

The Residue Is the Resource: Exact Fresh-State Complexity, Reconstruction Limits, and Challenge Design in Folding

Xiang Wang¹ and Ji Xiang²

¹ Guizhou University

ee.wangx24@gzu.edu.cn

² Chongqing University

20240674@stu.cqu.edu.cn

Abstract. Folding protocols cross public-randomness boundaries: input-dependent information must be fixed before a challenge is known, while the required output may depend on it. We ask how much state must be retained after accounting for information already available at the boundary. With unrestricted prechallenge preprocessing and linear postchallenge readout, a retained representation is feasible exactly when its coordinates span, modulo this available information, every output direction induced by the supported challenges. Thus the minimum retained width is the dimension of this residual output space.

The width depends on both the admissible source domain and the actual challenge support. Coordinatewise multiplication of length- n^2 vectors has width n^2 on the full product domain but $2n - 1$ when both inputs lie in the dimension- n Reed–Solomon code evaluated at n^2 distinct points. Below the exact-width threshold, for independent uniform inputs to a bilinear map, we characterize the optimal exact-reconstruction probability by a rank-weighted subspace optimization. Bilinear side information couples the missing directions through a common linear correction, so optimizing them independently can fail.

Applying the framework to direct homogeneous quadratic folding on a full product domain, an algebraic separation condition on the challenge family forces width at least rk , where r is the dimension of the span of mixed bilinear outputs. Scaled-Cauchy challenges attain this bound on any support of at least $2k + 1$ distinct non-pole values and yield an explicit matching rk -coordinate polynomial representation.

1 Introduction

Folding and accumulation protocols repeatedly cross public-randomness boundaries: some input-dependent information must be fixed before a verifier challenge is known, while the output required afterward may depend on that challenge. We study how many fresh field elements must remain available across such a boundary after accounting for information already present at the chosen cut.

The state carried by a concrete protocol need not be minimal. Some coordinates may be irrelevant to the later challenge, while other information may

already be supplied by the incoming state and should not be charged again. We therefore fix a transcript cut and ask only for the additional scalar information that must be retained before the challenge is sampled.

The answer can depend strongly on the admissible source domain. Consider coordinatewise multiplication of vectors of length n^2 , and suppose the challenge support contains a nonzero value. If both sources range freely over $\mathbb{F}^{n^2}$, then, after treating functions of either source alone as free, the retained width is n^2 . Now assume that n^2 distinct evaluation points are available and restrict both sources to the dimension- n Reed–Solomon code evaluated at those points. Coordinatewise products then lie in the evaluation space of polynomials of degree below $2n - 1$, so the exact width drops to $2n - 1$. The coordinatewise relation is unchanged; only the admissible source tuples have changed.

To formalize this phenomenon, let X be the admissible input domain at a chosen boundary, let P and Y be finite-dimensional $\mathbb{F}$ -vector spaces, let $K : X \rightarrow P$ denote the incoming state, and let $H_c : X \rightarrow Y$ denote the output required after a supported challenge c . Before the challenge, an arbitrary map $\Psi : X \rightarrow \mathbb{F}^m$ may retain m field elements; no linearity is imposed on this preprocessing, and the same Ψ must serve every supported challenge. Afterward, challenge-dependent linear maps use the incoming and retained states. We also fix a subspace $\mathcal{F} \leq \mathbb{F}^X$ of scalar functions declared free in the boundary accounting, so contributions whose scalarizations lie in $\mathcal{F}$ are not charged.

The corresponding algebraic object removes information already available at the cut. Let $Q := \mathbb{F}^X / \mathcal{F}$. Linear scalar observations of the incoming state generate $U(K) \leq Q$, while scalarizations of the outputs over the actual challenge support generate $V(H) \leq Q$. The residual target $T_{H|K}$ is the image of $V(H)$ in $Q/U(K)$.

Our characterization applies to a fixed retained state, not only to the optimum. If $\Psi = (\psi_1, \dots, \psi_m)$, then Ψ is feasible exactly when the residual classes of its coordinates span a subspace containing $T_{H|K}$. Consequently,

$$m_{\min}(H \mid K; \mathcal{F}) = \dim_{\mathbb{F}} T_{H|K} = \dim_{\mathbb{F}} \frac{U(K) + V(H)}{U(K)}. \quad (1.1)$$

Although Ψ may be nonlinear, its later linear use depends only on the residual scalar classes of its coordinates. Nonlinear preprocessing therefore cannot reduce the width below the residual dimension. At minimum width, these classes form a basis of $T_{H|K}$; more generally, every feasible state can be compressed by a fixed linear map to minimum width for the same target and boundary interface.

The quantity in (1.1) is relative to the chosen cut. It depends on the admissible inputs, the actual challenge support, the functions declared free, and the incoming state. It counts only fresh $\mathbb{F}$ -coordinates carried across that boundary; it is not a measure of total communication or computation, and a width statement by itself is not a soundness statement. We use this characterization to study restricted source domains, reconstruction below the exact-width threshold, and challenge design.

1.1 Our Results

Exact retained width and restricted domains Beyond (1.1), Section 3 gives a finite-support refinement separating directions absent from the actual challenge support from generated directions already removed by free or incoming information. Section 4 specializes the framework to quadratic folding on restricted domains. For restricted R1CS, the exact direct width is the dimension of a corresponding matrix space. For coded multiplication with multiplicative input spaces U, V and linear-term space W , the exact direct width is $\dim_{\mathbb{F}}(U \star V + W)$, where $\star$ denotes the Schur product, recovering the $2n - 1$ Reed–Solomon width above.

Reconstruction below the exact width Section 5 asks what remains possible below the exact-width threshold. For independent uniform inputs to a bilinear map over a finite field, we characterize the optimal probability of exact reconstruction by a rank-weighted subspace optimization. The reconstruction profile therefore depends not only on the number of missing residual directions, but also on the ranks of the corresponding scalar bilinear forms.

With bilinear side information, several missing quotient directions cannot in general be optimized independently: their representatives must arise from one common linear use of the incoming side information. This compatibility constraint allows tasks with the same exact retained width to have different sub-threshold reconstruction profiles. For polynomial multiplication, the scalar forms are symmetric Hankel matrices; we determine the optimum for deficits $s \leq \min\{3, k\}$, with the equality cases characterized in Appendix D.

Challenge design from retained width Section 6 turns the same accounting into a design problem. In the direct full-product quadratic model with one accumulator and k fresh sources, let r be the mixed-output dimension and d_a the dimension, on the actual challenge support, of the span of the challenge functions a_i and their pairwise products $a_i a_j$. We prove that the exact direct width is rd_a . Diagonal separation forces $d_a \geq k$ and $|\mathcal{C}| \geq d_a + k + 1 \geq 2k + 1$, hence width at least rk . Scaled-Cauchy weights with distinct poles and nonzero scales attain $d_a = k$ on every support of at least $2k + 1$ distinct non-pole values and admit a matching rk -coordinate polynomial carrier.

1.2 Related Work

Folding and accumulation provide the cryptographic setting for our question. Nova gives the relaxed-R1CS folding relation used in one of our main examples [12]. Subsequent work develops broader relations, multi-instance folding, and other accumulation mechanisms; representative examples include HyperNova, ProtoStar, ProtoGalaxy, Mova, Accumulation Without Homomorphism, and WARP [10, 2, 9, 7, 4, 3]. These works study how recursive state is constructed and verified. Our question is complementary: at a fixed public-randomness cut, we ask how much new scalar information is necessary after accounting for information already available there.

Static data structures provide a related preprocessing/query viewpoint. In the linear setting, a field-valued database is preprocessed before a later linear query, and one studies space–query tradeoffs [8]. Our cost measure instead charges fresh field elements retained across one boundary, relative to free and incoming information. In the linear full-domain specialization, retained width reduces to a matrix-rank difference.

Bilinear complexity measures the number of rank-one bilinear products needed to realize a bilinear map [1]. Retained width measures a different resource: a retained coordinate may be any function computed before the challenge and need not be a rank-one bilinear product. The matrix spaces appearing in our restricted quadratic calculations therefore measure independent mixed information, not multiplication complexity.

Our coded-multiplication calculations use the classical componentwise, or Schur, product of linear codes [17,15]. In particular, the Reed–Solomon product dimensions used in Section 4 are classical. Their role here is operational: the corresponding product space determines the exact retained width at the chosen boundary.

Section 5 is related to matrix spaces over finite fields. In the rank-metric literature, the rank distribution records how many matrices in a linear space have each possible rank [6,18]. Our reconstruction problem uses these ranks in a different optimization, while bilinear side information imposes a joint compatibility constraint. For polynomial multiplication, we also use Meshulam’s theorem on finite-field Hankel spaces [14].

A separate manuscript, *CauchyFold* [19], instantiates the scaled-Cauchy carrier of Section 6 in a concrete lattice folding protocol. The width and challenge-design results used there are established here; *CauchyFold* develops the commitment, reduction, security, parameter, and implementation layers.

Section 2 gives a technical overview. Sections 3–6 develop the exact-width characterization, restricted-domain calculations, sub-threshold reconstruction, and challenge design, respectively. The appendices give the ring-valued extension, supporting proofs, multi-boundary accounting, selected protocol-boundary identifications, and a related adaptive-query recovery model.

2 Technical Overview

This section explains why the main results take the forms stated above. The same boundary viewpoint underlies all four parts: retained information is measured only after free and incoming directions have been removed, while the admissible source domain and the actual challenge support determine which directions survive to be measured.

2.1 What Must Cross the Boundary?

Let $Q = \mathbb{F}^X / \mathcal{F}$, and let $\rho_K : \mathbb{F}^X \rightarrow Q/U(K)$ be the composite quotient map. For a prechallenge state $\Psi = (\psi_1, \dots, \psi_m)$, write $R_K(\Psi)$ for the span of the residual classes $\rho_K(\psi_1), \dots, \rho_K(\psi_m)$.

The feasibility criterion is $T_{H|K} \subseteq R_K(\Psi)$. Modulo free and incoming directions, every scalarization of every supported output must lie in the span of the retained coordinate classes; conversely, such scalar decompositions assemble into the required postchallenge readouts. Hence each retained coordinate contributes at most one residual direction, while a basis of $T_{H|K}$ attains the minimum width.

Challenge support enters before dimension is counted: unsupported coefficient directions never appear, while supported ones may still vanish modulo free and incoming information.

2.2 The Domain Determines the Mixed Directions

In the direct two-source model on a product domain $X_0 \times X_1$, we treat every scalar function of either source alone as free. Rectangle differences cancel these source-local terms and isolate the mixed interaction. For a quadratic fold, what remains after scalarization is a bilinear form evaluated on variations of the two sources.

The admissible domain determines how much of this interaction remains visible. Restricting the source sets restricts the variation spaces on which the bilinear forms are tested, so ambiently distinct mixed directions can become equal or vanish on the admissible domain.

Under the distinct-evaluation-point assumptions from Section 1, Reed–Solomon sources of degree below n produce a crossed term in the degree- $< 2n - 1$ evaluation space, giving an optimal $2n - 1$ -coordinate retained representation.

Restricted R1CS exhibits the same principle without relying on polynomial degree. Let Δ_0 and Δ_1 be the variation spaces generated by differences of admissible assignments on the two sides. Rectangle differences restrict each scalarized cross term to a bilinear form on $\Delta_0 \times \Delta_1$. After bases are chosen, these forms span a matrix space whose dimension is exactly the direct retained width. What matters is the dimension of this matrix space, not the rank of an individual matrix.

2.3 Reconstruction Below the Exact Width

Below the exact-width threshold, dimension alone no longer determines performance. Section 5 considers independent uniform source records over $\mathbb{F}_q$, with one width- D prechallenge state serving all supported challenges in $\mathcal{C} \subseteq \mathbb{F}_q^\times$, one-source functions free, and performance measured in the worst case over $c \in \mathcal{C}$.

Let $\mu : U \times V \rightarrow Y$ be bilinear, with independent uniform $u \in U$ and $v \in V$, and assume that $\mu(U, V)$ spans Y . Then

$$P_\mu^*(D) = \max_{\substack{W \leq Y \\ \dim W \leq D}} \Pr[\mu(u, v) \in W] = \max_{\substack{W \leq Y \\ \dim W \leq D}} q^{-\text{codim } W} \sum_{\lambda \in W^\perp} q^{-\text{rank}(\lambda \circ \mu)}.$$

Here $W^\perp \leq Y^\vee$ is the annihilator of W .

The first expression is attained by choosing a fixed projection $\Pi : Y \rightarrow W$ and retaining the coordinates of ΠT_μ , where T_μ is the crossed mixed term.

Reconstruction succeeds exactly when the folded bilinear value lies in W . Under the nonzero-challenge model of Section 5, the folded factors remain independent and uniform for every supported challenge.

The second expression explains why rank becomes relevant. If a scalar bilinear form β has rank r , then a nontrivial additive character has average q^{-r} on independent uniform inputs. Character orthogonality therefore turns concentration in W into a sum weighted by the ranks of the missing scalar forms. Conversely, the bilinear character estimate gives a uniform upper bound on agreement against arbitrary source-local corrections of the form $f(R_0) + g(R_1)$, yielding the same optimum.

Bilinear side information changes which representatives of the missing directions may be chosen. If the incoming mixed state is generated by $\nu : U \times V \rightarrow P$, a linear use of it replaces μ by $\mu - M\nu$ for some $M : P \rightarrow Y$. For one missing direction this permits a favorable representative modulo the side-information space. For several directions, however, all corrections must arise from the same M . Dually, the representative chosen for each missing scalar direction must vary linearly with that direction; equivalently, feasible representatives arise from one common linear section on the missing subspace. Individually minimum-rank representatives need not be jointly attainable.

For polynomial multiplication on degree- $< k$ polynomials, $d_Y = 2k - 1$ and the scalar bilinear forms are symmetric Hankel matrices. Writing the deficit as $s := d_Y - D$, the problem becomes a rank-distribution optimization over low-dimensional subspaces of the Hankel matrix space. Section 5 determines the optimum for $s \leq \min\{3, k\}$, and Appendix D characterizes the equality cases.

2.4 Using Width as a Design Criterion

The same framework can be used in the opposite direction: instead of fixing a challenge family and measuring its width, we ask which challenge functions make the relevant coefficient span small.

Consider the direct full-product model with one accumulator and k fresh sources, all ranging over a vector space S . One-source information is free, and there is no additional incoming mixed state. Let $B : S \times S \rightarrow Y$ be the bilinear polarization of the homogeneous quadratic map being folded. Write $Y_B := \text{span}_{\mathbb{F}} B(S, S)$ and $r := \dim_{\mathbb{F}} Y_B$. On the finite challenge support $\mathcal{C}$, let $V_a \leq \mathbb{F}^{\mathcal{C}}$ be the span of the functions a_i and $a_i a_j$ for $i < j$, and let $d_a = \dim_{\mathbb{F}} V_a$. The exact direct width is rd_a .

The two factors have separate origins: r measures the possible mixed output values, whereas d_a measures the independent challenge-coefficient directions. Because the source domain is the full product, rectangle differences can isolate every mixed source pair, so no further domain-induced collapse occurs.

To obtain a nontrivial design problem, we impose diagonal separation: the classes of $1, a_1^2, \dots, a_k^2$ must be linearly independent modulo V_a . This algebraic condition keeps the constant and diagonal coefficient functions distinguishable from the mixed span; it is not, by itself, a general soundness condition. Diagonal

separation forces $d_a \geq k$ and $|\mathcal{C}| \geq d_a + k + 1 \geq 2k + 1$, hence direct retained width at least rk .

Scaled-Cauchy weights

$$a_i(c) = \frac{\lambda_i}{c - \xi_i},$$

with distinct poles and nonzero scales, attain this lower bound on every support of at least $2k + 1$ distinct non-pole values. Partial fractions give $a_i a_j \in \text{span}\{a_i, a_j\}$ for $i \neq j$, so pairwise products introduce no new challenge directions. Diagonal separation then gives $d_a = k$ and direct width rk .

Writing $D(T) = \prod_i (T - \xi_i)$, multiplying the mixed contribution by $D(c)$ clears the Cauchy denominators and yields the evaluation of a Y_B -valued polynomial of degree below k . Its k coefficients form an explicit minimum-width retained state of exactly rk field coordinates.

3 Exact Retained Width

We formalize retained width at a single public-randomness boundary and prove its exact field-level characterization. Throughout this section, $\mathbb{F}$ is a field. The ring-valued extension appears in Appendix A.

3.1 The Boundary Model

Fix a public-randomness boundary. Let X be the set of admissible semantic inputs at this boundary, and let $\mathcal{C}$ be a finite nonempty challenge support. Let P and Y be finite-dimensional $\mathbb{F}$ -vector spaces. The incoming state is a map $K : X \rightarrow P$, and for each $c \in \mathcal{C}$, the required postchallenge output is a map $H_c : X \rightarrow Y$. No algebraic structure on X is assumed.

Fix a subspace $\mathcal{F} \leq \mathbb{F}^X$ of scalar functions credited as free at this cut, where $\mathbb{F}^X$ denotes the vector space of all functions from X to $\mathbb{F}$. For an $\mathbb{F}$ -vector space M , write $M^\vee := \text{Hom}_{\mathbb{F}}(M, \mathbb{F})$.

For an integer $m \geq 0$, a width- m realization consists of a prechallenge map $\Psi : X \rightarrow \mathbb{F}^m$ and, for every $c \in \mathcal{C}$, a map $G_c : X \rightarrow Y$ together with linear maps $L_{K,c} : P \rightarrow Y$ and $L_{\Psi,c} : \mathbb{F}^m \rightarrow Y$ such that

$$H_c(x) = G_c(x) + L_{K,c}(K(x)) + L_{\Psi,c}(\Psi(x)) \quad \text{for every } x \in X. \quad (3.1)$$

The term G_c represents the contribution credited as free: we require $\eta \circ G_c \in \mathcal{F}$ for every $\eta \in Y^\vee$.

No linearity assumption is imposed on Ψ . It is computed before the challenge is sampled, and the same Ψ must work for every $c \in \mathcal{C}$. After the challenge is known, $L_{K,c}$ and $L_{\Psi,c}$ may depend on c and on fixed public parameters of the boundary, but not on the realized values of x , $K(x)$, or $\Psi(x)$.

We charge only the m coordinates of $\Psi(x)$; the incoming state and the free contribution are not charged as fresh retained state. Define $m_{\min}(H \mid K; \mathcal{F})$ as the minimum $m \geq 0$ for which a width- m realization exists.




Fig. 1. A public-randomness boundary for (3.1). Arbitrary preprocessing produces $\Psi(x)$; after c is sampled, the readout is linear in the incoming and retained states. Only the m newly retained field elements are charged.

3.2 The Residual Target and the Exact Minimum

We now remove, at the scalar level, the information already credited at the boundary. Put $\mathcal{Q} := \mathbb{F}^X / \mathcal{F}$, and let $\pi_{\mathcal{F}} : \mathbb{F}^X \rightarrow \mathcal{Q}$ be the quotient map. Linear scalar observations of the incoming state define $U(K)$, while scalarizations of the supported outputs generate $V(H)$:

$$U(K) := \{\pi_{\mathcal{F}}(\omega \circ K) : \omega \in P^\vee\} \leq \mathcal{Q}, \quad (3.2)$$

$$V(H) := \text{span}_{\mathbb{F}}\{\pi_{\mathcal{F}}(\eta \circ H_c) : c \in \mathcal{C}, \eta \in Y^\vee\} \leq \mathcal{Q}. \quad (3.3)$$

Let $\pi_K : \mathcal{Q} \rightarrow \mathcal{Q}/U(K)$ be the quotient map and set $\rho_K := \pi_K \circ \pi_{\mathcal{F}}$. The residual target is the image of the supported scalar-output space after the incoming scalar information is also quotiented out:

$$T_{H|K} := \pi_K(V(H)) \leq \mathcal{Q}/U(K), \quad (3.4)$$

$$T_{H|K} \cong \frac{U(K) + V(H)}{U(K)}, \quad (3.5)$$

$$\ker \rho_K = \mathcal{F} + \{\omega \circ K : \omega \in P^\vee\}. \quad (3.6)$$

Thus $T_{H|K}$ is generated by the classes $\rho_K(\eta \circ H_c)$ as c ranges over $\mathcal{C}$ and η ranges over $Y^\vee$. Its definition uses only the boundary data, whereas $m_{\min}(H | K; \mathcal{F})$ is defined operationally by minimizing over prechallenge retained states.

For a candidate retained state $\Psi = (\psi_1, \dots, \psi_m) : X \rightarrow \mathbb{F}^m$, define its residual coordinate span by

$$R_K(\Psi) := \text{span}_{\mathbb{F}}\{\rho_K(\psi_1), \dots, \rho_K(\psi_m)\} \leq \mathcal{Q}/U(K).$$

Theorem 3.1 (Feasibility and exact retained width).

In the boundary model of Section 3.1, a candidate state $\Psi : X \rightarrow \mathbb{F}^m$ can serve as the retained state of a width- m realization if and only if

$$T_{H|K} \subseteq R_K(\Psi).$$

Consequently,

$$m_{\min}(H \mid K; \mathcal{F}) = \dim_{\mathbb{F}} T_{H|K}. \quad (3.7)$$

Writing $d := \dim_{\mathbb{F}} T_{H|K}$, a candidate state with d coordinates is feasible, and hence minimum-width, exactly when its residual coordinate classes form a basis of $T_{H|K}$.

Proof. Suppose $\Psi = (\psi_1, \dots, \psi_m)$ is feasible. Fix $c \in \mathcal{C}$ and $\eta \in Y^\vee$. Applying η to (3.1) expresses $\eta \circ H_c$ as the sum of a free scalar function, a linear scalar observation of K , and a linear combination of $\psi_1, \dots, \psi_m$. The first two terms vanish under ρ_K . Hence every generator of $T_{H|K}$ lies in $R_K(\Psi)$, proving the required inclusion.

Conversely, suppose $T_{H|K} \subseteq R_K(\Psi)$. Choose a basis $y_1, \dots, y_N$ of Y , with dual basis $\eta_1, \dots, \eta_N$. For every $c \in \mathcal{C}$ and $j \in \{1, \dots, N\}$, choose coefficients $a_{c,j,i} \in \mathbb{F}$ expressing $\rho_K(\eta_j \circ H_c)$ as a linear combination of $\rho_K(\psi_1), \dots, \rho_K(\psi_m)$. The remaining difference lies in $\ker \rho_K$; hence, by (3.6), there exist $f_{c,j} \in \mathcal{F}$ and $\omega_{c,j} \in P^\vee$ such that

$$\eta_j \circ H_c = f_{c,j} + \omega_{c,j} \circ K + \sum_{i=1}^m a_{c,j,i} \psi_i. \quad (3.8)$$

Define $G_c := \sum_j y_j f_{c,j}$ and $L_{K,c}(p) := \sum_j y_j \omega_{c,j}(p)$. For $z \in \mathbb{F}^m$, set $L_{\Psi,c}(z) := \sum_j y_j \sum_i a_{c,j,i} z_i$. Both readout maps are linear. For every $\eta \in Y^\vee$, the scalar function $\eta \circ G_c$ is a linear combination of the $f_{c,j}$ and therefore lies in $\mathcal{F}$. Multiplying (3.8) by y_j , summing over j , and using the dual-basis identity gives (3.1), proving feasibility of the given Ψ .

Let $d := \dim_{\mathbb{F}} T_{H|K}$. Every feasible state satisfies $d \leq \dim R_K(\Psi) \leq m$. Conversely, lifting a basis of $T_{H|K}$ through ρ_K gives a feasible state with d coordinates, proving (3.7). When $m = d$, the inclusion $T_{H|K} \subseteq R_K(\Psi)$ holds exactly when the d residual coordinate classes form a basis of the target. $\square$

Linear compression of a feasible state For the fixed target and boundary interface, let $\Psi : X \rightarrow \mathbb{F}^m$ be feasible and put $d := \dim_{\mathbb{F}} T_{H|K}$. There is a rank- d linear map $A : \mathbb{F}^m \rightarrow \mathbb{F}^d$ such that $A\Psi$ is a minimum-width retained state, with the free contributions and readout maps chosen for the compressed state. Indeed, choose a basis $\beta_1, \dots, \beta_d$ of $T_{H|K}$ and express $\beta_\ell = \sum_i A_{\ell i} \rho_K(\psi_i)$. These independent residual classes force the rows of A to be independent, and $R_K(A\Psi) = T_{H|K}$ gives feasibility by Theorem 3.1. Any two minimum-width states differ, modulo free and incoming information, by an invertible linear change of coordinates, since their residual coordinate classes are bases of the same target.

Dependence on the postchallenge interface Let $X = \{x_0, x_1\}$, with no incoming or free information, and let two supported challenges require the two indicator functions of x_0 and x_1 . Their residual target has dimension two, so Theorem 3.1 gives width two. If arbitrary nonlinear postchallenge decoding were allowed, one retained field element would suffice: encode x_0 and x_1 by two distinct field

values and decode the requested indicator after the challenge. Thus the residual dimension is the exact storage cost for the linear-readout interface of Section 3.1, but not for an unrestricted postchallenge decoder.

Linear output families When the boundary maps are linear, the residual dimension has a familiar matrix form. Let $X = \mathbb{F}^n$, take $\mathcal{F} = \{0\}$, and suppose $K(x) = Kx$ and $H_c(x) = L_c x$ for fixed matrices K and L_c . Let L_{all} be the vertical stack of the matrices L_c over $c \in \mathcal{C}$. Then

$$m_{\min} = \text{rank} \begin{pmatrix} K \\ L_{\text{all}} \end{pmatrix} - \text{rank } K = \dim L_{\text{all}}(\ker K). \quad (3.9)$$

The rank difference measures the row directions contributed by the supported outputs beyond those already supplied by K . For the second equality, consider the map $x \mapsto (Kx, L_{\text{all}}x)$. Projection of its image onto $\text{im } K$ is surjective, and the kernel of this projection is naturally identified with $L_{\text{all}}(\ker K)$. Rank-nullity gives the second equality.

Appendix G considers a different recovery interface for a uniform source over $\mathbb{F}_q^n$, where an arbitrary bounded-range stored state is supplemented by adaptive same-source scalar linear queries. The same rank difference appears there, but the operational interfaces are different.

3.3 Supported Challenge Directions

Suppose the supported outputs admit an expansion

$$H_c = G_c + L_{K,c} \circ K + \sum_{j=1}^{\ell} p_j(c) Z_j, \quad c \in \mathcal{C}, \quad (3.10)$$

where $p_j : \mathcal{C} \rightarrow \mathbb{F}$, $Z_j : X \rightarrow Y$, each $L_{K,c} : P \rightarrow Y$ is linear, and $\eta \circ G_c \in \mathcal{F}$ for every $\eta \in Y^\vee$.

The formal coefficient space is $\mathcal{D} := (Y^\vee)^\ell$, but only certain directions in $\mathcal{D}$ are generated on the actual challenge support. Define the supported coefficient space

$$D_{\mathcal{C}} := \text{span}_{\mathbb{F}} \{ (p_1(c)\eta, \dots, p_\ell(c)\eta) : c \in \mathcal{C}, \eta \in Y^\vee \} \leq \mathcal{D}. \quad (3.11)$$

Not every supported coefficient tuple contributes new information at the boundary. Define

$$S_Z : \mathcal{D} \longrightarrow \mathcal{Q}/U(K), \quad S_Z(\lambda_1, \dots, \lambda_\ell) := \sum_{j=1}^{\ell} \rho_K(\lambda_j \circ Z_j), \quad (3.12)$$

and let $N_Z := \ker S_Z$. Thus N_Z consists of coefficient combinations whose contribution vanishes after free and incoming information have been removed.

Proposition 3.2 (Supported coefficient directions).

For the family in (3.10),

$$T_{H|K} = S_Z(D_C), \quad (3.13)$$

$$m_{\min}(H | K; \mathcal{F}) = \dim_{\mathbb{F}} D_C - \dim_{\mathbb{F}}(D_C \cap N_Z). \quad (3.14)$$

Proof. Fix $c \in \mathcal{C}$ and $\eta \in Y^\vee$. Applying η to (3.10) and then applying ρ_K removes the free and incoming terms, giving

$$\rho_K(\eta \circ H_c) = S_Z(p_1(c)\eta, \dots, p_\ell(c)\eta). \quad (3.15)$$

As c and η vary, the left-hand side generates $T_{H|K}$, while the tuples on the right generate D_C . Hence $T_{H|K} = S_Z(D_C)$. Applying rank-nullity to $S_Z|_{D_C}$, followed by Theorem 3.1, gives (3.14). $\square$

The supported challenges generate the same residual image as the formal coefficient space exactly when

$$S_Z(D_C) = S_Z(\mathcal{D}) \iff D_C + N_Z = \mathcal{D}. \quad (3.16)$$

Equivalently, every formal coefficient direction has a supported representative modulo N_Z . The size of $\mathcal{C}$ alone does not determine this property: two supports of the same cardinality can generate coefficient subspaces in different positions relative to N_Z , and hence give different retained widths. Appendix B gives polynomial-evaluation examples of this phenomenon.

4 Quadratic Folding on Restricted Domains

We specialize the boundary model to quadratic folds of two source records. Throughout this section, $\mathbb{F}$ is a field and X_0, X_1 are nonempty admissible source domains. We use a direct boundary, in which every scalar function of one source alone is free:

$$\mathcal{F} = \text{Sep}(X_0, X_1) := \{a(x_0) + b(x_1) : a : X_0 \rightarrow \mathbb{F}, b : X_1 \rightarrow \mathbb{F}\}. \quad (4.1)$$

We instantiate Section 3 with $\mathcal{F} = \text{Sep}(X_0, X_1)$ and $K = 0$, and write $m_{\min}^{\text{dir}}$ for the resulting retained width.

4.1 Source-Local Functions and Rectangle Differences

Call $f : X_0 \times X_1 \rightarrow \mathbb{F}$ source-local if $f \in \text{Sep}(X_0, X_1)$. For $x_0^+, x_0^- \in X_0$ and $x_1^+, x_1^- \in X_1$, define

$$\begin{aligned} \square f(x_0^+, x_0^-; x_1^+, x_1^-) &:= f(x_0^+, x_1^+) - f(x_0^+, x_1^-) \\ &\quad - f(x_0^-, x_1^+) + f(x_0^-, x_1^-). \end{aligned} \quad (4.2)$$

Lemma 4.1 (Rectangle criterion). *A function $f : X_0 \times X_1 \rightarrow \mathbb{F}$ is source-local if and only if all of its rectangle differences vanish.*

Proof. If $f(x_0, x_1) = a(x_0) + b(x_1)$, the four terms in (4.2) cancel.

Conversely, suppose all rectangle differences vanish. Fix $x_0^* \in X_0$ and $x_1^* \in X_1$, and set $a(x_0) := f(x_0, x_1^*) - f(x_0^*, x_1^*)$ and $b(x_1) := f(x_0^*, x_1)$. Applying (4.2) to $(x_0, x_0^*; x_1, x_1^*)$ gives $f(x_0, x_1) = a(x_0) + b(x_1)$. $\square$

Write $[f]_{\text{sep}}$ for the class of f modulo $\text{Sep}(X_0, X_1)$; by Lemma 4.1, this class vanishes exactly when all rectangle differences vanish.

4.2 Restricted Quadratic Interactions

Direct quadratic folds. Let Y be a finite-dimensional $\mathbb{F}$ -vector space and consider

$$H_c(x_0, x_1) = E_0(x_0) + cT(x_0, x_1) + c^2E_1(x_1), \quad c \in \mathcal{C} \subseteq \mathbb{F}, \quad (4.3)$$

where $E_0 : X_0 \rightarrow Y$, $E_1 : X_1 \rightarrow Y$, and $T : X_0 \times X_1 \rightarrow Y$. The first and last terms depend on one source only. Hence, for every $\eta \in Y^\vee$,

$$[\eta \circ H_c]_{\text{sep}} = c[\eta \circ T]_{\text{sep}}. \quad (4.4)$$

Define $\Gamma_T : Y^\vee \rightarrow \mathbb{F}^{X_0 \times X_1} / \text{Sep}(X_0, X_1)$ by $\Gamma_T(\eta) := [\eta \circ T]_{\text{sep}}$.

Proposition 4.2 (Direct quadratic folding). *If the challenge support $\mathcal{C}$ contains a nonzero element, then*

$$m_{\min}^{\text{dir}} = \text{rank } \Gamma_T. \quad (4.5)$$

Proof. Equation (4.4) places every supported residual direction in $\text{im } \Gamma_T$. Conversely, for any $c_0 \in \mathcal{C} \setminus \{0\}$, the classes $[\eta \circ H_{c_0}]_{\text{sep}}$, as η varies over $Y^\vee$, span the same space as the classes $[\eta \circ T]_{\text{sep}}$. Thus the residual target is $\text{im } \Gamma_T$, and Theorem 3.1 gives (4.5). $\square$

Restricted R1CS. We next apply this calculation to the relaxed-R1CS fold used by Nova [12]. Let $A, B, C \in \mathbb{F}^{m \times n}$, and let an assignment $Z \in \mathbb{F}^n$ contain the relaxation coordinate u . Write $e_u \in \mathbb{F}^n$ for the corresponding standard basis vector, so $u = e_u^\top Z$. The relaxed residual is

$$E(Z) = AZ \odot BZ - uCZ, \quad (4.6)$$

where $\odot$ denotes coordinatewise multiplication. For two assignments Z_0, Z_1 ,

$$E(Z_0 + cZ_1) = E(Z_0) + cT_R(Z_0, Z_1) + c^2E(Z_1), \quad (4.7)$$

where

$$T_R(Z_0, Z_1) = AZ_0 \odot BZ_1 + AZ_1 \odot BZ_0 - u_0CZ_1 - u_1CZ_0. \quad (4.8)$$

Identify $(\mathbb{F}^m)^\vee$ with $\mathbb{F}^m$ through the standard pairing. For $\eta \in \mathbb{F}^m$, let $D_\eta := \text{diag}(\eta)$ and define

$$M_\eta := A^\top D_\eta B + B^\top D_\eta A - e_u \eta^\top C - C^\top \eta e_u^\top. \quad (4.9)$$

By construction, $\eta^\top T_R(Z_0, Z_1) = Z_0^\top M_\eta Z_1$, so every scalarized cross term is represented by a bilinear form.

Let $Z_0, Z_1 \subseteq \mathbb{F}^n$ be the nonempty admissible assignment sets, with $X_0 = Z_0$ and $X_1 = Z_1$. Define their variation spaces by

$$\Delta_b := \text{span}_{\mathbb{F}}\{z - z' : z, z' \in Z_b\}, \quad b \in \{0, 1\}. \quad (4.10)$$

Choose $P_b \in \mathbb{F}^{n \times d_b}$ whose columns form a basis of Δ_b . For an admissible rectangle, bilinearity gives the rectangle difference $(Z_0^+ - Z_0^-)^\top M_\eta (Z_1^+ - Z_1^-)$. Since the admissible differences span Δ_0 and Δ_1 , Lemma 4.1 yields

$$[\eta^\top T_R]_{\text{sep}} = 0 \iff P_0^\top M_\eta P_1 = 0. \quad (4.11)$$

Define $\Gamma_{\text{R1CS}} : \mathbb{F}^m \rightarrow \mathbb{F}^{d_0 \times d_1}$ by $\Gamma_{\text{R1CS}}(\eta) := P_0^\top M_\eta P_1$.

Theorem 4.3 (Restricted-R1CS width). *Assume that the challenge support contains a nonzero element. Then*

$$m_{\min}^{\text{dir}} = \text{rank } \Gamma_{\text{R1CS}} = \dim_{\mathbb{F}} \text{span} \{P_0^\top M_\eta P_1 : \eta \in \mathbb{F}^m\}. \quad (4.12)$$

Proof. By Proposition 4.2, $m_{\min}^{\text{dir}}$ is the rank of $\eta \mapsto [\eta^\top T_R]_{\text{sep}}$. Equation (4.11) shows that this map and Γ_{R1CS} have the same kernel, so rank-nullity gives (4.12). $\square$

Thus the retained width is the dimension of the restricted matrix space in (4.12), not the maximum rank of one of its matrices.

4.3 Computing a Minimum Representation

The same calculation also yields a minimum-width retained representation. Let $e_1, \dots, e_m$ be the standard basis of $\mathbb{F}^m$, and set $N_i := P_0^\top M_{e_i} P_1$, so $P_0^\top M_\eta P_1 = \sum_{i=1}^m \eta_i N_i$. Vectorizing gives

$$\mathbf{N} := [\text{vec}(N_1) \cdots \text{vec}(N_m)], \quad m_{\min}^{\text{dir}} = \text{rank } \mathbf{N}. \quad (4.13)$$

Let $d := \text{rank } \mathbf{N}$, and choose pivot indices $i_1, \dots, i_d$ such that $N_{i_1}, \dots, N_{i_d}$ form a basis of $\text{span}\{N_1, \dots, N_m\}$. For each $\ell \in \{1, \dots, m\}$, write

$$N_\ell = \sum_{j=1}^d \alpha_{\ell,j} N_{i_j}. \quad (4.14)$$

Corollary 4.4 (Minimum retained coordinates). *For every $\ell \in \{1, \dots, m\}$, there exist source-local functions $a_\ell : Z_0 \rightarrow \mathbb{F}$ and $b_\ell : Z_1 \rightarrow \mathbb{F}$ such that*

$$(T_R)_\ell = \sum_{j=1}^d \alpha_{\ell,j} (T_R)_{i_j} + a_\ell(Z_0) + b_\ell(Z_1). \quad (4.15)$$

Consequently, $(T_R)_{i_1}, \dots, (T_R)_{i_d}$ form a retained state of minimum width.

Proof. For each ℓ , let $\eta_\ell = e_\ell - \sum_j \alpha_{\ell,j} e_{i_j}$. By (4.14), $P_0^\top M_{\eta_\ell} P_1 = 0$; hence (4.11) and Lemma 4.1 imply that $\eta_\ell^\top T_R$ is source-local, which is exactly (4.15). Theorem 4.3 gives optimality. $\square$

Given P_0, P_1 , Gaussian elimination on $\mathbf{N}$ computes the width and selects pivot coordinates for a minimum-width representation.

4.4 Coded Multiplication and Reed–Solomon Domains

Consider $E(x, y, z, u) = x \odot y - uz$ on the product domain $\mathcal{X}(U, V, W) := U \times V \times W \times \mathbb{F}$, where $U, V, W \leq \mathbb{F}^m$, and instantiate the direct-boundary model with $X_0 = X_1 = \mathcal{X}(U, V, W)$. The components vary freely over the product domain; this is a semantic-domain condition, not a probabilistic independence assumption.

For $r_0 = (x_0, y_0, z_0, u_0)$ and $r_1 = (x_1, y_1, z_1, u_1)$,

$$\begin{aligned} E(r_0 + cr_1) &= E(r_0) + cT_{\text{cross}}(r_0, r_1) + c^2E(r_1), \\ T_{\text{cross}}(r_0, r_1) &= x_0 \odot y_1 + x_1 \odot y_0 - u_0z_1 - u_1z_0. \end{aligned} \quad (4.16)$$

For subspaces $U, V \leq \mathbb{F}^m$, write $U \star V := \text{span}_{\mathbb{F}}\{x \odot y : x \in U, y \in V\}$ for their Schur product. Every value of T_{cross} lies in $U \star V + W$.

Theorem 4.5 (Coded-multiplication width). *Assume that the challenge support contains a nonzero element. On the product domain $\mathcal{X}(U, V, W)$,*

$$m_{\min}^{\text{dir}} = \dim_{\mathbb{F}}(U \star V + W). \quad (4.17)$$

Proof. By Proposition 4.2, it suffices to identify the kernel of $\eta \mapsto [\eta^\top T_{\text{cross}}]_{\text{sep}}$. Let $S := U \star V + W$. If $\eta \in S^\perp$, then $\eta^\top T_{\text{cross}} \equiv 0$. Conversely, suppose $[\eta^\top T_{\text{cross}}]_{\text{sep}} = 0$. Since $T_{\text{cross}}(r, 0) = T_{\text{cross}}(0, r) = 0$, Lemma 4.1 applied to the rectangle $(r_0, 0; r_1, 0)$ gives $\eta^\top T_{\text{cross}}(r_0, r_1) = 0$ for all r_0, r_1 .

Taking $r_0 = (x, 0, 0, 0)$ and $r_1 = (0, y, 0, 0)$ gives $\eta \perp U \star V$, while $r_0 = (0, 0, 0, 1)$ and $r_1 = (0, 0, z, 0)$ gives $\eta \perp W$. Hence the kernel is $(U \star V + W)^\perp$, so the map has rank $\dim_{\mathbb{F}}(U \star V + W)$. Proposition 4.2 gives (4.17). $\square$

Retaining the coordinates of T_{cross} with respect to a basis of $U \star V + W$ gives a matching realization.

Reed–Solomon domains. Fix distinct evaluation points $\alpha = (\alpha_1, \dots, \alpha_m) \in \mathbb{F}^m$, and write $\text{ev}_\alpha(f) := (f(\alpha_1), \dots, f(\alpha_m))$. For $1 \leq r \leq m$, define

$$\text{RS}_r(\alpha) := \{\text{ev}_\alpha(f) : \deg f < r\}. \quad (4.18)$$

For positive integers a, b satisfying $a + b - 1 \leq m$, the standard Reed–Solomon star-product identity [17] gives

$$\text{RS}_a(\alpha) \star \text{RS}_b(\alpha) = \text{RS}_{a+b-1}(\alpha). \quad (4.19)$$

Taking $U = \text{RS}_a(\alpha)$, $V = \text{RS}_b(\alpha)$, and $W = \text{RS}_{a+b-1}(\alpha)$, Theorem 4.5 and (4.19) give

$$m_{\min}^{\text{dir}} = a + b - 1. \quad (4.20)$$

This minimum has a matching polynomial representation. For $j \in \{0, 1\}$, write $x_j = \text{ev}_\alpha(f_j)$, $y_j = \text{ev}_\alpha(g_j)$, and $z_j = \text{ev}_\alpha(h_j)$, where $\deg f_j < a$, $\deg g_j < b$, and $\deg h_j < a + b - 1$. Then $T_{\text{cross}} = \text{ev}_\alpha(t)$, where

$$t(X) = f_0(X)g_1(X) + f_1(X)g_0(X) - u_0h_1(X) - u_1h_0(X). \quad (4.21)$$

Since $\deg t < a + b - 1$, its $a + b - 1$ coefficients form a valid retained state: evaluation at the fixed points is linear in these coefficients, and (4.20) proves optimality.

On the full product domain $U = V = W = \mathbb{F}^m$, one has $U \star V = \mathbb{F}^m$, so the direct width is m . Now suppose k^2 distinct evaluation points are available, and set $m = k^2$ and $a = b = k$. Then

$$m_{\min}^{\text{dir}} = \begin{cases} k^2, & \text{on the full product domain,} \\ 2k - 1, & \text{on the Reed-Solomon domain.} \end{cases} \quad (4.22)$$

The defining coordinatewise equation is the same in both cases; only the admissible source domain changes.

5 Reconstruction Below the Exact Width

Below the exact width, reconstruction becomes a distributional problem. Throughout this section, let $\mathbb{F} = \mathbb{F}_q$. Let U, V, Y be finite-dimensional $\mathbb{F}$ -vector spaces and let $\mu : U \times V \rightarrow Y$ be bilinear. Replacing Y by $\text{span}_{\mathbb{F}} \mu(U, V)$ if necessary, assume that $\mu(U, V)$ spans Y , and write $d_Y := \dim_{\mathbb{F}} Y$. Fix a nonempty challenge support $\mathcal{C} \subseteq \mathbb{F}_q^\times$.

5.1 The Reconstruction Model

Draw $x_0, x_1 \in U$ and $y_0, y_1 \in V$ independently and uniformly, and write $R_0 = (x_0, y_0)$ and $R_1 = (x_1, y_1)$. For $c \in \mathcal{C}$, the required folded output is

$$H_c^\mu := \mu(x_0 + cx_1, y_0 + cy_1), \quad (5.1)$$

with expansion

$$H_c^\mu = \mu(x_0, y_0) + cT_\mu + c^2\mu(x_1, y_1), \quad T_\mu := \mu(x_0, y_1) + \mu(x_1, y_0). \quad (5.2)$$

We use the direct boundary on the two source records R_0, R_1 : arbitrary functions of either record alone are free. For $D \geq 0$, a width- D evaluator has one prechallenge map $\Psi : (U \times V)^2 \rightarrow \mathbb{F}^D$ that must work for the entire challenge

support. After receiving c , it may use arbitrary one-source maps $g_{0,c}, g_{1,c} : U \times V \rightarrow Y$ and a linear map $L_c : \mathbb{F}^D \rightarrow Y$, producing

$$\hat{H}_c = g_{0,c}(R_0) + g_{1,c}(R_1) + L_c\Psi(R_0, R_1). \quad (5.3)$$

For an evaluator E , let $p_c(E) := \Pr[\hat{H}_c = H_c^\mu]$, and define

$$P_\mu^*(D) := \max_E \min_{c \in \mathcal{C}} p_c(E). \quad (5.4)$$

The map Ψ may be computed arbitrarily before the challenge; only its post-challenge readout is required to be linear.

Since $\mu(U, V)$ spans Y , Proposition 4.2 gives exact direct width d_Y . The uniform distribution on the source records has full support; hence probability-one reconstruction at every supported challenge is equivalent to exact correctness on the whole semantic domain. Thus d_Y is also the first budget at which probability-one reconstruction is possible.

For a scalar bilinear form $B : U \times V \rightarrow \mathbb{F}$, write $\text{rank } B$ for the rank of the associated map $U \rightarrow V^\vee$. For $W \leq Y$, write $W^\perp := \{\lambda \in Y^\vee : \lambda(w) = 0 \text{ for every } w \in W\}$.

5.2 Reconstruction Without Side Information

The next estimate controls the one-source corrections allowed by the reconstruction model.

Lemma 5.1 (Bilinear character estimate). *Let U', V' be finite-dimensional $\mathbb{F}_q$ -vector spaces, let $\beta : U' \times V' \rightarrow \mathbb{F}_q$ be bilinear of rank r , and fix a nontrivial additive character $\chi : \mathbb{F}_q \rightarrow \mathbb{C}^\times$. For arbitrary functions $\phi : U' \rightarrow \mathbb{C}$ and $\psi : V' \rightarrow \mathbb{C}$ satisfying $|\phi|, |\psi| \leq 1$,*

$$|\mathbb{E}_{u,v} \phi(u) \psi(v) \chi(\beta(u, v))| \leq q^{-r/2}. \quad (5.5)$$

In particular,

$$\mathbb{E}_{u,v} \chi(\beta(u, v)) = q^{-r}. \quad (5.6)$$

Consequently, for $W \leq Y$ and independent uniform $u \in U, v \in V$,

$$\Pr[\mu(u, v) \in W] = q^{-\text{codim } W} \sum_{\lambda \in W^\perp} q^{-\text{rank}(\lambda \circ \mu)}. \quad (5.7)$$

Proof. Put $S := \mathbb{E}_{u,v} \phi(u) \psi(v) \chi(\beta(u, v))$. Cauchy–Schwarz in u gives

$$|S|^2 \leq \mathbb{E}_u |\mathbb{E}_v \psi(v) \chi(\beta(u, v))|^2 = \mathbb{E}_{v,v'} \psi(v) \overline{\psi(v')} \mathbb{E}_u \chi(\beta(u, v - v')).$$

Character orthogonality makes the inner average vanish unless $v - v'$ lies in the right kernel of β . Taking absolute values on the remaining terms gives

$|S|^2 \leq \Pr[v - v' \in \ker_R \beta] = q^{-r}$, proving (5.5). Taking $\phi = \psi = 1$ and averaging first over u gives (5.6). Finally, character orthogonality on Y/W gives $1_W(y) = q^{-\text{codim } W} \sum_{\lambda \in W^\perp} \chi(\lambda(y))$. Substituting $y = \mu(u, v)$ and applying (5.6) proves (5.7). $\square$

Theorem 5.2 (Reconstruction without side information). *For every $D \geq 0$,*

$$P_\mu^*(D) = \max_{\substack{W \leq Y \\ \dim W \leq D}} \Pr[\mu(u, v) \in W], \quad (5.8)$$

where $u \in U$ and $v \in V$ are independent and uniform. An optimum is attained by retaining coordinates of a fixed linear projection of T_μ .

Proof. Fix a width- D evaluator E and $c \in \mathcal{C}$, and put $W := \text{im } L_c$, so $\dim W \leq D$. Whenever the evaluator succeeds,

$$H_c^\mu - g_{0,c}(R_0) - g_{1,c}(R_1) \in W. \quad (5.9)$$

For $\lambda \in W^\perp$, put $\beta_\lambda := \lambda \circ \mu$, and let B_λ be its matrix in fixed bases of U and V . The part coupling the two source records is $c(\beta_\lambda(x_0, y_1) + \beta_\lambda(x_1, y_0))$. As a bilinear form between $R_0, R_1 \in U \oplus V$,

$$\text{rank} \begin{pmatrix} 0 & cB_\lambda \\ cB_\lambda^\top & 0 \end{pmatrix} = 2 \text{rank } B_\lambda = 2 \text{rank } \beta_\lambda. \quad (5.10)$$

Because $c \neq 0$, this identity holds in every characteristic. After scalarization, the one-source terms in (5.9) contribute bounded phases. Expanding the indicator of membership in W over $W^\perp$ and applying Lemma 5.1 termwise gives

$$p_c(E) \leq q^{-\text{codim } W} \sum_{\lambda \in W^\perp} q^{-\text{rank}(\lambda \circ \mu)} = \Pr[\mu(u, v) \in W]. \quad (5.11)$$

This holds for every supported c , proving the upper bound in (5.8).

Conversely, fix a maximizing $W \leq Y$, $\dim W \leq D$, and choose a projection $\Pi : Y \rightarrow W$. Retain the coordinates of ΠT_μ . After receiving c , output $\Pi \mu(x_0, y_0) + c \Pi T_\mu + c^2 \Pi \mu(x_1, y_1) = \Pi H_c^\mu$. Reconstruction succeeds exactly when $H_c^\mu \in W$. For every supported c , the folded factors $x_0 + cx_1$ and $y_0 + cy_1$ remain independent and uniform, so the success probability is $\Pr[\mu(u, v) \in W]$. $\square$

For $0 \leq D \leq d_Y$, a maximizing W may be taken to have dimension exactly D . Taking $N = W^\perp$ in (5.7) gives

$$P_\mu^*(D) = \max_{\substack{N \leq Y^\vee \\ \dim N = d_Y - D}} q^{-(d_Y - D)} \sum_{\lambda \in N} q^{-\text{rank}(\lambda \circ \mu)}. \quad (5.12)$$

For $D \geq d_Y$, exact reconstruction gives $P_\mu^*(D) = 1$.

5.3 Reconstruction with Bilinear Side Information

Suppose now that the incoming state contains mixed information built from the same source factors. Let $\nu : U \times V \rightarrow P$ be bilinear, and suppose the actual incoming mixed state is

$$K := T_\nu = \nu(x_0, y_1) + \nu(x_1, y_0). \quad (5.13)$$

For $c \in \mathcal{C}$, define $H_c^\nu := \nu(x_0 + cx_1, y_0 + cy_1)$. Since $H_c^\nu = \nu(x_0, y_0) + cK + c^2\nu(x_1, y_1)$, it is computable from the incoming state and one-source information.

A width- D evaluator may additionally apply a c -dependent linear map $M_c : P \rightarrow Y$ to K , producing $\hat{H}_c = g_{0,c}(R_0) + g_{1,c}(R_1) + M_c K + L_c \Psi(R_0, R_1)$. Let $P_{\mu|\nu}^*(D)$ denote the corresponding optimum.

Define $T := \{\lambda \circ \mu : \lambda \in Y^\vee\}$, $S := \{\eta \circ \nu : \eta \in P^\vee\}$, and

$$\mathcal{B} := T + S, \quad h := \dim(\mathcal{B}/S) = \dim T - \dim(T \cap S). \quad (5.14)$$

Theorem 5.3 (Reconstruction with bilinear side information). *For every $D \geq 0$,*

$$P_{\mu|\nu}^*(D) = \max_{\substack{M \in \text{Hom}_{\mathbb{F}}(P, Y) \\ W \leq Y, \dim W \leq D}} \Pr[(\mu - M\nu)(u, v) \in W], \quad (5.15)$$

where u, v are independent and uniform.

The exact remaining width is h . For $0 \leq D \leq h$, put $s := h - D$. Then

$$P_{\mu|\nu}^*(D) = \max_{\substack{N \leq \mathcal{B} \\ \dim N = s, N \cap S = \{0\}}} q^{-s} \sum_{B \in N} q^{-\text{rank } B}. \quad (5.16)$$

For $D \geq h$, $P_{\mu|\nu}^*(D) = 1$.

Proof. Primal form. Fix an evaluator and $c \in \mathcal{C}$. Let $W := \text{im } L_c$, let $\tau : Y \rightarrow Y/W$ be the quotient map, and put $M := c^{-1}M_c$. By the expansion of H_c^ν , the difference between $M_c K$ and MH_c^ν consists only of one-source terms. Thus, whenever the evaluator succeeds,

$$\tau(H_c^\mu - MH_c^\nu) = a_c(R_0) + b_c(R_1) \quad (5.17)$$

for suitable one-source maps $a_c, b_c : U \times V \rightarrow Y/W$.

The image of $\tau \circ (\mu - M\nu)$ need not span all of Y/W . Restricting its codomain to the span of its image gives a bilinear map that spans its codomain, while (5.17) remains a valid zero-width reconstruction whenever the original evaluator succeeds. Applying the $D = 0$ case of Theorem 5.2 therefore gives $p_c(E) \leq \Pr[(\mu - M\nu)(u, v) \in W]$. Taking the minimum over c and then maximizing over evaluators proves the upper bound in (5.15).

Conversely, fix $M : P \rightarrow Y$, a subspace $W \leq Y$ with $\dim W \leq D$, and a projection $\Pi : Y \rightarrow W$. Retain the coordinates of $\Pi(T_\mu - MT_\nu)$. After receiving c , output

$$MH_c^\nu + \Pi(H_c^\mu - MH_c^\nu). \quad (5.18)$$

The difference $H_c^\mu - MH_c^\nu$ equals $(\mu - M\nu)(x_0, y_0) + c(T_\mu - MT_\nu) + c^2(\mu - M\nu)(x_1, y_1)$, so the projection in (5.18) is computable from one-source information and the retained state. Reconstruction succeeds exactly when $(\mu - M\nu)(x_0 + cx_1, y_0 + cy_1) \in W$. The folded factors remain independent and uniform, so this state attains the probability in (5.15).

Exact remaining width. For a scalar bilinear form $B : U \times V \rightarrow \mathbb{F}$, define $C_B(R_0, R_1) := B(x_0, y_1) + B(x_1, y_0)$. For $u \in U$ and $v \in V$, take the rectangle with $R_0 \in \{0, (u, 0)\}$ and $R_1 \in \{0, (0, v)\}$. Its rectangle difference on C_B is $B(u, v)$. Hence C_B is source-local only if $B = 0$. The target contributes T , while the incoming crossed state contributes S . The residual mixed space is therefore $\mathcal{B}/S$, and Theorem 3.1 gives exact remaining width h .

Compatible representatives. We now place the target and side-information scalar forms in a common ambient space. This separates the target directions, those already supplied through ν , and the missing directions that must be represented compatibly.

Put $Z := \mathcal{B}^\vee$, and define

$$\zeta(u, v)(B) := B(u, v), \quad B \in \mathcal{B}. \quad (5.19)$$

The vectors $\zeta(u, v)$ span Z : a form in $Z^\vee = \mathcal{B}$ that annihilates all of them vanishes identically.

Define $J_\mu : Z \rightarrow Y$ and $J_\nu : Z \rightarrow P$ by

$$\lambda(J_\mu z) = z(\lambda \circ \mu), \quad \eta(J_\nu z) = z(\eta \circ \nu). \quad (5.20)$$

Then $\mu = J_\mu \zeta$ and $\nu = J_\nu \zeta$. For $J := (J_\mu, J_\nu) : Z \rightarrow Y \oplus P$, the dual map has image $\mathcal{B} = Z^\vee$, so J is injective.

Let $A := J_\nu$, define $H_c^\zeta := \zeta(x_0 + cx_1, y_0 + cy_1)$, and let $T_\zeta := \zeta(x_0, y_1) + \zeta(x_1, y_0)$. Then

$$H_c^\mu = J_\mu H_c^\zeta, \quad H_c^\nu = A H_c^\zeta, \quad A T_\zeta = T_\nu = K. \quad (5.21)$$

Because J is injective, fix a linear left inverse $R_J : Y \oplus P \rightarrow Z$ with $R_J J = I_Z$. The folded ν -output H_c^ν is exactly computable from K and one-source information. Hence an evaluator for H_c^μ yields one for H_c^ζ with the same retained budget and at least the same success probability, by applying R_J to $(\hat{H}_c^\mu, H_c^\nu)$. Conversely, applying J_μ to an evaluator for H_c^ζ gives an evaluator for H_c^μ , again without decreasing the success probability. The two reconstruction optima therefore coincide.

Applying the primal form to target ζ with side information $A\zeta$ gives

$$P_{\mu|\nu}^*(D) = \max_{\substack{\widetilde{M}: P \rightarrow Z \\ W_Z \leq Z, \dim W_Z \leq D}} \Pr[(I_Z - \widetilde{M}A)\zeta(u, v) \in W_Z]. \quad (5.22)$$

For fixed $(\widetilde{M}, W_Z)$, put $F := I_Z - \widetilde{M}A$ and $L := F^{-1}(W_Z)$. Since F is the identity on $\ker A$, we have $L \cap \ker A = W_Z \cap \ker A$, and hence $\dim(L \cap \ker A) \leq D$.

Conversely, let $L \leq Z$ satisfy $\dim(L \cap \ker A) \leq D$, and put $J_0 := L \cap \ker A$. Choose a complement C of J_0 in L , and then a complement C' such that $Z = C \oplus \ker A \oplus C'$. Define $F : Z \rightarrow Z$ to be zero on C and the identity on $\ker A \oplus C'$. Then $F^{-1}(J_0) = L$. Moreover, $I_Z - F$ vanishes on $\ker A$, so it factors through $A : Z \rightarrow \text{im } A$. Extending the resulting map from $\text{im } A$ to P gives $F = I_Z - \widetilde{M}A$ for some $\widetilde{M} : P \rightarrow Z$. Taking $W_Z = J_0$ realizes L . Therefore

$$P_{\mu|\nu}^*(D) = \max_{\substack{L \leq Z \\ \dim(L \cap \ker A) \leq D}} \Pr[\zeta(u, v) \in L]. \quad (5.23)$$

Assume $0 \leq D \leq h$. Starting from any feasible L , first adjoin vectors from $\ker A$ until $\dim(L \cap \ker A) = D$, and then adjoin lifts of a complement of $A(L)$ in $\text{im } A$. The second step leaves $L \cap \ker A$ unchanged, and neither enlargement decreases the probability mass. We may therefore assume $\dim(L \cap \ker A) = D$ and $A(L) = \text{im } A$.

Since $\text{im } A^\vee = S$, we have $\text{rank } A = \dim S$, and hence $\dim L = \dim S + D$. Put $N := L^\perp \leq Z^\vee = \mathcal{B}$. Then $\dim N = h - D = s$. Moreover, $S = (\ker A)^\perp$, while $A(L) = \text{im } A$ is equivalent to $L + \ker A = Z$; hence $N \cap S = \{0\}$. Conversely, every $N \leq \mathcal{B}$ with $\dim N = s$ and $N \cap S = \{0\}$ gives such an $L = N^\perp$. Character orthogonality now gives the rank sum in (5.16), completing the proof. $\square$

The condition $N \cap S = \{0\}$ has a useful interpretation. Let $p : \mathcal{B} \rightarrow \mathcal{B}/S$ be the quotient map. For feasible N , $p|_N$ is injective, so its inverse on $p(N)$ chooses representatives of several missing quotient directions simultaneously. These representatives must come from one linear section on the missing subspace $p(N)$; their ranks cannot in general be optimized independently.

If $h > 0$, let $r_{\min}(\mu | \nu) := \min_{B \in \mathcal{B} \setminus S} \text{rank } B$. For one missing direction, (5.16) gives

$$P_{\mu|\nu}^*(h-1) = \frac{1}{q} + \left(1 - \frac{1}{q}\right) q^{-r_{\min}(\mu|\nu)}. \quad (5.24)$$

For later use, write $\vartheta_q := (2q-1)/q^2$, the zero probability of a nonzero rank-one scalar bilinear form on independent uniform inputs. Several missing directions need not be determined by $r_{\min}$ alone.

5.4 Same Exact Width, Different Profiles

Example 5.4 (Same exact width, different profiles).

Let $U = V = \mathbb{F}_q^2$, let $\mu(u, v) = (u_1 v_1, u_2 v_2)$, and compare

$$\nu_S(u, v) = u_1 v_2 + u_2 v_1, \quad \nu_T(u, v) = (u_1 + u_2) v_2. \quad (5.25)$$

Using the standard bases, identify scalar bilinear forms with 2×2 matrices. For ν_S ,

Table 1. Reconstruction profiles with the same exact remaining width 2.

Side information	$D = 0$	$D = 1$	$D = 2$
ν_S	ϑ_q^2	ϑ_q	1
ν_T	ξ_q	ϑ_q	1

$$\mathcal{B}_S = \left\{ \begin{pmatrix} \alpha & \beta \\ \beta & \gamma \end{pmatrix} : \alpha, \beta, \gamma \in \mathbb{F}_q \right\}, \quad S_S = \mathbb{F}_q \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad (5.26)$$

whereas for ν_T ,

$$\mathcal{B}_T = \left\{ \begin{pmatrix} \alpha & \beta \\ 0 & \gamma + \beta \end{pmatrix} : \alpha, \beta, \gamma \in \mathbb{F}_q \right\}, \quad S_T = \mathbb{F}_q \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}. \quad (5.27)$$

Both quotient spaces have dimension two. Set

$$\xi_q := \frac{q^2 + q - 1}{q^3}. \quad (5.28)$$

Their optimal reconstruction profiles are:

For the symmetric space $\mathcal{B}_S$, a two-dimensional subspace contains at most two projective rank-one directions. Indeed, if two distinct such directions are represented by $\gamma aa^\top$ and $\delta bb^\top$, then a and b are linearly independent, and every linear combination using both directions has rank two. The diagonal plane is disjoint from S_S and attains this bound, giving $P_{\mu|\nu_S}^*(0) = \vartheta_q^2$.

For $\mathcal{B}_T$, the first-row plane is disjoint from S_T , and every nonzero matrix in it has rank one. It therefore gives $P_{\mu|\nu_T}^*(0) = \xi_q$, which is optimal because every nonzero bilinear form has rank at least one.

Both examples contain a rank-one form outside their side-information space, so (5.24) gives the same value ϑ_q at $D = 1$. Exact reconstruction gives value 1 at $D = 2$. Since $\xi_q - \vartheta_q^2 = (q-1)^3/q^4 > 0$, equal exact remaining width does not determine the below-threshold reconstruction profile.

In characteristic two, every nonzero quotient class in both examples has a rank-one representative, yet only the triangular example admits a linear section attaining these ranks simultaneously. Appendix C gives the complete rank distributions and the full compatibility analysis.

5.5 Polynomial Multiplication Near the Threshold

Consider polynomial multiplication $\mu(f, g) = fg$ for $f, g \in \mathbb{F}_q[t]_{<k}$, with $Y = \mathbb{F}_q[t]_{<2k-1}$ and $d_Y = 2k - 1$. For $\lambda \in Y^\vee$, let B_λ be the matrix of the scalar bilinear form $(f, g) \mapsto \lambda(fg)$ in the monomial bases. Its entries are $(B_\lambda)_{ij} = \lambda(t^{i+j})$ for $0 \leq i, j < k$. These matrices form the full $(2k - 1)$ -dimensional space of symmetric $k \times k$ Hankel matrices. Formula (5.12) therefore reduces reconstruction near the exact threshold to a rank-sum optimization over low-dimensional Hankel subspaces.

Theorem 5.5 (Near-threshold polynomial multiplication).

For every prime power q and $0 \leq s \leq \min\{3, k\}$,

$$P_\mu^*(d_Y - s) = \vartheta_q^s. \quad (5.29)$$

Proof. The case $s = 0$ is exact reconstruction. If $k = 1$, the only remaining case is $P_\mu^*(0) = \Pr[fg = 0] = \vartheta_q$. Assume henceforth $k \geq 2$. By (5.12), for deficit s we must maximize $q^{-s} \sum_{\lambda \in N} q^{-\text{rank } B_\lambda}$ over s -dimensional subspaces $N \leq Y^\vee$.

For $\alpha \in \mathbb{F}_q$, define $\ell_\alpha(p) := p(\alpha)$, and define $\ell_\infty(p) := [t^{2k-2}]p$. Then

$$B_{\ell_\alpha} = v_\alpha v_\alpha^\top, \quad v_\alpha = (1, \alpha, \dots, \alpha^{k-1})^\top, \quad (5.30)$$

$$B_{\ell_\infty} = e_{k-1} e_{k-1}^\top. \quad (5.31)$$

These are exactly the projective rank-one Hankel directions; Appendix D proves the converse.

If $\alpha_1, \dots, \alpha_s \in \mathbb{F}_q \cup \{\infty\}$ are distinct and $s \leq k$, the corresponding vectors are linearly independent. Hence

$$\text{rank} \left(\sum_{i=1}^s b_i B_{\ell_{\alpha_i}} \right) = |\{i : b_i \neq 0\}|. \quad (5.32)$$

For their span N ,

$$q^{-s} \sum_{\lambda \in N} q^{-\text{rank } B_\lambda} = q^{-s} \sum_{b \in \mathbb{F}_q^s} q^{-|\text{supp } b|} = \vartheta_q^s. \quad (5.33)$$

Thus evaluation directions attain the value in (5.29).

For $s = 1$, every nonzero matrix has rank at least one, so ϑ_q is optimal.

For $s = 2$, a two-dimensional space of symmetric matrices contains at most two projective rank-one directions. Indeed, if two independent rank-one directions are represented by $\gamma a a^\top$ and $\delta b b^\top$, then a, b are linearly independent, and every combination with both coefficients nonzero has rank two. The resulting rank count gives an upper bound of ϑ_q^2 , matching (5.33).

It remains to treat $s = 3$, so $k \geq 3$. We use only the following consequence of Meshulam's theorem [14]: the $k \times k$ Hankel space contains a $(2k-3)$ -dimensional subspace whose nonzero matrices all have rank at least three. Since the ambient Hankel space has dimension $2k-1$, every three-dimensional Hankel subspace intersects this subspace nontrivially. Thus every three-dimensional Hankel subspace contains a matrix of rank at least three.

Let N be three-dimensional. If N contains at least three projective rank-one directions, the preceding two-dimensional bound implies that any three distinct such directions are linearly independent and hence form a basis of N . Write representatives as $\gamma_i v_i v_i^\top$, $1 \leq i \leq 3$. The vectors v_1, v_2, v_3 must also be linearly independent; otherwise every matrix in N would have rank at most two, contradicting the preceding paragraph. Thus the rank of a linear combination is exactly the number of nonzero coefficients, and (5.33) gives ϑ_q^3 .

Now suppose that N contains at most two projective rank-one directions. Choose a matrix in N of rank at least three and a 3×3 minor that is nonzero on it. The determinant of this minor restricts to a nonzero homogeneous cubic on the projective plane of N . Every projective direction of rank at most two is a zero of this cubic. Appendix D gives the required point bound: there are at most $3q + 1$ such directions.

Let $t \leq 2$ be the number of projective rank-one directions and $u \leq 3q + 1$ the number of projective directions of rank at most two. Since $|\mathbb{P}^2(\mathbb{F}_q)| = q^2 + q + 1$,

$$q^{-3} \sum_{\lambda \in N} q^{-\text{rank } B_\lambda} \leq q^{-3} \left[1 + (q-1)(tq^{-1} + (u-t)q^{-2} + (q^2 + q + 1 - u)q^{-3}) \right]. \quad (5.34)$$

The right-hand side minus ϑ_q^3 is $(q-1)^2(qt - 6q + u)/q^6$. Since $t \leq 2$ and $u \leq 3q + 1$, we have $qt - 6q + u \leq 1 - q < 0$. This case is therefore strictly below ϑ_q^3 , while three distinct evaluation directions attain ϑ_q^3 . This proves (5.29). $\square$

Appendix D gives the equality characterization and explicit attaining projections.

6 Challenge Design for Direct High-Arity Quadratic Folding

We now apply the supported-direction framework of Section 3 to the challenge family itself.

Throughout this section, $\mathbb{F}$ is a field, and S and Y are finite-dimensional $\mathbb{F}$ -vector spaces. Let $Q : S \rightarrow Y$ be homogeneous quadratic, meaning that $Q(\alpha z) = \alpha^2 Q(z)$ for every $\alpha \in \mathbb{F}$, and assume that its polarization $B(x, y) := Q(x + y) - Q(x) - Q(y)$ is bilinear.

Fix $k \geq 1$. We take $X = S^{k+1}$, with one accumulator $z_0 \in S$ and k fresh sources $z_1, \dots, z_k \in S$, all varying freely. At the direct boundary, arbitrary functions of one source at a time are free. Thus

$$\text{Sep}_{k+1}(S) := \left\{ (z_0, \dots, z_k) \mapsto \sum_{i=0}^k g_i(z_i) : g_i : S \rightarrow \mathbb{F} \right\}. \quad (6.1)$$

We instantiate Section 3 with $\mathcal{F} = \text{Sep}_{k+1}(S)$ and $K = 0$. Thus individual-source information is already credited as free, and no additional joint mixed information is supplied by the incoming state. We write $m_{\min}^{\text{dir}}$ for the resulting retained width.

6.1 Fixed-Family Width

Let $\mathcal{C}$ be a finite nonempty challenge support, and let $a_i : \mathcal{C} \rightarrow \mathbb{F}$, $1 \leq i \leq k$, be the challenge weights. For $c \in \mathcal{C}$, set $z(c) := z_0 + \sum_{i=1}^k a_i(c)z_i$ and $H_c(z_0, \dots, z_k) := Q(z(c))$. Expanding the quadratic map gives

$$\begin{aligned}
Q(z(c)) &= Q(z_0) + \sum_{i=1}^k a_i(c)^2 Q(z_i) \\
&\quad + \sum_{i=1}^k a_i(c) B(z_0, z_i) + \sum_{1 \leq i < j \leq k} a_i(c) a_j(c) B(z_i, z_j).
\end{aligned} \tag{6.2}$$

The first two groups depend on one source at a time and are therefore free. Only the last two sums contribute to the residual target.

Let $Y_B := \text{span}_{\mathbb{F}}\{B(x, y) : x, y \in S\}$ and $r := \dim_{\mathbb{F}} Y_B$. Every mixed term in (6.2) lies in Y_B , so the mixed family may be regarded as Y_B -valued. This does not change its scalar residual span: every functional on Y_B extends to a functional on Y .

The challenge dependence of the mixed terms is described by

$$\begin{aligned}
V_a &:= \text{span}_{\mathbb{F}}(\{a_i : 1 \leq i \leq k\} \cup \{a_i a_j : 1 \leq i < j \leq k\}) \leq \mathbb{F}^{\mathcal{C}}, \\
d_a &:= \dim_{\mathbb{F}} V_a.
\end{aligned} \tag{6.3}$$

Theorem 6.1 (Direct high-arity width).

Under the direct-boundary assumptions above, $m_{\min}^{\text{dir}} = r d_a$.

Proof. Choose a basis $v_1, \dots, v_{d_a}$ of V_a . Expanding the functions a_i and $a_i a_j$ in this basis and regrouping the mixed terms in (6.2) writes their sum as $\sum_{\beta=1}^{d_a} v_{\beta}(c) M_{\beta}(z_0, \dots, z_k)$, where each M_{β} takes values in Y_B and is independent of c . Retaining the r coordinates of each M_{β} in a fixed basis of Y_B gives an explicit width- $r d_a$ realization.

It remains to show that no smaller state can realize the same family. Let $\mathcal{P} := \{(0, i) : 1 \leq i \leq k\} \cup \{(i, j) : 1 \leq i < j \leq k\}$ index the mixed source pairs. After reducing the output space to Y_B , the formal coefficient space of Proposition 3.2 is $\mathcal{D} = (Y_B^{\vee})^{\mathcal{P}}$. Because $K = 0$, its residual map is

$$\begin{aligned}
\Phi((\ell_{pq})_{(p,q) \in \mathcal{P}}) &:= \left[\sum_{(p,q) \in \mathcal{P}} \ell_{pq}(B(z_p, z_q)) \right] \\
&\in \mathbb{F}^{S^{k+1}} / \text{Sep}_{k+1}(S).
\end{aligned} \tag{6.4}$$

The map Φ is injective. Suppose the sum in (6.4) is source-local, and fix one pair (p, q) . Hold all other sources fixed and take a four-point rectangle difference in z_p and z_q . Every source-local term vanishes. Every mixed term other than the (p, q) term also vanishes because it depends on at most one of the two varied sources. The remaining expression is $\ell_{pq}(B(z_p^+ - z_p^-, z_q^+ - z_q^-))$. Since the source domain is the full product S^{k+1} , both differences range over all of S . As $B(S, S)$ spans Y_B , this forces $\ell_{pq} = 0$. Repeating the argument for every pair proves injectivity.

For each $c \in \mathcal{C}$, let $p(c) \in \mathbb{F}^{\mathcal{P}}$ have coordinates $p(c)_{0i} = a_i(c)$ and $p(c)_{ij} = a_i(c) a_j(c)$, and let $R_a := \text{span}_{\mathbb{F}}\{p(c) : c \in \mathcal{C}\}$. The vectors $p(c)$ form the rows

of the challenge-evaluation matrix whose columns are the evaluations of the functions a_i and $a_i a_j$. Its column space is V_a , so $\dim R_a = \dim V_a = d_a$.

Under the natural identification $\mathcal{D} \cong \mathbb{F}^P \otimes Y_B^\vee$, the supported coefficient space of Proposition 3.2 is $D_C = R_a \otimes Y_B^\vee$. Since Φ is injective, its null space is zero. Proposition 3.2 therefore gives $m_{\min}^{\text{dir}} = \dim D_C = d_a r$, matching the explicit width- rd_a realization above. $\square$

For a fixed quadratic map, r is fixed. Challenge design at this boundary therefore reduces to making d_a as small as possible. The design objective is therefore to compress the mixed challenge span while keeping the constant and diagonal coefficient functions distinguishable from it.

6.2 Lower Bounds from Diagonal Separation

We impose one additional condition on the challenge family. We say that it has diagonal separation if $b_0 + \sum_{i=1}^k b_i a_i^2 \in V_a$ implies $b_0 = b_1 = \dots = b_k = 0$. Equivalently, the classes of $1, a_1^2, \dots, a_k^2$ are linearly independent modulo V_a . Diagonal separation is an algebraic condition on the challenge family used here; by itself it is not a general protocol soundness statement.

Proposition 6.2 (Challenge-space lower bounds). *If the challenge family has diagonal separation, then $d_a \geq k$ and $|\mathcal{C}| \geq d_a + k + 1 \geq 2k + 1$. Consequently, $m_{\min}^{\text{dir}} \geq rk$.*

Proof. Suppose $\sum_{i=1}^k \alpha_i a_i = 0$. Squaring pointwise gives

$$\sum_i \alpha_i^2 a_i^2 = -2 \sum_{i < j} \alpha_i \alpha_j a_i a_j \in V_a;$$

in characteristic two, the right-hand side is zero. Diagonal separation therefore forces $\alpha_i^2 = 0$, and hence $\alpha_i = 0$, for every i . Thus $a_1, \dots, a_k$ are linearly independent, so $d_a \geq k$.

For the support bound, diagonal separation says that the $k + 1$ classes of $1, a_1^2, \dots, a_k^2$ are linearly independent in $\mathbb{F}^C/V_a$. Hence $|\mathcal{C}| - d_a \geq k + 1$, which gives $|\mathcal{C}| \geq d_a + k + 1 \geq 2k + 1$. The width bound follows from Theorem 6.1. $\square$

6.3 Scaled-Cauchy Attainment

Choose distinct poles $\xi_1, \dots, \xi_k \in \mathbb{F}$ and nonzero scales $\lambda_1, \dots, \lambda_k \in \mathbb{F}^\times$, and define $a_i(c) := \lambda_i / (c - \xi_i)$.

Theorem 6.3 (Scaled-Cauchy attainment). *Let $\mathcal{C}_{\text{Cau}} \subseteq \mathbb{F} \setminus \{\xi_1, \dots, \xi_k\}$ have at least $2k + 1$ elements. Then the scaled-Cauchy family has diagonal separation, $V_a = \text{span}_{\mathbb{F}}\{a_1, \dots, a_k\}$, $d_a = k$, and $m_{\min}^{\text{dir}} = rk$. If $|\mathcal{C}_{\text{Cau}}| = 2k + 1$, the support-cardinality lower bound of Proposition 6.2 is attained as well.*

Proof. For $i \neq j$, partial fractions give $a_i a_j = \frac{\lambda_j}{\xi_i - \xi_j} a_i - \frac{\lambda_i}{\xi_i - \xi_j} a_j$. Every pairwise product therefore lies in the span of the original Cauchy functions. Hence $V_a = \text{span}\{a_1, \dots, a_k\}$ and $d_a \leq k$. To obtain equality and invoke Proposition 6.2, it remains to verify diagonal separation on the finite support.

Let $D(T) := \prod_{i=1}^k (T - \xi_i)$ and $P_i(T) := D(T)/(T - \xi_i)$. If diagonal separation failed, the span description above would give coefficients $b_0, b_1, \dots, b_k, u_1, \dots, u_k \in \mathbb{F}$, with at least one of $b_0, \dots, b_k$ nonzero, such that $b_0 + \sum_i b_i a_i(c)^2 = \sum_i u_i a_i(c)$ for every $c \in \mathcal{C}_{\text{Cau}}$. Multiplying by $D(c)^2$ shows that every supported challenge is a root of

$$G(T) := b_0 D(T)^2 + \sum_{i=1}^k b_i \lambda_i^2 P_i(T)^2 - \sum_{i=1}^k u_i \lambda_i D(T) P_i(T). \quad (6.5)$$

The degree of G is at most $2k$, while $\mathcal{C}_{\text{Cau}}$ contains at least $2k + 1$ distinct points. Hence $G \equiv 0$. Fix i and evaluate at $T = \xi_i$. We have $D(\xi_i) = 0$, $P_j(\xi_i) = 0$ for $j \neq i$, and $P_i(\xi_i) = \prod_{j \neq i} (\xi_i - \xi_j) \neq 0$. Thus the only surviving term gives $b_i = 0$.

This holds for every i . Once the b_i 's vanish, $b_0 D(T)^2$ is the only term of degree $2k$ in G ; all remaining terms have degree at most $2k - 1$. Therefore $b_0 = 0$, contradicting the assumed failure of diagonal separation.

Proposition 6.2 now gives $d_a \geq k$, while the partial-fraction identity gave $d_a \leq k$. Thus $d_a = k$, and Theorem 6.1 gives $m_{\min}^{\text{dir}} = rk$. If $|\mathcal{C}_{\text{Cau}}| = 2k + 1$, the support bound is attained as well. $\square$

Over a finite field, choosing k distinct poles together with $2k + 1$ distinct supported non-pole values requires $|\mathbb{F}| \geq 3k + 1$.

6.4 The Cauchy Carrier

The scaled-Cauchy family also admits an explicit retained state of minimum width. For $1 \leq i < j \leq k$, let $P_{ij}(T) := D(T)/((T - \xi_i)(T - \xi_j))$, and define

$$H_{\text{src}}(T) := \sum_{i=1}^k \lambda_i P_i(T) B(z_0, z_i) + \sum_{1 \leq i < j \leq k} \lambda_i \lambda_j P_{ij}(T) B(z_i, z_j). \quad (6.6)$$

Since $\deg P_i = k - 1$ and $\deg P_{ij} = k - 2$, we have $\deg H_{\text{src}} < k$. For every $c \in \mathcal{C}_{\text{Cau}}$, the mixed contribution in (6.2) is $D(c)^{-1} H_{\text{src}}(c)$.

Fix a basis of Y_B . The k Y_B -valued coefficients of H_{src} occupy rk field coordinates. These coefficients are determined before the challenge, while evaluation at c and multiplication by the public nonzero scalar $D(c)^{-1}$ are linear in them. They therefore form a width- rk retained state, and Theorem 6.3 shows that this width is minimum.

The carrier also gives a cleared polynomial identity for the quadratic relation. To extend Q to S -valued polynomials without restricting the characteristic, choose a basis $s_1, \dots, s_n$ of S . For $w(T) = \sum_{j=1}^n f_j(T) s_j$, define

$$Q_T(w(T)) := \sum_{j=1}^n f_j(T)^2 Q(s_j) + \sum_{1 \leq i < j \leq n} f_i(T) f_j(T) B(s_i, s_j). \quad (6.7)$$

This extension agrees with Q on S , and its polarization is the $\mathbb{F}[T]$ -bilinear extension of B .

Proposition 6.4 (Cauchy carrier identity).

Set $W(T) := D(T)z_0 + \sum_{i=1}^k \lambda_i P_i(T)z_i$. Then

$$Q_T(W(T)) = D(T)^2 Q(z_0) + \sum_{i=1}^k \lambda_i^2 P_i(T)^2 Q(z_i) + D(T) H_{\text{src}}(T). \quad (6.8)$$

Proof. Expanding $Q_T(W(T))$ using (6.7), the diagonal terms give the first two terms in (6.8). The accumulator-source terms are $D(T)\lambda_i P_i(T)B(z_0, z_i)$. For $i < j$, $P_i(T)P_j(T) = D(T)P_{ij}(T)$, so the source-source terms are

$$D(T)\lambda_i \lambda_j P_{ij}(T)B(z_i, z_j).$$

Together the mixed terms are exactly $D(T)H_{\text{src}}(T)$. □

The separate CauchyFold construction [19] commits the coefficients of H_{src} before the scaled-Cauchy challenge and evaluates the resulting carrier afterward inside a concrete lattice folding protocol.

7 Conclusion

At a fixed public-randomness boundary, the relevant resource is not the size of a particular protocol state, but the scalar information that remains after free and incoming directions have been removed. Under arbitrary prechallenge preprocessing and linear postchallenge readout, the exact retained width is the dimension of this residual target. This viewpoint separates three effects: the admissible source domain determines which mixed directions exist, the actual challenge support determines which coefficient directions are generated, and the boundary accounting determines which generated directions remain genuinely new. Below the exact-width threshold, dimension alone no longer determines optimal reconstruction: the rank geometry of the missing scalar forms matters, and bilinear side information couples their representatives through a common linear section. Under diagonal separation in the direct full-product quadratic model, scaled-Cauchy challenges attain the resulting width lower bound and admit a matching minimum-width carrier.

Although motivated by folding, the boundary question is more general. It arises whenever input-dependent information must be fixed before a later public challenge or query and is subsequently used linearly together with side information. Understanding which residual invariants continue to govern such staged preprocessing and recovery problems under richer later-use interfaces is a natural direction beyond the setting studied here.

References

1. Bläser, M.: Fast Matrix Multiplication. No. 5 in Graduate Surveys, Theory of Computing Library (2013). <https://doi.org/10.4086/toc.gs.2013.005>, <https://theoryofcomputing.org/articles/g005/>
2. Bünz, B., Chen, B.: ProtoStar: Generic Efficient Accumulation/Folding for Special-Sound Protocols. In: Advances in Cryptology—ASIACRYPT 2023, Part II. Lecture Notes in Computer Science, vol. 14439, pp. 77–110. Springer (2023). https://doi.org/10.1007/978-981-99-8724-5_3
3. Bünz, B., Chiesa, A., Fenzi, G., Wang, W.: Linear-Time Accumulation Schemes. In: Theory of Cryptography—TCC 2025, Part I. Lecture Notes in Computer Science, vol. 16268, pp. 369–399. Springer (2026). https://doi.org/10.1007/978-3-032-12287-2_13, first published online 2 December 2025
4. Bünz, B., Mishra, P., Nguyen, W., Wang, W.: Accumulation Without Homomorphism. In: 16th Innovations in Theoretical Computer Science Conference (ITCS 2025). Leibniz International Proceedings in Informatics (LIPIcs), vol. 325, pp. 23:1–23:25. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2025). <https://doi.org/10.4230/LIPIcs.ITCS.2025.23>
5. Cohn, P.M.: Free Ideal Rings and Localization in General Rings, New Mathematical Monographs, vol. 3. Cambridge University Press (2006). <https://doi.org/10.1017/CB09780511542794>
6. Delsarte, P.: Bilinear Forms over a Finite Field, with Applications to Coding Theory. Journal of Combinatorial Theory, Series A **25**(3), 226–241 (1978). [https://doi.org/10.1016/0097-3165\(78\)90015-8](https://doi.org/10.1016/0097-3165(78)90015-8)
7. Dimitriou, N., Garreta, A., Manzur, I., Vlasov, I.: Mova: Nova Folding without Committing to Error Terms. Cryptology ePrint Archive, Paper 2024/1220 (2024), <https://eprint.iacr.org/2024/1220>
8. Dvir, Z., Golovnev, A., Weinstein, O.: Static Data Structure Lower Bounds Imply Rigidity. In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing. pp. 967–978. STOC '19, ACM (2019). <https://doi.org/10.1145/3313276.3316348>
9. Eagen, L., Gabizon, A.: ProtoGalaxy: Efficient ProtoStar-style Folding of Multiple Instances. Cryptology ePrint Archive, Paper 2023/1106 (2023), <https://eprint.iacr.org/2023/1106>
10. Kothapalli, A., Setty, S.: HyperNova: Recursive Arguments for Customizable Constraint Systems. In: Advances in Cryptology—CRYPTO 2024, Part X. Lecture Notes in Computer Science, vol. 14929, pp. 345–379. Springer (2024). https://doi.org/10.1007/978-3-031-68403-6_11
11. Kothapalli, A., Setty, S.: NeutronNova: Group-Based Folding Done Right. In: Advances in Cryptology—CRYPTO 2026, Part IX. Lecture Notes in Computer Science, vol. 16808, pp. 524–556. Springer (2026). https://doi.org/10.1007/978-3-032-35424-2_18
12. Kothapalli, A., Setty, S., Tzialla, I.: Nova: Recursive Zero-Knowledge Arguments from Folding Schemes. In: Advances in Cryptology—CRYPTO 2022, Part IV. Lecture Notes in Computer Science, vol. 13510, pp. 359–388. Springer (2022). https://doi.org/10.1007/978-3-031-15985-5_13
13. Lund, C., Fortnow, L., Karloff, H., Nisan, N.: Algebraic Methods for Interactive Proof Systems. Journal of the ACM **39**(4), 859–868 (1992). <https://doi.org/10.1145/146585.146605>

14. Meshulam, R.: Spaces of Hankel Matrices over Finite Fields. *Linear Algebra and its Applications* **218**, 73–76 (1995). [https://doi.org/10.1016/0024-3795\(93\)00158-V](https://doi.org/10.1016/0024-3795(93)00158-V)
15. Mirandola, D., Zémor, G.: Critical Pairs for the Product Singleton Bound. *IEEE Transactions on Information Theory* **61**(9), 4928–4937 (2015). <https://doi.org/10.1109/TIT.2015.2450207>
16. Osadnik, M.: PikkuFold: Efficient Folding in a Few Kilobytes. *Cryptology ePrint Archive*, Paper 2026/1809 (2026), <https://eprint.iacr.org/2026/1809>
17. Randriambololona, H.: On Products and Powers of Linear Codes under Component-wise Multiplication. In: *Algorithmic Arithmetic, Geometry, and Coding Theory*, Contemporary Mathematics, vol. 637, pp. 3–78. American Mathematical Society (2015). <https://doi.org/10.1090/conm/637/12749>
18. Ravagnani, A.: Rank-Metric Codes and Their Duality Theory. *Designs, Codes and Cryptography* **80**(1), 197–216 (2016). <https://doi.org/10.1007/s10623-015-0077-3>
19. Wang, X.: CauchyFold: Residue-Optimal High-Arity Lattice Folding via Scaled Cauchy Challenges. *Cryptology ePrint Archive*, Paper 2026/2011 (2026), <https://eprint.iacr.org/2026/2011>

A Retained Width over Commutative Rings

Section 3 characterizes retained width over a field. Over a commutative ring, the corresponding quantity depends on the residual target together with its embedding in the ambient quotient. We describe it by an ambient cover width and derive matrix formulas for free and finitely presented ambient modules.

A.1 The Ring-Valued Boundary Model

Let A be a commutative ring with identity, let X be the admissible semantic domain, and let $\mathcal{C}$ be a finite nonempty challenge support. We declare an A -submodule $\mathcal{F} \leq A^X$ of scalar functions to be free at the boundary.

The incoming state is a map $K : X \rightarrow P$, where P is an A -module. For each $c \in \mathcal{C}$, the required output is a map $H_c : X \rightarrow Y$, where Y is a finitely generated projective A -module. For an A -module M , write $M^\vee := \text{Hom}_A(M, A)$.

A width- m realization consists of one retained-state map $\Psi : X \rightarrow A^m$, maps $G_c : X \rightarrow Y$, and A -linear maps $L_{K,c} : P \rightarrow Y$ and $L_{\Psi,c} : A^m \rightarrow Y$ satisfying

$$H_c(x) = G_c(x) + L_{K,c}(K(x)) + L_{\Psi,c}(\Psi(x)) \quad (\text{A.1})$$

for every $x \in X$ and $c \in \mathcal{C}$. We require $\eta \circ G_c \in \mathcal{F}$ for every $\eta \in Y^\vee$. The same prechallenge state Ψ must work for all supported challenges. The readout maps may depend on c and fixed public parameters, but not on the realized values of x , $K(x)$, or $\Psi(x)$. The prechallenge map Ψ is otherwise unrestricted.

Let $\pi_{\mathcal{F}} : A^X \rightarrow \mathcal{Q}_A := A^X / \mathcal{F}$ be the quotient map. Define

$$U_A(K) := \{\pi_{\mathcal{F}}(\omega \circ K) : \omega \in P^\vee\}, \quad M_K := \mathcal{Q}_A / U_A(K). \quad (\text{A.2})$$

Let $\rho_K : A^X \rightarrow M_K$ be the composite quotient map. Its kernel is $\mathcal{F} + \{\omega \circ K : \omega \in P^\vee\}$. The residual target is the embedded submodule

$$T_{H|K} := \text{span}_A \{\rho_K(\eta \circ H_c) : c \in \mathcal{C}, \eta \in Y^\vee\} \leq M_K. \quad (\text{A.3})$$

It is finitely generated because $\mathcal{C}$ is finite and $Y^\vee$ is finitely generated.

A.2 Exact Width via Ambient Cover

For a finitely generated submodule $T \leq M$, define its ambient cover width by

$$\text{cov}_A(T \subseteq M) := \min \left\{ m \geq 0 : \begin{array}{l} \text{there is an } A\text{-linear map } \varphi : A^m \rightarrow M \\ \text{such that } T \subseteq \text{im } \varphi \end{array} \right\}. \quad (\text{A.4})$$

The covering image may be strictly larger than T . Thus the resource concerns the embedded target $T \subseteq M$, rather than only the minimum number of generators of the abstract module T .

For a candidate state $\Psi = (\psi_1, \dots, \psi_m) : X \rightarrow A^m$, define the A -linear map $\widehat{\Psi} : A^m \rightarrow M_K$ by $\widehat{\Psi}(e_i) = \rho_K(\psi_i)$.

Theorem A.1 (Feasibility and exact retained width over a ring). *In the model of Appendix A.1, a candidate state $\Psi : X \rightarrow A^m$ can serve as the retained state of a width- m realization if and only if*

$$T_{H|K} \subseteq \text{im } \widehat{\Psi}.$$

Consequently,

$$m_{\min}(H \mid K; \mathcal{F}) = \text{cov}_A(T_{H|K} \subseteq M_K). \quad (\text{A.5})$$

Proof. Suppose $\Psi = (\psi_1, \dots, \psi_m)$ is feasible. Scalarizing (A.1) and applying ρ_K removes the free and incoming terms. Every class $\rho_K(\eta \circ H_c)$ therefore lies in $\text{span}_A \{\rho_K(\psi_1), \dots, \rho_K(\psi_m)\} = \text{im } \widehat{\Psi}$. These classes generate $T_{H|K}$, proving the required inclusion.

Conversely, suppose $T_{H|K} \subseteq \text{im } \widehat{\Psi}$. Finite projectivity of Y gives a finite dual basis $(y_j, \eta_j)_{j=1}^N$ satisfying

$$y = \sum_{j=1}^N y_j \eta_j(y) \quad \text{for every } y \in Y.$$

For each c, j , choose coefficients $a_{c,j,i} \in A$ representing $\rho_K(\eta_j \circ H_c)$ in the classes $\rho_K(\psi_i)$ of the given state. The remaining difference lies in $\ker \rho_K$, so there are $f_{c,j} \in \mathcal{F}$ and $\omega_{c,j} \in P^\vee$ such that

$$\eta_j \circ H_c = f_{c,j} + \omega_{c,j} \circ K + \sum_{i=1}^m a_{c,j,i} \psi_i. \quad (\text{A.6})$$

Assemble these scalar decompositions by defining

$$\begin{aligned} G_c(x) &:= \sum_{j=1}^N y_j f_{c,j}(x), \\ L_{K,c}(p) &:= \sum_{j=1}^N y_j \omega_{c,j}(p), \\ L_{\Psi,c}(z) &:= \sum_{j=1}^N y_j \left(\sum_{i=1}^m a_{c,j,i} z_i \right). \end{aligned}$$

Both readout maps are A -linear and fixed independently of the realized input. For every $\eta \in Y^\vee$, we have $\eta \circ G_c = \sum_j \eta(y_j) f_{c,j} \in \mathcal{F}$, so the free condition holds. Multiplying (A.6) by y_j , summing over j , and using the finite dual-basis identity yields (A.1), proving feasibility of the given Ψ .

Every feasible state induces an ambient cover of the same width. Conversely, any ambient cover $\varphi : A^m \rightarrow M_K$ lifts through the surjection ρ_K : choose $\psi_i : X \rightarrow A$ with $\rho_K(\psi_i) = \varphi(e_i)$. Then $\hat{\Psi} = \varphi$, so the criterion just proved makes Ψ feasible. Taking minima gives (A.5). $\square$

The construction requires no projectivity assumption on P . Over a field, ambient cover width equals the dimension of the target subspace, so Theorem A.1 specializes to Theorem 3.1.

A.3 Computing the Cover Width

For $R \in A^{n \times \ell}$, define its inner rank by

$$\text{irank}_A(R) := \min \{ m \geq 0 : R = EF, \quad E \in A^{n \times m}, \quad F \in A^{m \times \ell} \}. \quad (\text{A.7})$$

Equivalently, this is the least free rank through which the represented map factors [5, §0.1 and Proposition 5.4.3].

Proposition A.2 (Free ambient module). *For $T = \text{im } R \leq A^n$,*

$$\text{cov}_A(T \subseteq A^n) = \text{irank}_A(R). \quad (\text{A.8})$$

Proof. A factorization $R = EF$ gives an ambient cover $E : A^m \rightarrow A^n$. Conversely, let $\varphi : A^m \rightarrow A^n$ cover T , and let E be its matrix. Each column of R has a preimage under E . Collecting these preimages as the columns of F gives $R = EF$. Taking minima proves (A.8). $\square$

Now let

$$A^p \xrightarrow{L} A^n \xrightarrow{\pi} M \rightarrow 0$$

be a finite presentation, with $L \in A^{n \times p}$. Let $R \in A^{n \times \ell}$ represent target generators, and put $T := \text{im}(\pi R)$.

Proposition A.3 (Presented ambient module). *With this notation,*

$$\text{cov}_A(T \subseteq M) = \min_{Z \in A^{p \times \ell}} \text{irank}_A(R - LZ). \quad (\text{A.9})$$

Equivalently, this is the least m for which $R = EF + LZ$, with $E \in A^{n \times m}$, $F \in A^{m \times \ell}$, and $Z \in A^{p \times \ell}$.

Proof. If $R - LZ = EF$, then $\pi R = (\pi E)F$, so $\pi E : A^m \rightarrow M$ covers T .

Conversely, let $\varphi : A^m \rightarrow M$ cover T . Lift each $\varphi(e_i)$ to A^n , and collect these lifts in a matrix E , so $\pi E = \varphi$. Choose F with $\pi R = \varphi F$. Every column of $R - EF$ belongs to $\ker \pi = \text{im } L$, so $R - EF = LZ$ for some Z . This proves the reverse inequality. $\square$

The minimization over Z accounts for changes of target-generator representatives by presentation relations. Over a field, the formula reduces to

$$\text{cov}_{\mathbb{F}}(T \subseteq M) = \text{rank}[L \ R] - \text{rank } L, \quad (\text{A.10})$$

because $T \cong (\text{im } L + \text{im } R) / \text{im } L$.

A.4 A Ring Example

Let $A = \mathbb{Z}/4\mathbb{Z}$, and take

$$L = \begin{pmatrix} 2 \\ 0 \end{pmatrix}, \quad R = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}.$$

Let $M = A^2 / \text{im } L$, with quotient map π . The first column of R vanishes in M , while the second remains nonzero. Hence $T = \text{im}(\pi R)$ has ambient cover width one.

In contrast, $\text{irank}_A(R) = 2$. Indeed, a width-one factorization $R = pq$, with $p \in A^{2 \times 1}$ and $q \in A^{1 \times 2}$, would give $p_2 q_2 = 1$. Thus p_2 is a unit. The lower-left entry then forces $q_1 = 0$, contradicting the upper-left entry $p_1 q_1 = 2$.

Taking $Z = (1, 0)$ gives

$$R - LZ = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix},$$

whose inner rank is one. The presentation relations must therefore be incorporated before measuring the factorization width.

B Finite Challenge Supports and Evaluation Families

Section 3.3 distinguishes formal coefficient directions from those generated by the actual challenge support. We apply that distinction to polynomial evaluation and give two supports of equal cardinality with different retained widths.

B.1 Evaluation-Family Width

Work over a field $\mathbb{F}$, and let $g_x(T) = \sum_{j=0}^d A_j(x)T^j$ be a prechallenge polynomial, where each $A_j : X \rightarrow \mathbb{F}$ is a scalar function. Using the declared-free space $\mathcal{F}$ and quotient $\mathcal{Q} = \mathbb{F}^X / \mathcal{F}$, define

$$\Gamma : \mathbb{F}^{d+1} \rightarrow \mathcal{Q}, \quad \Gamma(a_0, \dots, a_d) := \pi_{\mathcal{F}} \left(\sum_{j=0}^d a_j A_j \right). \quad (\text{B.1})$$

For $r \in \mathbb{F}$, write $v(r) := (1, r, \dots, r^d)$. The scalar class of $g_x(r)$ is $\Gamma(v(r))$.

Suppose the incoming state contains a root state K_0 and coefficient measurements corresponding to a subspace $E_{\text{in}} \leq \mathbb{F}^{d+1}$. For a basis $b_1, \dots, b_t$ of E_{in} , these measurements are $\kappa_i(x) := \sum_{j=0}^d (b_i)_j A_j(x)$. Thus, for $K = (K_0, \kappa_1, \dots, \kappa_t)$, the incoming scalar span is $U(K) = U(K_0) + \Gamma(E_{\text{in}})$. For a finite nonempty evaluation support $\Omega \subseteq \mathbb{F}$, let $E_{\Omega} := \text{span}_{\mathbb{F}}\{v(r) : r \in \Omega\}$.

Proposition B.1 (Evaluation-family width). *The minimum retained width needed to produce $g_x(r)$ for every $r \in \Omega$ is*

$$m_{\min} = \dim_{\mathbb{F}} \frac{U(K_0) + \Gamma(E_{\text{in}} + E_{\Omega})}{U(K_0) + \Gamma(E_{\text{in}})}. \quad (\text{B.2})$$

Proof. The incoming coefficient measurements contribute $\Gamma(E_{\text{in}})$, and the supported evaluations contribute $\Gamma(E_{\Omega})$. Theorem 3.1 gives (B.2). $\square$

The kernel of Γ captures coefficient relations after quotienting by the declared-free space. The denominator in (B.2) then credits both the root state and the incoming coefficient measurements.

On the full polynomial domain $\mathbb{F}[T]_{\leq d}$, with no free information or root state, Γ is injective. The formula becomes

$$m_{\min} = \dim_{\mathbb{F}} \frac{E_{\text{in}} + E_{\Omega}}{E_{\text{in}}}. \quad (\text{B.3})$$

Although the Vandermonde argument gives $\dim E_{\Omega} = \min\{|\Omega|, d+1\}$, its intersection with E_{in} can depend on the supported points.

B.2 Equal Cardinality, Different Width

Work over $\mathbb{F}_7$, let g range over all polynomials of degree at most two, and let the incoming measurement be $\ell(g) := g(0) + g(1)$. Its coefficient direction is $s_0 = v(0) + v(1) = (2, 1, 1)$.

For a two-point support Ω , the evaluation vectors are independent. Equation (B.3) gives

$$m_{\min}(\Omega) = 2 - \mathbf{1}[s_0 \in E_{\Omega}]. \quad (\text{B.4})$$

Let $Z_{\Omega}(T) := \prod_{r \in \Omega} (T - r)$. The kernel of evaluation on Ω , restricted to degree-at-most-two polynomials, is spanned by Z_{Ω} . Row-space/kernel duality therefore gives $s_0 \in E_{\Omega}$ if and only if $\ell(Z_{\Omega}) = 0$.

For $\Omega_1 = \{2, 3\}$, the values $Z_{\Omega_1}(0) = 6$ and $Z_{\Omega_1}(1) = 2$ give $\ell(Z_{\Omega_1}) = 1$ in $\mathbb{F}_7$, so $m_{\min}(\Omega_1) = 2$. For $\Omega_2 = \{2, 5\}$, the corresponding values are 3 and 4, giving $\ell(Z_{\Omega_2}) = 0$ and $m_{\min}(\Omega_2) = 1$. A matching realization follows from

$$\ell(g) = 3g(2) + 6g(5), \quad g(5) = 3g(2) + 6\ell(g). \quad (\text{B.5})$$

Thus retaining $g(2)$ suffices to recover both supported evaluations.

The two supports have equal cardinality, but their evaluation spans have different intersections with the incoming direction.

C Rank Distributions in the Side-Information Example

Example 5.4 compares two bilinear side-information tasks with exact remaining width two. We parameterize every feasible plane in their dual optimizations, determine its rank distribution, and analyze the simultaneous choice of rank-one representatives.

Throughout this appendix, let $\mathbb{F} = \mathbb{F}_q$ and $U = V = \mathbb{F}_q^2$, with $\mu(u, v) = (u_1v_1, u_2v_2)$, $\nu_S(u, v) = u_1v_2 + u_2v_1$, and $\nu_T(u, v) = (u_1 + u_2)v_2$. We identify scalar bilinear forms with matrices through $B(u, v) = u^\top Bv$.

C.1 Parameterizing All Feasible Planes

The scalar-form spaces and side-information subspaces are

$$\begin{aligned} \mathcal{B}_S &= \left\{ \begin{pmatrix} \alpha & \beta \\ \beta & \gamma \end{pmatrix} : \alpha, \beta, \gamma \in \mathbb{F}_q \right\}, & S_S &= \mathbb{F}_q \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \\ \mathcal{B}_T &= \left\{ \begin{pmatrix} \alpha & \beta \\ 0 & \gamma + \beta \end{pmatrix} : \alpha, \beta, \gamma \in \mathbb{F}_q \right\}, & S_T &= \mathbb{F}_q \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}. \end{aligned} \quad (\text{C.1})$$

In both cases, the quotient by the side-information line is identified with $\mathbb{F}_q^2$ through the coordinates (α, γ) .

At retained width zero, Theorem 5.3 optimizes over two-dimensional subspaces N disjoint from the corresponding side-information line. The quotient map restricts to an isomorphism on each such N . Every feasible plane is therefore the graph of a unique linear function

$$\beta = a\alpha + b\gamma, \quad a, b \in \mathbb{F}_q. \quad (\text{C.2})$$

For a plane N , let $t(N)$ denote the number of its projective rank-one directions. There is one matrix of rank zero, $(q-1)t(N)$ matrices of rank one, and $(q-1)(q+1-t(N))$ matrices of rank two. Consequently,

$$q^{-2} \sum_{B \in N} q^{-\text{rank } B} = q^{-2} [1 + (q-1)t(N)q^{-1} + (q-1)(q+1-t(N))q^{-2}]. \quad (\text{C.3})$$

This expression is strictly increasing in $t(N)$.

C.2 Rank Distributions

Symmetric side information. The plane associated with (a, b) is

$$N_{a,b}^S = \left\{ \begin{pmatrix} \alpha & a\alpha + b\gamma \\ a\alpha + b\gamma & \gamma \end{pmatrix} : \alpha, \gamma \in \mathbb{F}_q \right\}. \quad (\text{C.4})$$

A nonzero member has rank one exactly when

$$\alpha\gamma - (a\alpha + b\gamma)^2 = 0. \quad (\text{C.5})$$

The left-hand side is a nonzero homogeneous quadratic: if its two square coefficients vanish, then $a = b = 0$, leaving the term $\alpha\gamma$. It therefore has at most two projective zeros.

For odd q , put $\Delta := 1 - 4ab$. Then

$$t(N_{a,b}^S) = \begin{cases} 2, & \Delta \text{ is a nonzero square,} \\ 1, & \Delta = 0, \\ 0, & \Delta \text{ is a nonsquare.} \end{cases} \quad (\text{C.6})$$

If $a = 0$, equation (C.5) factors as $\gamma(\alpha - b^2\gamma) = 0$, giving two distinct projective roots. If $a \neq 0$, there is no root with $\gamma = 0$, and the quadratic in α/γ has discriminant $1 - 4ab$.

For $q = 2^h$, let $\text{Tr} : \mathbb{F}_q \rightarrow \mathbb{F}_2$ be the absolute trace. Then

$$t(N_{a,b}^S) = \begin{cases} 2, & \text{Tr}(ab) = 0, \\ 0, & \text{Tr}(ab) = 1. \end{cases} \quad (\text{C.7})$$

When $a = 0$, the same factorization gives two roots. When $a \neq 0$, setting $x = \alpha/\gamma$ and $y = a^2x$ transforms (C.5) into $y^2 + y = (ab)^2$. The map $y \mapsto y^2 + y$ has kernel $\{0, 1\}$ and image equal to the trace-zero subspace. Since $\text{Tr}((ab)^2) = \text{Tr}(ab)$, equation (C.7) follows.

Thus every feasible symmetric plane has at most two projective rank-one directions, and the diagonal plane attains two.

Triangular side information. Every feasible plane has the form

$$N_{a,b}^T = \left\{ \begin{pmatrix} \alpha & a\alpha + b\gamma \\ 0 & a\alpha + (b+1)\gamma \end{pmatrix} : \alpha, \gamma \in \mathbb{F}_q \right\}. \quad (\text{C.8})$$

Its determinant is $\alpha(a\alpha + (b+1)\gamma)$. Hence

$$t(N_{a,b}^T) = \begin{cases} q+1, & a = 0, b = -1, \\ 1, & a \neq 0, b = -1, \\ 2, & b \neq -1. \end{cases} \quad (\text{C.9})$$

In the first case, every nonzero matrix has rank one. In the second, the only projective zero is $\alpha = 0$. In the third, the two linear factors give two distinct projective directions. Together with (C.3), these formulas determine the complete rank distribution of every feasible plane.

C.3 Compatibility of Rank-One Representatives

The maximum counts 2 and $q+1$ in (C.3) give the width-zero values ϑ_q^2 and ξ_q in Table 1. Both examples have a rank-one form outside the side-information line, so the one-deficit formula (5.24) gives the common width-one value ϑ_q . Width two permits exact reconstruction.

In characteristic two, every nonzero quotient class (α, γ) has a rank-one representative in the symmetric example. Indeed, bijectivity of squaring on $\mathbb{F}_q$ gives a unique β satisfying $\beta^2 = \alpha\gamma$. The resulting matrix is nonzero with determinant zero, hence has rank one.

These choices cannot form a linear section whose nonzero values all have rank one. Such a section would have a two-dimensional image containing $q+1$ projective rank-one directions, whereas every feasible symmetric plane has at most two.

For the triangular example, the choice $\beta = -\gamma$ is linear and gives the section

$$(\alpha, \gamma) \mapsto \begin{pmatrix} \alpha & -\gamma \\ 0 & 0 \end{pmatrix}. \quad (\text{C.10})$$

Every nonzero value of this section has rank one. Thus both examples have pointwise minimum rank one on every nonzero quotient class in characteristic two, but only the triangular example admits a linear section attaining all these minima simultaneously.

D Hankel Rank Geometry Near the Threshold

This appendix proves the rank-one classification, evaluation-span rank law, and cubic point bound used in Section 5.5. It then derives the equality cases and gives explicit optimal projections.

D.1 Hankel Scalarizations and Rank-One Directions

Let $U = V = \mathbb{F}_q[t]_{<k}$, $Y = \mathbb{F}_q[t]_{<2k-1}$, and $\mu(f, g) = fg$. For $\lambda \in Y^\vee$, define

$$(B_\lambda)_{ij} := \lambda(t^{i+j}), \quad 0 \leq i, j < k. \quad (\text{D.1})$$

The map $\lambda \mapsto B_\lambda$ is an isomorphism onto the full $k \times k$ Hankel space $\mathcal{H}_k$: every exponent from 0 to $2k-2$ occurs as $i+j$, and the values of λ on these monomials may be prescribed independently. In particular, $\dim \mathcal{H}_k = 2k-1$.

For $k=1$, the only nontrivial deficit is scalar multiplication. Assume $k \geq 2$ below. For $\alpha \in \mathbb{F}_q$, let $\ell_\alpha(p) := p(\alpha)$, and let $\ell_\infty(p) := [t^{2k-2}]p$. Set $v_\alpha := (1, \alpha, \dots, \alpha^{k-1})^\top$ and $v_\infty := e_{k-1}$. Then

$$B_{\ell_\alpha} = v_\alpha v_\alpha^\top, \quad B_{\ell_\infty} = v_\infty v_\infty^\top. \quad (\text{D.2})$$

Lemma D.1 (Rank-one Hankel directions). *Every nonzero rank-one Hankel matrix is a nonzero scalar multiple of exactly one matrix in (D.2).*

Proof. A nonzero symmetric rank-one matrix has the form $\gamma vv^\top$, with $\gamma \neq 0$ and $v \neq 0$. Indeed, starting from a rank-one factorization $vw^\top$, symmetry forces w to be proportional to v .

If $v_0 \neq 0$, the Hankel identities give $v_0 v_j = v_1 v_{j-1}$ for $1 \leq j < k$. Setting $\alpha = v_1/v_0$, induction yields $v_j = v_0 \alpha^j$, giving the direction B_{ℓ_α} .

If $v_0 = 0$, let j be the least index with $v_j \neq 0$. Were $j < k-1$, equality of the entries at (j, j) and $(j-1, j+1)$ would give $v_j^2 = v_{j-1} v_{j+1} = 0$, a contradiction. Thus only v_{k-1} can be nonzero, giving the direction B_{ℓ_∞} .

The finite parameter is determined by v_1/v_0 , and the infinite direction has zero first coordinate. Hence the listed projective directions are distinct. $\square$

D.2 The Evaluation-Span Rank Law

Lemma D.2 (Evaluation-span rank law). *Let $\alpha_1, \dots, \alpha_s \in \mathbb{F}_q \cup \{\infty\}$ be distinct, with $s \leq k$. Then*

$$\text{rank} \left(\sum_{i=1}^s b_i B_{\ell_{\alpha_i}} \right) = |\{i : b_i \neq 0\}| \quad (\text{D.3})$$

for every $b \in \mathbb{F}_q^s$.

Proof. The vectors v_{α_i} are linearly independent. If all points are finite, their first s coordinates give a nonsingular Vandermonde matrix. If one point is ∞ , the first $s-1$ coordinates force all finite-point coefficients to vanish, after which the coefficient of v_∞ also vanishes.

Let V_α be the matrix with columns v_{α_i} , and choose a left inverse L_α . Then

$$\sum_{i=1}^s b_i B_{\ell_{\alpha_i}} = V_\alpha \text{diag}(b) V_\alpha^\top.$$

This matrix has rank at most the number of nonzero entries of b . Multiplication on the left by L_α and on the right by $L_\alpha^\top$ recovers $\text{diag}(b)$, proving the reverse inequality. $\square$

For $N = \text{span}\{\ell_{\alpha_1}, \dots, \ell_{\alpha_s}\}$, the rank law gives

$$q^{-s} \sum_{\lambda \in N} q^{-\text{rank } B_\lambda} = q^{-s} \sum_{b \in \mathbb{F}_q^s} q^{-|\text{supp } b|} = \left(\frac{2q-1}{q^2} \right)^s. \quad (\text{D.4})$$

D.3 The Cubic Point Count

Lemma D.3 (Plane cubic bound). *A nonzero homogeneous cubic over $\mathbb{F}_q$ has at most $3q+1$ zeros in $\mathbf{P}^2(\mathbb{F}_q)$.*

Proof. Suppose first that the cubic has an $\mathbb{F}_q$ -rational line component. That line contains $q + 1$ rational points. After a projective change of coordinates taking it to the line at infinity, the remaining quadratic factor restricts to a nonzero polynomial of total degree at most two on the affine plane.

Such a polynomial has at most $2q$ affine zeros. To see this, write it as a polynomial in the second variable of degree d , with leading coefficient a_d in the first variable. At most $2 - d$ first-coordinate values make a_d vanish. Every other fiber contains at most d zeros; bounding an exceptional fiber by q gives at most $2q$ zeros in total. The projective cubic therefore has at most $3q + 1$ rational zeros.

Now suppose there is no rational line component. If there is a rational zero P , every rational line through P has a nonzero degree-three restriction and contains at most two further zeros. There are $q + 1$ such lines, and every other rational zero lies on exactly one of them. The total is therefore at most $1 + 2(q + 1) \leq 3q + 1$, since $q \geq 2$. If there is no rational zero, the claim is immediate. $\square$

D.4 Equality Cases Near the Threshold

For the $s = 3$ case, we use the consequence of Meshulam's theorem stated in the proof of Theorem 5.5: the $k \times k$ Hankel space contains a $(2k - 3)$ -dimensional subspace whose nonzero matrices all have rank at least three. The equality cases now follow from the upper-bound argument of Theorem 5.5.

Corollary D.4 (Equality characterization).

Let $k \geq 2$, let $1 \leq s \leq \min\{3, k\}$, and let $N \leq Y^\vee$ have dimension s . Then

$$q^{-s} \sum_{\lambda \in N} q^{-\text{rank } B_\lambda} = \vartheta_q^s \quad (\text{D.5})$$

if and only if N is spanned by s distinct projective evaluation functionals ℓ_α , with $\alpha \in \mathbb{F}_q \cup \{\infty\}$.

Proof. For $s = 1$, equality requires the nonzero generator to have rank one, so Lemma D.1 applies.

For $s = 2$, the argument used in Example 5.4 allows at most two projective rank-one directions. The normalized rank sum reaches ϑ_q^2 only when both are present. They span N , and Lemma D.1 identifies them as evaluation directions.

For $s = 3$, the case with at most two projective rank-one directions in Theorem 5.5 gives a strict inequality. Equality therefore requires at least three. Any three are linearly independent by the two-dimensional argument and hence form a basis of N . Lemma D.1 identifies them as distinct evaluation directions.

Conversely, Lemma D.2 and (D.4) give equality for every such evaluation span. $\square$

D.5 Explicit Optimal Projections

Let $N = \text{span}\{\ell_{\alpha_1}, \dots, \ell_{\alpha_s}\}$ be an attaining space, and put $W = N^\perp$. Choose $h_1, \dots, h_s \in Y$ with $\ell_{\alpha_i}(h_j) = \delta_{ij}$; such polynomials exist because the evaluation

functionals are linearly independent. Then

$$\Pi_N p := p - \sum_{i=1}^s \ell_{\alpha_i}(p) h_i \quad (\text{D.6})$$

is a projection onto W : each ℓ_{α_i} vanishes on its image, and Π_N fixes every element of W . Retaining the coordinates of $\Pi_N T_\mu$ therefore realizes the optimum of Theorem 5.2.

For explicit choices, write $p(t) = \sum_{j=0}^{m_0} p_j t^j$, where $m_0 := 2k - 2$. For one deficit, take $N = \text{span}\{\ell_0\}$ and $\Pi_1 p = p - p_0$. The retained coordinates are the coefficients of T_μ in degrees $1, \dots, m_0$.

For two deficits, take $N = \text{span}\{\ell_0, \ell_\infty\}$ and $\Pi_2 p = p - p_0 - p_{m_0} t^{m_0}$. The retained coordinates are the coefficients in degrees $1, \dots, m_0 - 1$.

For three deficits, assume $k \geq 3$ and take $N = \text{span}\{\ell_0, \ell_1, \ell_\infty\}$. An explicit projection is

$$\begin{aligned} \Pi_3 p &= p - p(0)(1 - t) - p(1)t - p_{m_0}(t^{m_0} - t) \\ &= \sum_{j=2}^{m_0-1} p_j(t^j - t). \end{aligned} \quad (\text{D.7})$$

Its image consists exactly of the polynomials with zero evaluations at 0 and 1 and zero leading coefficient. Thus the coefficients of T_μ in degrees $2, \dots, m_0 - 1$, read out in the basis $\{t^j - t\}_{j=2}^{m_0-1}$, provide exactly $2k - 4$ retained coordinates.

For independent uniform $f, g \in \mathbb{F}_q[t]_{<k}$, success of this projection is equivalent to

$$f(0)g(0) = 0, \quad f(1)g(1) = 0, \quad f_{k-1}g_{k-1} = 0. \quad (\text{D.8})$$

When $k \geq 3$, the three linear observations of each polynomial are independent, so all six values are jointly uniform. The three success events are independent, each with probability ϑ_q , giving success probability ϑ_q^3 .

E Accounting Across Transcript Boundaries

A protocol may contain several public-randomness boundaries, with information generated at one cut available at later cuts. Local retained widths must therefore be compared relative to both the actual incoming state and the scalar directions already included in the cumulative accounting.

Throughout this appendix, work over a field $\mathbb{F}$ and fix one public-transcript branch. Challenges sampled before a selected boundary are treated as fixed public parameters. Source-dependent transcript messages remain functions of the semantic input; fixing public randomness does not turn them into constants.

E.1 Common Quotients and Width Preservation

At boundary v_j , let X_j be the local semantic domain, let $\mathcal{F}_j \leq \mathbb{F}^{X_j}$ be the declared-free scalar space, and write $Q_j := \mathbb{F}^{X_j} / \mathcal{F}_j$. Choose a common parameter

space X_* , maps $p_j : X_* \rightarrow X_j$, and a subspace $\mathcal{F}_* \leq \mathbb{F}^{X_*}$ satisfying

$$f \in \mathcal{F}_j \implies f \circ p_j \in \mathcal{F}_*. \quad (\text{E.1})$$

Pullback induces a linear map

$$\theta_j : Q_j \rightarrow Q_*, \quad [f] \mapsto [f \circ p_j], \quad Q_* := \mathbb{F}^{X_*} / \mathcal{F}_*. \quad (\text{E.2})$$

Let $I_j^{\text{loc}}, V_j^{\text{loc}} \leq Q_j$ be the local incoming and supported target spans. Put $I_j := \theta_j(I_j^{\text{loc}})$ and $V_j := \theta_j(V_j^{\text{loc}})$. The span of these finitely many images, together with the root and final-target images used below, gives a finite-dimensional accounting space $Q \leq Q_*$. The whole quotient Q_* need not be finite-dimensional.

Condition (E.1) makes the pullback map well-defined. It does not by itself guarantee preservation of the local residual width.

Lemma E.1 (Preservation of a residual quotient). *Put $E_j := I_j^{\text{loc}} + V_j^{\text{loc}}$. Then*

$$\dim \frac{E_j}{I_j^{\text{loc}}} - \dim \frac{I_j + V_j}{I_j} = \dim \frac{E_j \cap \ker \theta_j}{I_j^{\text{loc}} \cap \ker \theta_j}. \quad (\text{E.3})$$

Consequently, the local width is preserved exactly when

$$E_j \cap \ker \theta_j \subseteq I_j^{\text{loc}}. \quad (\text{E.4})$$

Proof. The induced map $E_j/I_j^{\text{loc}} \rightarrow \theta_j(E_j)/\theta_j(I_j^{\text{loc}})$ is surjective. Its kernel is

$$\frac{I_j^{\text{loc}} + (E_j \cap \ker \theta_j)}{I_j^{\text{loc}}} \cong \frac{E_j \cap \ker \theta_j}{I_j^{\text{loc}} \cap \ker \theta_j}.$$

Rank-nullity gives (E.3). The two widths are equal precisely when this kernel is zero, which is equivalent to (E.4). $\square$

Injectivity of θ_j on all of E_j is sufficient but not necessary. The exact requirement is that no direction newly needed relative to the local incoming span disappear under the common-space identification.

The final target requires the analogous condition relative to the root state. Let I_0^{loc} and V_G^{loc} be the local root and final-target spans, with identification θ_G satisfying $\theta_G(I_0^{\text{loc}}) = U(K_0)$ and $\theta_G(V_G^{\text{loc}}) = V(G)$ in the common accounting space. Preservation of the original final-target width holds when

$$(I_0^{\text{loc}} + V_G^{\text{loc}}) \cap \ker \theta_G \subseteq I_0^{\text{loc}}. \quad (\text{E.5})$$

Without this condition, the common-space calculation still applies to the image problem.

For the remainder of this appendix, assume (E.4) at each selected boundary. When the final-target width is interpreted in its original local model, also assume (E.5). For a state S , write $U(S) \leq Q$ for its scalar span, and for a target family H , write $V(H) \leq Q$ for its supported scalar target span.

E.2 Relative Width

For states K, S , define

$$w(S \mid K) := \dim_{\mathbb{F}} \frac{U(K) + U(S)}{U(K)}, \quad (\text{E.6})$$

and for a target family H , define

$$w(H \mid K) := \dim_{\mathbb{F}} \frac{U(K) + V(H)}{U(K)}. \quad (\text{E.7})$$

Theorem 3.1 gives $m_{\min}(H \mid K) = w(H \mid K)$ at a field-valued boundary. For several states, use the convention $U(S_1, \dots, S_t) := \sum_i U(S_i)$.

Proposition E.2 (Chain rule). *For states K, S_1, S_2 ,*

$$w(S_1, S_2 \mid K) = w(S_1 \mid K) + w(S_2 \mid K, S_1). \quad (\text{E.8})$$

Proof. Set $B_0 := U(K)$, $B_1 := U(K) + U(S_1)$, and $B_2 := U(K) + U(S_1) + U(S_2)$. The identity $\dim(B_2/B_0) = \dim(B_1/B_0) + \dim(B_2/B_1)$ is exactly (E.8). $\square$

In particular, $w(S \mid K) = 0$ whenever $U(S) \subseteq U(K)$. If $U(K) \subseteq U(K')$, then $w(S \mid K') \leq w(S \mid K)$ and $w(H \mid K') \leq w(H \mid K)$.

E.3 Local and Cumulative Accounting

Consider selected public-randomness boundaries $v_1, \dots, v_N$ along the fixed branch. Let K_0 be the root state used as the accounting baseline. At v_j , let K_j be the actual incoming state and H_j the target family to be supported after the challenge. Write $I_j := U(K_j)$ and $V_j := V(H_j)$.

The local retained width is

$$r_j := w(H_j \mid K_j) = \dim \frac{I_j + V_j}{I_j}. \quad (\text{E.9})$$

Define the cumulative target spans by

$$B_0 := U(K_0), \quad B_j := B_{j-1} + V_j \quad (1 \leq j \leq N), \quad (\text{E.10})$$

their increments by

$$\delta_j := \dim \frac{B_{j-1} + V_j}{B_{j-1}} = \dim \frac{B_j}{B_{j-1}}, \quad (\text{E.11})$$

and the incoming novelty by

$$e_j := \dim \frac{B_{j-1} + I_j}{B_{j-1}}. \quad (\text{E.12})$$

The quantity r_j measures what is required beyond the actual incoming span I_j ; δ_j measures what the current target adds to B_{j-1} . The spaces B_j contain the root baseline and earlier target directions, but they need not correspond to any stored protocol state. They are not enlarged merely because a direction appears in an incoming state.

Accordingly, e_j measures novelty relative to B_{j-1} , not relative to the entire transcript history. A direction can contribute to e_j at more than one boundary if it has not entered the cumulative target span.




Fig. 2. Local and cumulative accounting at v_j . The upper lane measures V_j relative to the actual incoming span I_j , giving r_j . The lower lane measures it relative to B_{j-1} , giving δ_j . The cumulative span need not be stored. Under $I_j \subseteq B_{j-1}$, one has $\delta_j \leq r_j$.

Proposition E.3 (Local and cumulative increments). *At every selected boundary,*

$$\delta_j \leq e_j + r_j. \quad (\text{E.13})$$

If $I_j \subseteq B_{j-1}$, then $e_j = 0$ and $\delta_j \leq r_j$. If $I_j = B_{j-1}$, then $\delta_j = r_j$.

Proof. Since $B_{j-1} + V_j \subseteq B_{j-1} + I_j + V_j$,

$$\begin{aligned} \delta_j &\leq \dim \frac{B_{j-1} + I_j + V_j}{B_{j-1}} \\ &= e_j + \dim \frac{B_{j-1} + I_j + V_j}{B_{j-1} + I_j}. \end{aligned}$$

The natural map

$$\frac{I_j + V_j}{I_j} \twoheadrightarrow \frac{B_{j-1} + I_j + V_j}{B_{j-1} + I_j}$$

is surjective, so the second term is at most r_j . This proves (E.13). The two special cases follow from the definitions. $\square$

We call $I_j \subseteq B_{j-1}$ the inherited-state condition. It permits the actual incoming state to retain only part of the previously accounted information.

E.4 Global Accounting

Theorem E.4 (Multi-boundary accounting). *For the selected sequence,*

$$\dim \frac{B_N}{B_0} = \sum_{j=1}^N \delta_j \leq \sum_{j=1}^N (e_j + r_j). \quad (\text{E.14})$$

If a final target family G satisfies $V(G) \subseteq B_N$, then

$$w(G \mid K_0) \leq \dim \frac{B_N}{B_0} \leq \sum_{j=1}^N (e_j + r_j). \quad (\text{E.15})$$

Under the inherited-state condition at every selected boundary, all e_j vanish and both upper bounds reduce to $\sum_{j=1}^N r_j$.

Proof. The nested chain $B_0 \subseteq \dots \subseteq B_N$ gives

$$\dim(B_N/B_0) = \sum_j \dim(B_j/B_{j-1}) = \sum_j \delta_j.$$

Applying Proposition E.3 termwise proves (E.14). If $V(G) \subseteq B_N$, then $w(G \mid K_0) = \dim((B_0 + V(G))/B_0) \leq \dim(B_N/B_0)$, yielding (E.15). Under inherited state, every $e_j = 0$. $\square$

The telescope in (E.14) is exact. Local widths bound the cumulative increments, with the exact correction derived below. This accounting concerns exact scalar spans; it does not provide a composition rule for the distributional reconstruction probabilities of Section 5.

E.5 Exact Local Slack

Let $P, U, V \leq Q$ represent the previous cumulative span, actual incoming span, and current target span. Define

$$e := \dim \frac{P+U}{P}, \quad r := \dim \frac{U+V}{U}, \quad \delta := \dim \frac{P+V}{P}. \quad (\text{E.16})$$

Proposition E.5 (Exact local correction). *With this notation,*

$$e + r - \delta = \dim \frac{U + (P \cap V)}{(P \cap U) + (U \cap V)}. \quad (\text{E.17})$$

Proof. The denominator is contained in U , hence in the numerator. The dimension formula gives

$$\begin{aligned} e + r - \delta &= \dim U - \dim(P \cap U) - \dim(U \cap V) + \dim(P \cap V), \\ \dim(U + (P \cap V)) &= \dim U + \dim(P \cap V) - \dim(P \cap U \cap V), \\ \dim((P \cap U) + (U \cap V)) &= \dim(P \cap U) + \dim(U \cap V) - \dim(P \cap U \cap V). \end{aligned}$$

Subtracting the last two identities proves (E.17). $\square$

Corollary E.6 (Slack under inherited state). *If $U \subseteq P$, then $e = 0$ and*

$$r - \delta = \dim \frac{P \cap V}{U \cap V}. \quad (\text{E.18})$$

Proof. Under $U \subseteq P$, the quotient in (E.17) becomes

$$\frac{U + (P \cap V)}{U} \cong \frac{P \cap V}{U \cap V}.$$

This gives (E.18). $\square$

The numerator consists of current target directions already present in the cumulative span. The denominator consists of those still carried by the actual incoming state. Their quotient measures the excess local charge over the cumulative increment.

E.6 Global Accounting with the Exact Correction

Using the notation of Appendix E.3, define

$$\sigma_j := \dim \frac{I_j + (B_{j-1} \cap V_j)}{(B_{j-1} \cap I_j) + (I_j \cap V_j)}. \quad (\text{E.19})$$

Proposition E.5 gives $\delta_j = e_j + r_j - \sigma_j$. Telescoping therefore yields

$$\dim \frac{B_N}{B_0} = \sum_{j=1}^N (e_j + r_j - \sigma_j). \quad (\text{E.20})$$

If $V(G) \subseteq B_N$, then

$$w(G \mid K_0) \leq \sum_{j=1}^N (e_j + r_j - \sigma_j) \leq \sum_{j=1}^N (e_j + r_j). \quad (\text{E.21})$$

When the left-hand side is interpreted in the original final-target model, condition (E.5) applies.

Under inherited state, $\sigma_j = \dim((B_{j-1} \cap V_j)/(I_j \cap V_j))$, and the cumulative dimension is $\dim(B_N/B_0) = \sum_j (r_j - \sigma_j)$.

Incoming novelty cannot be omitted in general. Take a one-dimensional Q , baseline $B_0 = 0$, and one boundary with $I_1 = V_1 = Q$. Then $r_1 = 0$, but $e_1 = \delta_1 = 1$: the target is available at the selected cut yet contributes one direction relative to the root baseline.

For a concrete protocol, the source partition and boundary convention determine the free information, the actual incoming state determines I_j , and the selected postchallenge outputs determine V_j . Appendix F applies these identifications to selected cuts in two folding protocols.

F Protocol Boundary Identifications

We apply the scalar-span framework to two different transcript effects. In the selected PikkuFold cut, the later folded statement is source-local with respect to the accumulator/fresh-source partition. In the selected NeutronNova sequence, later scalar claims are inherited linear observations of root claims or earlier polynomial messages.

F.1 PikkuFold

Consider the folding protocol Π_{fold} of PikkuFold [16, §3.2, Fig. 2]. Place the selected boundary after the terminal RingSC evaluation has been checked using the claims $(t'_i)_{i \in [k]}$ and t'_{acc} , and before the final challenge vector

$$z = (z_i)_{i \in [k]} \leftarrow \mathcal{C}^k \quad (\text{F.1})$$

is sampled.

All preceding verifier randomness is fixed at this cut, including the layered-projection, trace, and batching randomness, the RingSC batching randomness, and the terminal RingSC point. Write the latter as $r_d = (r_{d,0}, r_{d,1})$.

Take the complete accumulator opening as the left source x_0 , and the tuple of all fresh openings as the right source x_1 . On the semantic domain of source openings, the terminal claims are given by

$$t'_{\text{acc}} = \text{MLE}[w_{\text{acc}}](r_{d,1}), \quad t'_i = \text{MLE}[w_i](r_{d,1}), \quad i \in [k], \quad (\text{F.2})$$

where MLE denotes multilinear extension. Thus t'_{acc} depends only on x_0 , while each t'_i depends only on x_1 .

After sampling z , the folded quantities are

$$\begin{aligned} w_{\text{fold}} &= w_{\text{acc}} + \sum_{i \in [k]} z_i w_i, & y_{\text{fold}} &= y_{\text{acc}} + \sum_{i \in [k]} z_i y_i, \\ s_{\text{fold}} &= r_{d,1}, & t_{\text{fold}} &= t'_{\text{acc}} + \sum_{i \in [k]} z_i t'_i. \end{aligned} \quad (\text{F.3})$$

For the retained-width calculation, take the target family to be the folded statement update

$$H_z := (y_{\text{fold}}, s_{\text{fold}}, t_{\text{fold}}), \quad z \in \mathcal{C}^k. \quad (\text{F.4})$$

Let $H_{\text{late}} := (H_z)_{z \in \mathcal{C}^k}$.

The witness update has the same source decomposition.

For scalar accounting, regard the ring-valued statement coordinates as a finite-dimensional $\mathbb{F}_q$ -vector space and use all $\mathbb{F}_q$ -linear scalarizations. Multiplication by a fixed ring element z_i is $\mathbb{F}_q$ -linear. For any scalarization η , fixed public z and $r_{d,1}$ therefore give

$$\eta(H_z(x_0, x_1)) = a_{\eta,z}(x_0) + b_{\eta,z}(x_1) + \kappa_{\eta,z}, \quad (\text{F.5})$$

where $\kappa_{\eta,z}$ is public. The accumulator term collects observations of y_{acc} and t'_{acc} ; the fresh-source term collects observations of the y_i, t'_i with the fixed fold coefficients. The coordinate $s_{\text{fold}} = r_{d,1}$ contributes a constant. Absorbing that constant into either one-source term shows that every supported scalarization belongs to the direct-boundary free space.

Proposition F.1 (PikkuFold later-cut width). *At the selected boundary,*

$$V(H_{\text{late}}) = 0 \quad (\text{F.6})$$

in the direct-boundary quotient, and hence

$$r_{\text{late}} = m_{\min}^{\text{dir}}(H_{\text{late}}) = 0. \quad (\text{F.7})$$

Proof. Equation (F.5) expresses every supported scalar target as a sum of a function of x_0 and a function of x_1 . Its class modulo the source-local space is therefore zero. Theorem 3.1 gives (F.7). $\square$

The projection, trace, batching, and RingSC computations precede the selected boundary; the target is the subsequent z -dependent statement update.

F.2 NeutronNova

We next consider selected scalar cuts in NeutronNova [11]. The relevant mechanisms are SumFold and the folding of nested sum-check instances. For the scalar accounting below, each degree-bounded polynomial message is linearly recoverable from its emitted field coordinates together with the preceding incoming claim. If evaluation coordinates are used, their points must form an information set for the relevant polynomial space.

SumFold. Consider Construction 1 of NeutronNova [11, §4, Construction 1]. Let $k = 2^\nu$ be the number of input instances, and write their scalar claims as T_i , indexed by $i \in \{0, 1\}^\nu$. Include these claims in the root state K_0 . After the verifier samples $\rho \in \mathbb{F}^\nu$, the packed initial claim is

$$T(\rho) = \sum_{i \in \{0,1\}^\nu} \text{eq}(\rho, i) T_i, \quad (\text{F.8})$$

where eq is the multilinear equality polynomial. For fixed ρ , its coefficients are public, so

$$[T(\rho)] \in U(K_0). \quad (\text{F.9})$$

Unroll the standard sum-check invocation in Construction 1, writing its round polynomials as $g_1, \dots, g_\nu$ and verifier challenges as $\beta_1, \dots, \beta_\nu$. Put

$$c_0 := T(\rho), \quad c_j := g_j(\beta_j). \quad (\text{F.10})$$

The round consistency relation is

$$g_j(0) + g_j(1) = c_{j-1}. \quad (\text{F.11})$$

Evaluation at a fixed field point is linear in the coefficients of an emitted round polynomial [13].

Assign the scalar coordinates representing g_j to the target family of the stage at which it is emitted. The preceding incoming claim may supply a coefficient eliminated through (F.11). Under the recoverability condition above, once β_j is sampled,

$$c_j = g_j(\beta_j) \quad (\text{F.12})$$

is a linear observation of those message coordinates together with previously accounted claims. At any later selected cut v_ℓ , it therefore satisfies

$$[c_j] \in U(K_0) + \sum_{i < \ell} V(H_i). \quad (\text{F.13})$$

Here j indexes the sum-check round, while ℓ indexes the selected accounting cuts.

Construction 1 returns a terminal pair (c, r_b) and sets

$$T' = c \text{ eq}(\rho, r_b)^{-1}. \quad (\text{F.14})$$

In the local round notation, $c = c_\nu$ and $r_b = (\beta_1, \dots, \beta_\nu)$. On branches where the public inverse in (F.14) is defined, this normalization is another public scalar multiple of already accounted information.

Nested sum-check claims. Consider Construction 4 [11, §5.2, Construction 4]. The verifier samples $\gamma \in \mathbb{F}$, and the two input claim families are combined as

$$\hat{T}_i = T_i + \gamma T_{\text{pc},i}. \quad (\text{F.15})$$

For fixed γ , these are linear observations of the root scalar claims.

SumFold reduces the combined instances to a folded claim T_γ . The prover then sends two scalar claims separately, denoted here by

$$T_{\text{out}}, \quad T_{\text{pc,out}}, \quad (\text{F.16})$$

and the verifier checks

$$T_\gamma = T_{\text{out}} + \gamma T_{\text{pc,out}}. \quad (\text{F.17})$$

These values are denoted T, T_{pc} in Steps 4–5 of Construction 4.

Include both coordinates in the target family of the stage where they are emitted. Their separate later availability is inherited from those two scalar messages. Relation (F.17) alone would provide only their linear combination.

Subsequent nested sum-check rounds use the same assignment: each emitted polynomial belongs to the current target family, and its evaluation at a later public challenge is a linear observation of that earlier message.

Inherited-state inclusion. Let $v_1, \dots, v_N$ be the selected scalar cuts. The root state K_0 contains the original scalar claims, K_j is the selected incoming scalar state at v_j , and H_j contains the scalar coordinates of messages and outputs emitted in the stage following v_j .

Proposition F.2 (Inherited scalar state in the selected NeutronNova sequence). *Assume that each degree-bounded polynomial message is linearly recoverable from the scalar representation included in its target family together with the preceding incoming claim. In particular, evaluation representations must be information sets for their polynomial spaces. Then, at every selected later cut,*

$$U(K_j) \subseteq U(K_0) + \sum_{i < j} V(H_i) = B_{j-1}. \quad (\text{F.18})$$

Proof. Proceed in the order of the selected cuts. Every original scalar claim belongs to $U(K_0)$, as does any public linear combination of root claims, including (F.8) and (F.15).

Every other permitted incoming scalar observation is derived from an earlier polynomial or scalar message. The selected coordinates of that message belong to an earlier target span $V(H_i)$. By the message-recovery hypothesis, a later evaluation of the polynomial is a linear observation of those coordinates together with its preceding incoming claim. The induction hypothesis accounts for the latter. Equations (F.12) and (F.14) are the SumFold instances of these observations.

The two terminal claims in (F.16) are emitted separately and both belong to the target family at that stage. Later observations of either claim therefore also lie in the cumulative span. This proves (F.18). $\square$

Under the common-space and preservation conditions of Appendix E.1, Proposition F.2 supplies the inherited-state hypothesis of Theorem E.4. Polynomial messages and separate terminal claims may contribute positive local width when first emitted, while their later linear evaluations contribute no new cumulative direction.

G Linear Recovery with Bounded Storage and Adaptive Queries

The linear specialization of Section 3.2 has a related recovery problem once a uniform source distribution and later source access are specified. Here the decoder seeks a complete linear target from bounded stored information and adaptive scalar linear responses from the same source.

G.1 The Recovery Model

Let b, t be nonnegative integers, let X be uniform in $\mathbb{F}_q^n$, and fix public matrices $K \in \mathbb{F}_q^{p \times n}$ and $L \in \mathbb{F}_q^{\ell \times n}$. The decoder receives $U = KX$ and aims to recover $Z = LX$.

An encoder computes

$$S_{\text{st}} = \text{Enc}(K, L, X; R_{\text{enc}}) \in \mathcal{S}_{\text{st}}, \quad |\mathcal{S}_{\text{st}}| \leq q^b. \quad (\text{G.1})$$

The encoder may be nonlinear and computationally unrestricted. Encoder and decoder randomness may be correlated with each other, but their joint value $R = (R_{\text{enc}}, R_{\text{dec}})$ is independent of X .

The decoder receives K, L, U, S_{st} and its random coins. It may issue adaptive row queries $a \in \mathbb{F}_q^{1 \times n}$ to the same source, receiving only

$$\mathcal{O}_X(a) = aX \in \mathbb{F}_q. \quad (\text{G.2})$$

Every execution path contains at most t queries. Query choices, stopping, and the final output may depend arbitrarily on the decoder's initial information, coins, and previous responses.

Define

$$\begin{aligned} p^*(b, t) &:= \sup_{\text{Enc, Dec}} \Pr[\widehat{Z} = Z], \\ d &:= \text{rank} \begin{pmatrix} K \\ L \end{pmatrix} - \text{rank } K. \end{aligned} \quad (\text{G.3})$$

G.2 Conditional Target Distribution

Lemma G.1 (Conditional target distribution).

For every $u \in \text{im } K$, choose x_u with $Kx_u = u$. Then

$$Z \mid U = u \sim \text{Unif}(Lx_u + L(\ker K)), \quad (\text{G.4})$$

and this affine space contains exactly q^d values.

Proof. Conditioned on $U = u$, the uniform source X is uniform on the affine fiber $x_u + \ker K$. The fibers of $L|_{\ker K}$ are cosets of $\ker K \cap \ker L$ and have equal size. Applying L therefore gives (G.4). Moreover,

$$\begin{aligned} \dim L(\ker K) &= \dim \ker K - \dim(\ker K \cap \ker L) \\ &= \text{rank} \begin{pmatrix} K \\ L \end{pmatrix} - \text{rank } K \\ &= d. \end{aligned}$$

Thus the conditional target is uniform on q^d values. $\square$

The same conclusion holds after conditioning on source-independent randomness. It need not hold after conditioning on the nonlinear stored state.

G.3 The Exact Storage–Response Tradeoff

Theorem G.2 (Exact storage–response tradeoff). *In the model of Appendix G.1,*

$$p^*(b, t) = q^{-(d-b-t)_+}, \quad (x)_+ := \max\{x, 0\}. \quad (\text{G.5})$$

In particular, perfect recovery is possible exactly when $b + t \geq d$.

Proof. Upper bound. Fix an attainable value $U = u$ and a value $R = \rho$ of all source-independent coins. By Lemma G.1, the target remains uniform on q^d values. The encoder and decoder are now deterministic.

There are at most q^b stored-state values. For each such value, the adaptive interaction forms a q -ary response tree of depth at most t , with at most q^t leaves. Query addresses introduce no additional branching: each address is fixed by the initial information and previous responses.

There are therefore at most q^{b+t} complete observations, each specifying a stored-state value and a response-tree leaf. At any observation τ , the decoder outputs a fixed value z_τ , and

$$\begin{aligned} \Pr[\tau \wedge Z = z_\tau \mid U = u, R = \rho] &\leq \Pr[Z = z_\tau \mid U = u, R = \rho] \\ &\leq q^{-d}. \end{aligned} \tag{G.6}$$

Summing over observations bounds the conditional success probability by $\min\{1, q^{b+t-d}\}$. Averaging over U, R proves the upper bound.

Matching strategy. Choose rows $r_1, \dots, r_d \in \text{row } L$ whose classes form a basis of $(\text{row } K + \text{row } L) / \text{row } K$. Let R_* be the matrix of these rows. Then $\text{rank}[K; R_*] = \text{rank } K + d$, and there are public matrices C, E with

$$L = CK + ER_*. \tag{G.7}$$

Thus KX and R_*X together determine LX .

Applying Lemma G.1 with target R_*X , the conditional vector $R_*X \mid U = u$ is uniform on all of $\mathbb{F}_q^d$. Set $s := \min\{b, d\}$ and $t_0 := \min\{t, d - s\}$. The encoder stores the first s coordinates of R_*X . The decoder queries the next t_0 coordinates using the corresponding rows of R_* , and guesses each remaining coordinate as zero.

Conditioned on the observed coordinates, the unobserved coordinates remain independent and uniform, so all guesses are correct with probability

$$q^{-(d-s-t_0)} = q^{-(d-b-t)_+}. \tag{G.8}$$

Whenever the guesses are correct, (G.7) gives the correct target. This matches the upper bound. $\square$

The optimum is therefore attained by a nonadaptive linear strategy even though nonlinear storage, unrestricted decoding, and adaptive queries are permitted.

G.4 Relation to Retained Width

In the $\mathcal{F} = \{0\}$ linear special case of Section 3.2, when L stacks the supported output maps, the residual dimension d in (G.3) equals the quantity in (3.9).

The operational interfaces differ. Section 3 requires one prechallenge state supporting an entire challenge-dependent output family through linear readout. Appendix G fixes a uniform source, asks for the complete target LX , allows

an arbitrary bounded-range stored state and unrestricted decoder, and permits additional same-source scalar linear queries after storage.

Under these distributional and access assumptions, the shared residual dimension determines the exact recovery probability in (G.5).