \leq j \leq n} \{\log |S_j|\}}{\log |S|}$.

By default, when we talk about the share size of secret-sharing schemes for an access structure, we consider schemes for one-bit secrets.

We next define linear secret-sharing schemes, which are schemes in which the mapping that the dealer uses to generate the shares is linear.

Definition A.3 (Linear Secret-Sharing Schemes). Let Π be a secret-sharing scheme with domain of secrets S . We say that Π is a *linear* secret-sharing scheme over a finite field $\mathbb{F}$ if there are integers $\ell_r, \ell_1, \dots, \ell_n$ such that $S = \mathbb{F}$, $R = \mathbb{F}^{\ell_r}$, $S_1 = \mathbb{F}^{\ell_1}, \dots, S_n = \mathbb{F}^{\ell_n}$, and the mapping Π is a *linear* mapping over $\mathbb{F}$ from $\mathbb{F}^{1+\ell_r}$ to $\mathbb{F}^{\ell_1+\dots+\ell_n}$. The *share length* of party P_i is ℓ_i , the *maximal share length* of the scheme is $\max_{i \in [n]} \ell_i$, and its *total share length* is $\sum_{i=1}^n \ell_i$. These quantities count elements of $\mathbb{F}$.

Definition A.4 (Shamir threshold secret sharing [35]). Let $1 \leq t \leq q$, let $P_1, \dots, P_q$ be parties, and let $\mathbb{F}$ be a finite field with $|\mathbb{F}| > q$. Shamir's construction gives a linear secret-sharing scheme over $\mathbb{F}$ for the t -out-of- q threshold access structure, with one element of $\mathbb{F}$ per share.

For an arbitrary finite field $\mathbb{F}$, let $m = \lceil \log_{|\mathbb{F}|}(q+1) \rceil$ and run Shamir's construction over the degree- m extension field $\mathbb{K}$ of $\mathbb{F}$, which has $|\mathbb{K}| = |\mathbb{F}|^m > q$. Restricting the secret to the subfield $\mathbb{F} \subseteq \mathbb{K}$ preserves correctness and privacy, and the $\mathbb{K}$ -linear sharing map is also $\mathbb{F}$ -linear. Since one element of $\mathbb{K}$ is represented by m elements of $\mathbb{F}$, this gives a linear t -out-of- q threshold scheme over $\mathbb{F}$ with share length m field elements per party. In particular, for $\mathbb{F} = \mathbb{F}_2$ the share size is $\lceil \log_2(q+1) \rceil$ bits per party.

Definition A.5 (Additive secret sharing). Let $U = \{P_1, \dots, P_q\}$ be parties, and let $a \in \mathbb{F}$. To share a additively among the parties of U , the dealer samples independent uniform elements $r_1, \dots, r_{q-1} \in \mathbb{F}$, gives r_j to party P_j for every $j < q$, and gives

$$a - \sum_{j=1}^{q-1} r_j$$

to party P_q . The parties in U can reconstruct a by summing their shares, while the shares held by every strict subset of U are independent uniform elements of $\mathbb{F}$, independently of a . Hence, additive secret sharing is a linear scheme for the q -out-of- q access structure.

Fact A.6 (Closure under OR and AND). Let $f_1, \dots, f_m$ be monotone functions over the same parties, each realized by a secret-sharing scheme. Then both $\bigvee_{i=1}^m f_i$ and $\bigwedge_{i=1}^m f_i$ can be realized by secret-sharing schemes whose total and maximal share sizes are at most the sums of the corresponding share sizes of the original schemes. If the original schemes are linear over the same field, then the resulting schemes are linear as well.

Complexity conventions [6]. Given a (monotone) access structure $f : \{0, 1\}^n \rightarrow \{0, 1\}$ we let $\text{SSize}(f) := \min_{\mathcal{D} \text{ realizes } f} |\mathcal{D}|$, where $|\mathcal{D}|$ denotes the total share size of a secret-sharing scheme $\mathcal{D}$ for one-bit secrets. The *exponent* of f is $n^{-1} \cdot \log_2 \text{SSize}(f)$, and it is denoted by $\mathbf{S}(f)$. If $\mathcal{F}$ is a collection of n -party access structures then

$$\text{SSize}(\mathcal{F}) := \max_{f \in \mathcal{F}} \text{SSize}(f), \quad \text{and} \quad \mathbf{S}(\mathcal{F}) := \max_{f \in \mathcal{F}} \mathbf{S}(f).$$

When $\mathcal{F} = \{\mathcal{F}_n\}$ is a sequence of collections $\mathcal{F}_n$ of n -party access structures we think of $\text{SSize}(\mathcal{F})$ as a function of n , and define the *secret-sharing exponent* $\mathbf{S}(\mathcal{F})$ to be $\mathbf{S}(\mathcal{F}) := \limsup_{n \rightarrow \infty} \mathbf{S}(\mathcal{F}_n)$. For an access structure F , let $\text{LSSize}(F)$ denote the minimum total share size of a linear scheme realizing F . All the above definitions extend to the linear setting by replacing SSize with LSSize .

We denote by $\mathbf{U}(a, n)$ (resp., $\mathbf{U}_\ell(a, n)$) the secret-sharing exponent (resp., the linear exponent) of (a, n) -upslices and by $\mathbf{U}(\alpha)$ (resp., $\mathbf{U}_\ell(\alpha)$) the secret-sharing exponent (resp., the linear exponent) of $(\alpha n, n)$ -upslices. Unless stated otherwise, exponents are defined for schemes with a one-bit secret.

A.2 Function families

We define slices, multislices, midslices and upslices, as defined in [6].

Definition A.7 (Slices, Multislices and δ -Midslices). For $a \leq b \in [n]$, we define the $(a : b)$ -*multislice* of f to be the access structure $F : \{0, 1\}^n \rightarrow \{0, 1\}$ for which

$$F(x) = \begin{cases} 0 & \text{if } \text{wt}(x) < a \\ f(x) & \text{if } \text{wt}(x) \in [a, b] \\ 1 & \text{if } \text{wt}(x) > b \end{cases}.$$

We say that F is $(a : b, n)$ -*multislice access structure* (or just $(a : b, n)$ -multislice) if F is an $(a : b)$ -multislice of some n -party access structure f . An $(a : a)$ -multislice is referred to as an *a-slice*.

For a constant $\delta \in (0, 1/2)$, the δ -*midslice* of an n -party access structure f is its $(a : b)$ -multislice with $a = \lceil (\frac{1}{2} - \delta)n \rceil$, $b = \lfloor (\frac{1}{2} + \delta)n \rfloor$. We call an access structure obtained this way a δ -*midslice access structure*.

Definition A.8 (Upslices). For $a \in [n]$, we define the a -*upslice* of f to be the access structure $F : \{0, 1\}^n \rightarrow \{0, 1\}$ for which

$$F(x) = \begin{cases} 0 & \text{if } \text{wt}(x) < a \\ f(x) & \text{if } \text{wt}(x) = a \\ 1 \iff \exists x' : \text{wt}(x') = a, x' \leq x, f(x') = 1 & \text{if } \text{wt}(x) > a \end{cases}.$$

We say that F is an (a, n) -*upslice access structure* (or just (a, n) -upslice) if F is an (a, n) -upslice of some n -party access structure f .

Observe that F is an (a, n) -upslice if and only if all its minterms are of weight a .

B Technical Lemmas

B.1 Binomial estimates

To evaluate the combinatorial bound of Theorem 1.2 at constant densities, we will use the final claim of this subsection, which in turn relies on the preceding claims.

We first note that changing either argument of a binomial coefficient by one changes the coefficient by at most a polynomial factor. We will use this fact repeatedly to absorb the $O(1)$ -rounding required when converting densities to integers.

Fact B.1 (Unit shifts). *For all integers $0 \leq r \leq m$, each of the ratios $\binom{m}{r \pm 1} / \binom{m}{r}$ and $\binom{m \pm 1}{r} / \binom{m}{r}$ that is defined and nonzero equals one of*

$$\frac{m-r}{r+1}, \quad \frac{r}{m-r+1}, \quad \frac{m+1}{m+1-r}, \quad \frac{m-r}{m},$$

and lies in $[1/(m+1), m+1]$. Hence, changing either argument by one changes $\log \binom{m}{r}$ by at most $\log(m+1)$.

Together with a comparison to the binomial distribution, Fact B.1 yields the following estimate for binomial coefficients when the normalized index j/m differs from a fixed density by $O(1/m)$.

Claim B.2 (Binomial coefficients near a fixed density). *Let $\beta \in (0, 1)$ and $c > 0$ be constants. There exists a constant $C > 0$, depending only on β and c , such that for all integers $m \geq 1$ and $0 \leq j \leq m$ with $|j - \beta m| \leq c$,*

$$\left| \log \binom{m}{j} - H_2(\beta) m \right| \leq C \log(m+1).$$

Proof. Consider the $m+1$ nonnegative numbers

$$p_i = \binom{m}{i} \beta^i (1-\beta)^{m-i}, \quad 0 \leq i \leq m,$$

which sum to one by the binomial theorem. Taking logarithms and rearranging we get for every i that

$$\log \binom{m}{i} = \log p_i + H_2(\beta) m + (i - \beta m) \log \frac{1-\beta}{\beta}. \quad (14)$$

For $i = j$ the last term has absolute value at most $c \left| \log \frac{1-\beta}{\beta} \right|$. Since $p_j \leq 1$, (14) therefore bounds $\log \binom{m}{j} - H_2(\beta) m$ from above by a constant depending only on β and c . For the lower bound, the ratio $p_{i+1}/p_i = \frac{m-i}{i+1} \cdot \frac{\beta}{1-\beta}$ is at least one precisely when $i \leq \beta(m+1) - 1$, so the sequence $(p_i)_i$ increases up to some index $i^* = \beta m + O(1)$ and decreases afterwards. As the $m+1$ numbers sum to one, $p_{i^*} \geq 1/(m+1)$, and (14) with $i = i^*$ gives

$$\log \binom{m}{i^*} - H_2(\beta) m \geq -\log(m+1) - O(1).$$

Since both j and i^* are a constant away from βm it holds that $|j - i^*| = O(1)$. Applying Fact B.1 $O(1)$ times yields

$$\log \binom{m}{j} - H_2(\beta) m \geq -O(\log(m+1)).$$

Combining the two bounds and choosing a sufficiently large constant C proves the claim. $\square$

We will also use the following uniform version, including indices near the endpoints:

$$\frac{2^{mH_2(j/m)}}{m+1} \leq \binom{m}{j} \leq 2^{mH_2(j/m)} \quad (0 \leq j \leq m, m \geq 1). \quad (15)$$

For $0 < j < m$, take $\beta = j/m$ in the binomial distribution above. Its probability at j is a maximum, so it lies between $1/(m+1)$ and 1, giving both inequalities. The endpoint cases are immediate. The resulting $O(\log(m+1))$ error is uniform in j .

Combining these estimates gives the exponent of the arbitrary-profile construction.

Claim B.3 (Evaluation of the layer bound). *Fix constants α, λ, τ satisfying the feasibility conditions preceding Corollary 6.4, and put*

$$k = \alpha n + O(1), \quad q = \lambda n + O(1), \quad t = \tau n + O(1).$$

With τ_L, τ_R, I as defined there, we have

$$p_{k,q,t}^{-1} = 2^{I(\alpha;\lambda,\tau)n + O(\log n)}$$

and

$$\binom{n-q}{k-t} + \sum_{j=t}^q \binom{q}{j} = 2^{\max\{(1-\lambda)H_2(\tau_L), \lambda\kappa(\tau_R)\}n + O(\log n)}.$$

Consequently, the exponent in (6) is $I(\alpha; \lambda, \tau) + \max\{(1-\lambda)H_2(\tau_L), \lambda\kappa(\tau_R)\}$.

Proof. The identity

$$p_{k,q,t} = \frac{\binom{q}{t} \binom{n-q}{k-t}}{\binom{n}{k}}$$

and Claim B.2 give

$$\log p_{k,q,t}^{-1} = [H_2(\alpha) - (1-\lambda)H_2(\tau_L) - \lambda H_2(\tau_R)]n + O(\log n).$$

Likewise, the masking-side term has logarithm $(1-\lambda)H_2(\tau_L)n + O(\log n)$.

For the disclosure-side term, let $r = \max(t, \lceil q/2 \rceil)$. The coefficient $\binom{q}{r}$ is largest among $\binom{q}{t}, \dots, \binom{q}{q}$, so

$$\binom{q}{r} \leq \sum_{j=t}^q \binom{q}{j} \leq (q+1) \binom{q}{r}.$$

Since $r/q = \max(\tau_R, 1/2) + O(1/n)$, the logarithm of this sum is $\lambda\kappa(\tau_R)n + O(\log n)$. The sum of the two per-partition terms is within a factor of two of the larger term. Polynomial overheads contribute only $O(\log n)$ to the logarithm. $\square$

Together with Lemma 5.1, the first estimate also shows that

$$C(n, k; q, t) = 2^{I(\alpha;\lambda,\tau)n + O(\log n)}.$$

Thus the covering loss is optimal at the exponential scale for a fixed profile. The original profile $\tau = \alpha\lambda$ has $I = 0$, and Lemma 5.2 gives an explicit cover in this case.

B.2 Quadratic Midslice Schemes

Proof of Theorem 6.2. We plug the quadratic RCDS protocol of Beimel, Othman, and Peter [16, Thm. 6.5] into the RCDS-to-midslice construction of Applebaum et al. [4, Lemmas 5.5 and 5.8]. We use the parameters from the proof of [4, Lem. 5.10], which performs the same procedure as we do here with their RCDS protocols.

Up to adding $o(n)$ dummy parties and rounding the parameters (which changes the share size by only a $2^{o(n)}$ factor), assume that $n = K\ell$ with $K = \ell = \sqrt{n}$. Set $\varepsilon = n^{-0.1}$ and

$$a = \left(\frac{1}{2} - \delta - \varepsilon\right)\ell, \quad b = \left(\frac{1}{2} + \delta + \varepsilon\right)\ell.$$

The RCDS input domain of each server has size $N = 2^\ell$, and the robustness required by the compiler is

$$T = \sum_{j=a}^b \binom{b}{j}.$$

If the RCDS messages have length m , Lemmas 5.5 and 5.8 of [4] give a δ -midslice scheme of total share size $m 2^{o(n)}$.

Theorem 6.5 of [16] gives a quadratic K -server T -RCDS protocol over $\mathbb{F}_2$ for T s that satisfy $T < \min\{N/2K, 2^{\sqrt{N}/K}\}$ with

$$m = O\left(N^{(K-1)/3} T^{2(K-1)/3+1} K^{2K} \log^2 N \log^{(4K-1)/3} T\right)$$

The requirement for T is satisfied for all sufficiently large n : indeed, $T \leq 2^b = N^{1/2+\delta+o(1)} = o(N/K)$ since $\delta < 1/2$, and $\log T = O(\ell) = o(\sqrt{N}/K)$.

We next compute the midslice share size, beginning with an estimate of T . The usual entropy bound for binomial sums yields

$$\log T = \tau(\delta)\ell + o(\ell), \quad \tau(\delta) = \begin{cases} \left(\frac{1}{2} + \delta\right) H_2\left(\frac{1/2 - \delta}{1/2 + \delta}\right), & 0 \leq \delta \leq \frac{1}{6}, \\ \frac{1}{2} + \delta, & \frac{1}{6} \leq \delta < \frac{1}{2}. \end{cases}$$

For the second range, the interval $[a, b]$ contains the middle binomial coefficient and the exponent is b , and for the first range the exponent is dominated by the expression $\binom{b}{a}$ with $a \geq b/2$.

Taking logarithms in the RCDS bound yields

$$\begin{aligned} \log m &\leq \frac{K-1}{3}\ell + \left(\frac{2(K-1)}{3} + 1\right) (\tau(\delta)\ell + o(\ell)) + o(n) \\ &= \left(\frac{1}{3} + \frac{2}{3}\tau(\delta)\right)n + o(n), \end{aligned}$$

where the equality is implied by $K\ell = n$, and by the fact that the factors K^{2K} and the logarithmic factors contribute only $O(\sqrt{n} \log n) = o(n)$. Substituting the two values of $\tau(\delta)$ gives

$$\frac{1}{3} + \frac{2}{3}\tau(\delta) = \begin{cases} \frac{1}{3} + \frac{1+2\delta}{3} H_2\left(\frac{1-2\delta}{1+2\delta}\right), & 0 \leq \delta \leq \frac{1}{6}, \\ \frac{2}{3} + \frac{2\delta}{3}, & \frac{1}{6} \leq \delta < \frac{1}{2}, \end{cases}$$

which is $Q(\delta)$.

Finally, the RCDS protocol of [16, Thm. 6.5] has quadratic sharing and reconstruction. The compiler of [4, Lemmas 5.5 and 5.8] distributes and recovers these messages using linear secret-sharing schemes. Composing these linear maps with the quadratic RCDS maps keeps the degree at most two, so the resulting midslice scheme has quadratic sharing and reconstruction. Its total share size is $m 2^{o(n)} = 2^{Q(\delta)n+o(n)}$, as claimed. $\square$

B.3 Simultaneous Leading Minors

Proof of Lemma 7.8. We fix the field $\mathbb{F}$ and the matrices $Y_1, \dots, Y_t$ over $\mathbb{F}$, and write $r_j = \text{rank}_{\mathbb{F}}(Y_j)$. For any field extension $E/\mathbb{F}$ and any matrices $L, R \in E^{(d+1) \times (d+1)}$, we have

$$\text{rank}_E(LY_jR) \leq \text{rank}_E(Y_j) = r_j,$$

where the equality follows from Lemma 7.5. Thus, if $k > r_j$, every $k \times k$ minor of LY_jR vanishes, including $\Delta_k(Y_j; L, R)$.

For the other direction, we now write $L = (\ell_{ab})$ and $R = (u_{ab})$, treating their $2(d+1)^2$ entries as independent formal variables and the entries of the Y_j 's as fixed coefficients in $\mathbb{F}$. We work in the polynomial ring

$$S = \mathbb{F}[\ell_{ab}, u_{ab} : a, b \in [d+1]].$$

Each $\Delta_k(Y_j; L, R)$ belongs to S . At the end of the proof, we assign values to the variables in a suitable finite extension of $\mathbb{F}$.

Fix one pair (j, k) with $k \leq r_j$. There is a nonzero $k \times k$ minor of Y_j . Choose permutation matrices $L_0, R_0 \in \mathbb{F}^{(d+1) \times (d+1)}$ that move the rows and columns of this minor to the first k positions. Then $\Delta_k(Y_j; L_0, R_0)$ equals that minor up to sign and is therefore nonzero. This assignment to the entries of L, R shows that $\Delta_k(Y_j; L, R)$ is a nonzero polynomial.

Next, we find matrices L^*, R^* that do not depend on (j, k) , and provide the same guarantee. In the polynomial ring S , we define the following polynomial

$$P(L, R) = \prod_{j=1}^t \prod_{k=1}^{\text{rank}_{\mathbb{F}}(Y_j)} \Delta_k(Y_j; L, R).$$

Every factor is a nonzero polynomial by the prior argument with L_0 and R_0 . A polynomial ring over a field has no zero divisors, so their product P is also a nonzero polynomial, with an empty product interpreted as 1. Hence, it holds for any particular assignment (L_*, R_*) that

$$P(L_*, R_*) \neq 0 \iff \text{every factor is nonzero with the assignment } (L_*, R_*).$$

Let $D_0 = \deg P$. If $\mathbb{F}$ is infinite, set $E = \mathbb{F}$. If $\mathbb{F}$ is finite, choose a finite extension $E/\mathbb{F}$ with $|E| > D_0$. In either case, $E/\mathbb{F}$ is a finite extension and there is a finite subset $\Omega \subseteq E$ with $|\Omega| > D_0$; when E is finite, we may take $\Omega = E$. The polynomial P remains nonzero when regarded as a polynomial over E , because the embedding $\mathbb{F} \hookrightarrow E$ is injective. By the Schwartz-Zippel lemma applied to Ω , there exist matrices $L_*, R_* \in E^{(d+1) \times (d+1)}$ such that $P(L_*, R_*) \neq 0$. For this choice, every required leading minor is nonzero. As shown in the first part of the proof, whenever $k > \text{rank}_{\mathbb{F}}(Y_j)$, the leading minor $\Delta_k(Y_j; L, R)$ is zero for all choices of L, R , including L_*, R_* . Therefore, taking $L = L_*$ and $R = R_*$ proves the lemma. $\square$