LEARNING LATENT ALGEBRAIC STRUCTURE FROM AMBIGUOUS SET OBSERVATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

We study when statistically learnable latent structure can also be recovered efficiently, and how membership queries change the answer. An unknown support $A \subseteq \mathbb{F}_2^n$ has small additive doubling and is observed through a fixed set B satisfying $|A \triangle B| \le \eta |A|$. We seek one linear subspace V such that every compatible support A is covered by few V-cosets and satisfies $|V| \le |A|$. For every $\eta < 1$, polynomially many uniform samples suffice statistically, with cost polynomial in the doubling constant and proportional to $(1-\eta)^{-1}$; this radius dependence is sharp. Under a specified hardness assumption for learning parities with noise (search-LPN), however, no polynomial-time sample-only learner achieves even constant covering cost, including when the latent support is unique. At fixed structural parameters and the same constant covering budget, adding exact membership queries to B permits polynomial-time recovery. The general query learner constructs a short structural list and uses fresh samples to select one common output through a majority-coverage rule. Persistent structured cores make this candidate construction possible. At doubling one, a complementary distinction appears at $\eta = 1/3$: coarse recovery remains polynomial time, while exact recovery requires exponentially many accesses in the worst case when latent cardinality is unknown.

1 INTRODUCTION

When observations leave several latent explanations, a learner can seek a structural description valid for all of them. Two questions then arise: do passive samples contain enough information to learn such a description, and can it be computed efficiently? We study these questions in a discrete algebraic model and compare passive sampling with membership access to the observed support.

Ambiguity already arises when every observed point is genuine. Let $U, X \le \mathbb{F}_2^n$ have trivial intersection. The observation $B = U$ may come from either $A_0 = U$ or $A_1 = U \oplus X$ after deletions. If $\dim X = t$, both explanations are compatible at deletion radius $1 - 2^{-t}$. Access to B cannot reveal which unseen directions were present. Our target is a compact algebraic description shared by all compatible explanations.

Our latent support is a nonempty set $A \subseteq G := \mathbb{F}_2^n$ with small additive doubling, $|A + A| \le K|A|$. The observed support $B \subseteq G$ satisfies

$$|A \triangle B| \le \eta |A|, \qquad 0 \le \eta < 1. \tag{1}$$

A concrete example is a union of s distinct cosets of a subspace V_0. Such a set has doubling at most s and consists of shared linear directions together with s offsets. Small doubling allows more general supports, including sets that occupy only part of each covering coset.

We seek a linear subspace V whose cosets cover every compatible support:

$$A \subseteq \bigcup_{j=1}^{q} (z_j + V), \qquad |V| \le |A|.$$

The size constraint rules out the ambient space as a trivial answer. After quotienting by V, the support occupies at most q positions; each coset may contain only part of the support. The learned V must work simultaneously for all supports compatible with the same B, while the offsets may depend on the support. We call this a *common structural hypothesis*.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Why should one hypothesis serve several latent explanations? Deletions and insertions consume the same budget in equation 1, which implies

$$|B| \le (1+\eta)|A \cap B|, \qquad \frac{|A \cap B|}{|B|} \ge \frac{1}{1+\eta} > \frac{1}{2}. \tag{2}$$

Every explanation therefore regards a strict majority of B as genuine. More strongly, any set covering half of B intersects each compatible A in at least $(1-\eta)|A|/2$ points. For a union of a few parallel cosets, small doubling turns this overlap into a cover of all of A. This majority-coverage principle underlies both of our positive learning results.

Samples suffice statistically. Clean polynomial Freiman–Ruzsa (PFR) structure supplies a bounded-coset cover for each latent support (Gowers et al., 2025). We search over such descriptions, retain those covering an empirical majority, and choose one with minimum subspace dimension. Uniform convergence makes the coverage test reliable; minimum dimension enforces the size constraint for every compatible support. Polynomially many samples suffice for every $\eta < 1$, with cost polynomial in K and proportional to $(1-\eta)^{-1}$. The exhaustive search takes exponential time.

Computational hardness despite unique identification. Could a different sample-only algorithm find a coarse cover efficiently? Under search-LPN hardness at a specified constant noise rate, no polynomial-time learner achieves even constant covering cost, already at $K = 1$ and a fixed positive radius below $1/3$, where the latent affine support is unique. The lower bound allows polynomial-size output lists and knowledge of the latent cardinality.

Queries provide a constructive route. Adding exact membership queries to B lets us generate a short family of candidates containing a serving entry for every compatible support. Fresh samples then apply the majority-coverage test and minimum-dimension rule to select one common output. The general learner has polynomial sample complexity and XP query/time complexity. Under moderate corruption, a single core suffices and the bound is FPT in K. For fixed structural parameters the query learner is polynomial time; its cover bound is larger than the statistical bound. Section 2 makes the access comparison at a common fixed quality budget.

Writing $\varepsilon = 1 - \eta$ and $p(K)$ for a clean PFR covering bound, the results are:

Access	Regime	Guarantee
Samples	All $\eta < 1$	**Statistical sufficiency:** common V with cost $2Kp(K)/\varepsilon$; polynomial samples, exhaustive time.
Samples	Fixed $K = 1$, $0 < \eta < 1/3$	**Computational barrier:** no polynomial-time constant-cost learner under the stated LPN assumption.
Samples + MQ	$\eta \le c_{\mathrm{rob}}/\sqrt{K}$	**FPT computation:** one common V; FPT in K.
Samples + MQ	All $\eta < 1$	**XP computation:** one common V; polynomial samples, XP queries and time.

Exact recovery sharpens this contrast. At $K = 1$ and $\eta = 1/3$, recovering the unique latent support can require exponentially many accesses when its cardinality is unknown, whereas a coarse cover is polynomial-time computable. Section 7 develops this boundary.

Technical contribution. Candidate generation requires reusable access to a structured core of a potentially exponentially large support. The downstream PFR routine must query and sample the same realized set, with size and doubling guarantees for the entire core. We construct such a core at the sharp $\sqrt{\alpha}$ retained-mass scale (Reiher & Schoen, 2024). Weighted proposals make useful candidates accessible; shared sketches and independent validation and cleaning produce a finite record defining one deterministic subset. The record supports membership queries and exact uniform sampling conditional on finite-block success. Crucially for residual peeling, the extractor bounds the probability of returning an invalid core even without an energy promise, while high energy guarantees success with high probability. Earlier implicit BSG algorithms provide randomized tests and sampling with set-sandwich guarantees (Tulsiani & Wolf, 2014). Section 6 develops the fixed-set interface used here; Appendix M compares the related constructions.

Related work. The common-target viewpoint connects to latent identifiability under partial observation (Yao et al., 2024; Syrota et al., 2025). Robust and list-decodable learning also address ambiguity (Charikar et al., 2017; Raghavendra & Yau, 2020; Bakshi & Kothari, 2021); our insertion–deletion budget yields the majority property above and permits a common output. Finite-field subspace learning already captures LPN in a nested-mixture setting (Chen et al., 2021), and noisy-parity assumptions underlie earlier query-access comparisons (Feldman & Shah, 2009). Here the reduction targets coarse quotient covers of a fixed corrupted support. On the algorithmic side, we combine persistent core extraction with clean algorithmic PFR (Arunachalam et al., 2026; Castro-Silva et al., 2026). Appendix M compares the targets and oracle guarantees in detail.

2 Problem Formulation and Main Results

For fixed K, η and observation B, define the identification set

$$\mathcal{I}_{K,\eta}(B) := \{A \subseteq G : A \neq \varnothing,\ |A + A| \le K|A|,\ |A \triangle B| \le \eta|A|\}. \tag{3}$$

For a subspace $V \le G$, let $\pi_V : G \to G/V$ be the quotient map and $\mathcal{N}_V(A) := |\pi_V(A)|$ the number of V-cosets meeting A, and set

$$q_A(V) := \begin{cases} \mathcal{N}_V(A), & |V| \le |A|, \\ +\infty, & |V| > |A|. \end{cases} \tag{4}$$

Thus $q_A(V) \le q$ implies

$$\log_2 |A| - \log_2 q \le \dim V \le \log_2 |A|. \tag{5}$$

Equivalently, a linear map W with $\ker W = V$ maps A to at most q distinct values, since $|W(A)| = \mathcal{N}_V(A)$. A common-output guarantee requires one realized V to satisfy the cost bound for all $A \in \mathcal{I}_{K,\eta}(B)$ on the same success event.

The objective has a direct interpretation for an affine support $A = a + H$. In this case,

$$\mathcal{N}_V(a + H) = 2^{\dim H - \dim(H \cap V)}.$$

Thus $q_A(V) \le q$ requires V to contain all but at most $\log_2 q$ dimensions of the latent direction space H, while $|V| \le |H|$ controls its total dimension. At cost one, these conditions force $V = H$. Constant-cost recovery therefore relaxes exact direction recovery by only a constant number of dimensions. For general small-doubling supports, the coset cover extends this objective to sets with several offsets and incomplete occupancy within each coset.

For nonempty finite S, let U_S be uniform on S. We compare

$$\mathsf{Samp}(S) = \{\mathsf{Sample}_S\}, \qquad \mathsf{SampMQ}(S) = \{\mathsf{Sample}_S, \mathsf{MQ}_S\},$$

where sample calls are independent and exact, and $\mathsf{MQ}_S(x) = \mathbf{1}[x \in S]$. The learner accesses the fixed observed support B; the hidden A is never directly queryable.

Membership access lets the learner test additive relations at points constructed from its samples. For example, given $u, v, w \in B$, it can determine whether $u + v + w \in B$; this test is used by the weighted proposal mechanism in Section 6. The answer concerns one fixed observed support, so repeated queries at the same point always agree. The computational resource is the ability to choose these test points.

Fix an integer-valued nondecreasing polynomial $p(K)$ dominating the clean PFR covering bound (Gowers et al., 2025). The query algorithm uses the covering bound P_{PFR} of the specific algorithmic routine in Appendix C. To express statistical sharpness, write

$$Q_{\mathrm{com}}(B; K, \eta) := \min_{V \le G} \max_{A \in \mathcal{I}_{K,\eta}(B)} q_A(V)$$

when the identification set is nonempty. We use FPT for bounds of the form $f(\kappa)\,\mathrm{poly}(n, \log(1/\rho))$ with parameter-independent polynomial degree; XP permits that degree to depend on the structural parameters κ.

Theorem 1 (Statistical common structural learning and sharp ambiguity cost)**.** *Let $\varepsilon = 1 - \eta$ and assume $\mathcal{I}_{K,\eta}(B) \neq \varnothing$. Given K, η, ρ and access only to $\mathsf{Samp}(B)$, a randomized learner outputs a basis of $\widehat{V}$ such that*

$$\Pr\left[\forall A \in \mathcal{I}_{K,\eta}(B) : q_A(\widehat{V}) \leq \frac{2Kp(K)}{\varepsilon}\right] \geq 1 - \rho. \tag{6}$$

A sufficient sample bound is

$$O\left(\frac{(n+1)p(K)\log(2p(K))\log(2/\varepsilon) + \log(2/\rho)}{\varepsilon^2}\right). \tag{7}$$

An explicit exhaustive implementation is finite but exponential in n.

Moreover, for every integer $t \geq 1$, there are $K = 1$, $\eta_t = 1 - 2^{-t}$, and an observation B for which

$$Q_{\mathrm{com}}(B; 1, \eta_t) = (1 - \eta_t)^{-1}.$$

Thus the worst-case ε^{-1} dependence of the common structural objective is sharp up to a constant factor already at doubling one.

Theorem 2 (Conditional sample-only computational barrier)**.** *Fix constants $Q \geq 1$ and $0 < \eta_0 < 1/3$. Put $k = 1 + \lfloor \log_2 Q \rfloor$, $\theta = \frac{\eta_0}{8k}$.*

Assume search-LPN with a uniformly random secret and noise rate θ has no randomized polynomial-time algorithm with constant success probability. Then there is no randomized polynomial-time learner which, for every valid $K = 1$ instance (A, B) with $|A \triangle B| \leq \eta_0 |A|$, receives only independent samples from U_B and with constant success probability outputs a polynomial-size list of explicitly represented subspaces containing some V with $q_A(V) \leq Q$. The conclusion still holds if the learner is given $|A|$ or $\dim A$. Since $\eta_0 < 1/3$, every valid observation in this theorem has at most one compatible doubling-one support.

Theorem 3 (Query-assisted common structural learning)**.** *There are absolute constants and computable nondecreasing functions with the following properties.*

Moderate corruption. *If $\mathcal{I}_{K,c_{\mathrm{rob}}/\sqrt{K}}(B) \neq \varnothing$, then from $\mathsf{SampMQ}(B)$ a randomized learner outputs one V such that, with probability at least $1 - \rho$,*

$$\forall A \in \mathcal{I}_{K,c_{\mathrm{rob}}/\sqrt{K}}(B) : q_A(V) \leq P_{\mathrm{rob}}(K),$$

with sample, query, and time complexity $F_{\mathrm{rob}}(K)\,\mathrm{poly}(n, \log(1/\rho))$.

Arbitrary corruption. *Let $\varepsilon = 1 - \eta$ and assume $\mathcal{I}_{K,\eta}(B) \neq \varnothing$. From $\mathsf{SampMQ}(B)$ a randomized learner outputs one V such that, with probability at least $1 - \rho$,*

$$\forall A \in \mathcal{I}_{K,\eta}(B) : q_A(V) \leq CK^{5/2}\varepsilon^{-5/2} P_{\mathrm{PFR}}\left(C'K^4\varepsilon^{-4}\right). \tag{8}$$

The number of uniform samples is polynomial in $K, \varepsilon^{-1}, n, \log(1/\rho)$, while membership-query complexity and ordinary running time are at most

$$G(K, \varepsilon^{-1})\left(1 + n + \log(1/\rho)\right)^{O(\sqrt{K/\varepsilon})}. \tag{9}$$

Comparison at a common quality budget. Fix $K = 1$ and a constant $\eta_* \in (0, 1/3)$. Choose a constant Q_* dominating the cover bounds of Theorems 1 and 3. At this same quality, samples suffice statistically; under search-LPN hardness at noise $\eta_*/(8(1 + \lfloor \log_2 Q_* \rfloor))$, Theorem 2 excludes polynomial-time sample-only learning. Membership access permits polynomial-time learning because the structural parameters are fixed. This comparison holds where the latent support is uniquely identified.

3 Statistical Common Structure from Samples Alone

The selection rule has two parts: majority coverage controls the quotient cost, and minimum dimension enforces the size constraint. Let $p = p(K)$ and consider descriptions $(V, z_1, \ldots, z_p)$ with induced sets

$$Z = \bigcup_{j=1}^{p} (z_j + V), \qquad \mathcal{Z}_p := \{Z : V \leq G,\ z_j \in G\}.$$

At the population level, retain descriptions with $U_B(Z) \geq 1/2$ and choose one minimizing $\dim V$. Every compatible A has a clean PFR witness (V_A, Z_A) with $A \subseteq Z_A$ and $|V_A| \leq |A|$. By equation 2, $U_B(Z_A) \geq 1/(1+\eta) > 1/2$, so every witness is feasible. The selected V consequently satisfies $|V| \leq |A|$ for all compatible supports.

Now fix any retained description. Writing $D = A \setminus B$ and $O = B \setminus A$, its clean overlap obeys

$$|A \cap B \cap Z| \geq \frac{|B|}{2} - |O| = \frac{|A| - |D| - |O|}{2} \geq \frac{\varepsilon}{2}|A|. \tag{10}$$

One of the p cosets contains a slice $C \subseteq A$ of size at least $\varepsilon|A|/(2p)$. Since $C + C \subseteq V$ and $|A + C| \leq K|A|$, Ruzsa covering gives

$$\mathcal{N}_V(A) \leq \frac{|A + C|}{|C|} \leq \frac{2Kp}{\varepsilon}. \tag{11}$$

This argument applies to every compatible A for the same selected description.

The sample learner retains descriptions with empirical mass at least $1/2 + \varepsilon/8$. On the uniform event

$$\sup_{Z \in \mathcal{Z}_p} |\widehat{P}(Z) - U_B(Z)| \leq \varepsilon/8, \tag{12}$$

every clean witness passes because $1/(1+\eta) \geq 1/2 + \varepsilon/4$, and every survivor has population mass at least $1/2$. The preceding proof then applies unchanged. Affine subspaces have VC dimension at most $n + 1$, and unions of p affine subspaces have VC dimension $O((n+1)p\log(2p))$. Parallel unions form a subclass, giving the sample bound in equation 7. Appendix H supplies the full argument.

Sharpness has a simple example. Take $G = \mathbb{F}_2^t$, $B = \{0\}$, and $\varepsilon = 2^{-t}$. Both $A_0 = \{0\}$ and $A_1 = G$ are compatible at $K = 1$. The size constraint for A_0 forces a common output to be $V = \{0\}$, whose cost on A_1 is ε^{-1}. Conversely this V costs at most $|G|$ on any subset, proving equality. Appendix K also exhibits directional ambiguity among equally large latent supports.

Samples suffice, but the selection rule is computationally exhaustive. Can passive samples support an efficient learner?

4 A Conditional Computational Barrier from LPN

Search-LPN asks to recover s from examples $(x, \langle s, x\rangle \oplus e)$, with uniform x and independent $e \sim \text{Bernoulli}(\theta)$. The subspace $\mathbb{F}_2^d \times \{0\}$ covers any single parity graph with two cosets, independently of s. We use several blocks so that every cost-Q output reveals a short secret list, then couple noisy examples to one fixed support.

Even a coarse cover reveals a short secret list. Let $s \in \mathbb{F}_2^d$ be the LPN secret, put $\ell = \lfloor \log_2 Q \rfloor$ and $k = \ell + 1$, and define

$$A_s = \left\{(x_1, \ldots, x_k, \langle s, x_1\rangle, \ldots, \langle s, x_k\rangle) : x_i \in \mathbb{F}_2^d\right\}. \tag{13}$$

Its dimension is $D = kd$. If $q_{A_s}(V) \leq Q$, then

$$\dim(A_s \cap V) \geq D - \ell, \qquad D - \ell \leq \dim V \leq D.$$

Hence $\dim V^\perp \leq k + \ell$ and $\dim(A_s^\perp \cap V^\perp) \geq k - \ell = 1$. Moreover,

$$A_s^\perp = \{(a_1 s, \ldots, a_k s, a_1, \ldots, a_k) : a \in \mathbb{F}_2^k\}.$$

Enumerating the at most $2^{k+\ell} \leq 2Q^2$ vectors of $V^\perp$ therefore gives at most $O(kQ^2)$ candidate secrets, including s.

Realizing the samples through one fixed support. For each complete tuple $x \in \mathbb{F}_2^D$, independently draw and freeze a k-bit noise vector $E(x)$ with Bernoulli(θ) coordinates. Let B_E be the resulting noisy graph. At $\theta = \eta_0/(8k)$, a tuple changes with probability at most $\eta_0/8$, so

$$\Pr_E[|A_s \triangle B_E| > \eta_0 |A_s|] \leq \exp(-\Omega(\eta_0^2 2^D)).$$

We draw the table without conditioning on the corruption bound and include invalid tables in the failure probability. Couple fresh grouped LPN examples to the experiment that first draws E and then samples uniformly from B_E. Until a complete input tuple repeats, the transcripts coincide. For at most T requests, the collision probability bounds their total variation distance by $\binom{T}{2}2^{-D}$. A polynomial-time structural learner would thus return a serving subspace with constant probability. Discard entries outside the prescribed dimension range, decode the remaining list as above, and validate the resulting secret candidates on fresh LPN samples. The true secret has mismatch rate θ and every false secret has rate $1/2$, yielding a polynomial-time search-LPN algorithm. Appendix I gives the amplification and error accounting.

The hard instances have a uniquely identified latent support, and the construction preserves hardness when its cardinality is supplied.

5 Query-Assisted Common Learning

Membership queries let us replace exhaustive search by a short candidate family. Once the family contains a serving entry for every compatible support, fresh samples apply the same majority-coverage and minimum-dimension rule as in Section 3. The main task is to construct those candidates when the observation itself can have large doubling.

Extracting enough structured candidates. For a finite set S, its additive energy

$$E(S) := |\{(a, b, c, d) \in S^4 : a + b = c + d\}|$$

counts repeated additive relations. Write $C_A = A \cap B$. If a residual $R \subseteq B$ retains at least half of C_A for some compatible A, then

$$E(R) \geq \frac{|A \cap R|^4}{|A + A|} \geq \frac{|C_A|^4}{16K|A|}.$$

Using $|R| \leq |B| \leq (1 + \eta)|C_A|$ and $|C_A| \geq \varepsilon|A|$ gives

$$\frac{E(R)}{|R|^3} \geq \frac{\varepsilon}{16K(1 + \eta)^3} \geq \frac{\varepsilon}{128K} =: \alpha_\star. \tag{14}$$

At this energy level, the persistent BSG procedure of Section 6 extracts a core of relative size $\gamma_\star = c_{\mathrm{BSG}}\sqrt{\alpha_\star} = \Theta(\sqrt{\varepsilon/K})$.

We repeatedly remove such cores. The residual-density test permits continuation only while the residual has constant density in B. Thus each valid peel removes an $\Omega(\sqrt{\varepsilon/K})$ fraction of B, and $M_\star = O(\sqrt{K/\varepsilon})$ rounds suffice. On one joint good event, each normal stopping route—a small residual, compiler failure under its completeness guarantee, or reaching the peel cap—implies

$$\forall A \in \mathcal{I}_{K,\eta}(B) : |A \cap R_{\mathrm{fin}}| < \tfrac{1}{2}|A \cap B|. \tag{15}$$

Rejection exhaustion and failed PFR calls are charged as algorithmic failures. Appendix E specifies the finite tests and budgets.

The removed cores are disjoint, so every compatible A has a serving core with

$$|A \cap Y_j| \geq \frac{\varepsilon}{2M_\star}|A|.$$

From observed cores to latent covers. Each extracted core has doubling at most $L_\star = C_{\mathrm{BSG}}\alpha_\star^{-4}$. Apply the clean algorithmic PFR routine to obtain a subspace H_j with $|H_j| \leq |Y_j|$ and $\mathcal{N}_{H_j}(Y_j) \leq P_\star$, where $P_\star = P_{\mathrm{PFR}}(L_\star)$.

Fix a compatible support A and one of its serving cores Y_j, and put $\delta = \varepsilon/(2M_\star)$. Among the at most $P_\star$ occupied H_j-cosets, one contains a slice $C \subseteq A \cap Y_j$ with $|C| \geq \delta|A|/P_\star$. Points of this slice differ only along H_j, so $C + C \subseteq H_j$. Small doubling now transfers the core cover to the whole latent support:

$$|A + C| \leq K|A|, \qquad \mathcal{N}_{H_j}(A) \leq \frac{|A + C|}{|C|} \leq \frac{KP_\star}{\delta},$$

where the middle inequality follows from Ruzsa covering. The clean slice is used only in the analysis; the learner never needs to identify which core points belong to A.

It remains to enforce the output-size constraint. The available comparison is $|H_j| \le |Y_j| \le |B| < 2|A|$, and it holds for every compatible support. If $H_j \neq \{0\}$, remove one vector from its returned basis to form an index-two subspace V_j; otherwise set $V_j = H_j$. This fixed rule guarantees $|V_j| \le |A|$ simultaneously for all compatible A, without knowing their cardinalities. Each H_j-coset splits into at most two V_j-cosets, so every support served by Y_j satisfies $q_A(V_j) \le 2KP_\star/\delta = 4KM_\star P_\star/\varepsilon$. Applying this construction to every extracted core therefore gives subspaces satisfying

$$m = O(\sqrt{K/\varepsilon}), \qquad \forall A \in \mathcal{I}_{K,\eta}(B)\ \exists i(A) : q_A(V_{i(A)}) \le Q_{\mathrm{list}}, \tag{16}$$

where we choose the integer budget

$$Q_{\mathrm{list}} := \left\lceil C_{\mathrm{list}} K^{3/2} \varepsilon^{-3/2} P_{\mathrm{PFR}}(C_{\mathrm{list,L}} K^4 \varepsilon^{-4}) \right\rceil. \tag{17}$$

Direct sample complexity is polynomial, but evaluating the implicit sets has a recursive cost. Membership in

$$R_j = B \setminus \bigcup_{i<j} Y_i$$

requires evaluating earlier core predicates. Each core predicate in turn calls the residual oracle from which it was constructed. If T_j bounds the cost of the jth core predicate, the implementation admits a bound of the form $T_j \le a(1 + \sum_{i<j} T_i)$, where a is polynomial in n and $\log(1/\rho)$ for fixed structural parameters. Over $M_\star = O(\sqrt{K/\varepsilon})$ rounds, this gives the parameter-dependent exponent in equation 9. The final majority filter adds polynomial work. Appendix E gives the resource bounds.

Selecting one common output. The serving entry in equation 16 may depend on A. The majority-coverage rule turns this family into one common hypothesis.

Lemma 4 (Fresh-sample conversion of a universal list)**.** *Let $Q \in \mathbb{N}_{\ge 1}$. Suppose a fixed explicit list $V_1, \dots, V_m$ satisfies*

$$\forall A \in \mathcal{I}_{K,\eta}(B)\ \exists i : q_A(V_i) \le Q.$$

Then a postprocessor using only fresh samples from B and no new membership queries outputs one list entry $V_\star$ such that, with probability at least $1 - \zeta$,

$$\forall A \in \mathcal{I}_{K,\eta}(B) : q_A(V_\star) \le \frac{2KQ}{\varepsilon}.$$

Its extra sample and time complexity is polynomial in $m, Q, \varepsilon^{-1}, \log(1/\zeta)$ and n; using VC uniform convergence, the sample count can be made independent of n apart from the representation of the returned subspaces.

Condition on the past that fixes the candidate list. For each V_i, group fresh samples by their V_i-coset and let Z_i be the union of the at most Q most frequent observed cosets. Retain indices with $\widehat{P}(Z_i) \ge 1/2 + \varepsilon/8$ and return a retained entry of minimum dimension, using fixed tie-breaking. If none survives, return failure. Computing these scores requires only coset identification and frequency counts.

Consider one event on which empirical mass is accurate to $\varepsilon/8$ for every union of at most Q cosets of every listed subspace. Fix any compatible A and its serving entry $V_{i(A)}$. Some union Z_A of at most Q of its cosets contains A. Since the top-frequency union maximizes empirical coverage for that entry,

$$\widehat{P}(Z_{i(A)}) \ge \widehat{P}(Z_A) \ge \frac{1}{1+\eta} - \frac{\varepsilon}{8} \ge \frac{1}{2} + \frac{\varepsilon}{8}.$$

Every compatible support therefore has a surviving serving entry. Minimum dimension gives, simultaneously,

$$|V_\star| \le |V_{i(A)}| \le |A| \qquad \text{for every } A \in \mathcal{I}_{K,\eta}(B).$$

The selected union has true mass at least $1/2$, so equation 10 and the covering argument of Section 3 give the common cost bound.

The concentration event ranges over coset unions of the fixed candidates. For each V_i, this class has VC dimension at most Q, so

$$O\big(\varepsilon^{-2}\,[Q\log(2/\varepsilon)+\log(2m/\zeta)]\big)$$

fresh samples suffice. The same event handles every compatible A. Applying the lemma with equation 17 proves the general query theorem; Appendix F gives the full accounting.

Under moderate corruption, one core already captures clean mass from every compatible support. Its retained fraction is $\Theta(K^{-1/2})$, larger by a constant factor than the allowed contamination. The size-based purity bound of Appendix D then supplies a common lifting guarantee directly. With one core there is no residual recursion, yielding the FPT bound; Appendix G gives the shared success event and resource analysis.

6 A Persistent Structured-Core Compiler

After extracting a core, the query learner passes it to the PFR routine, which needs membership queries and uniform samples from that same set. The key requirement is a size and doubling guarantee for the entire set defined by the returned record, valid throughout all later adaptive calls. Peeling can also reach residuals below the energy threshold, so the extraction routine must control invalid outputs on those inputs as well.

Theorem 5 (Persistent structured-core compiler)**.** *There are absolute constants $c_{\mathrm{BSG}}, C_{\mathrm{BSG}} > 0$ and a randomized algorithm which, given $\mathsf{SampMQ}(S)$ for a nonempty $S \subseteq \mathbb{F}_2^n$ and parameters $\alpha, \rho_{\mathrm{snd}}, \rho_{\mathrm{cmp}} \in (0,1)$, always terminates and returns either* FAIL *or a finite record D. Every nonfailure D defines one deterministic set $Y_D \subseteq S$. Without any energy promise,*

$$\Pr\big[D \neq \textsc{Fail} \text{ and } \big(|Y_D| < c_{\mathrm{BSG}}\sqrt{\alpha}|S| \text{ or } |Y_D + Y_D| > C_{\mathrm{BSG}}\alpha^{-4}|Y_D|\big)\big] \le \rho_{\mathrm{snd}}.$$

If $E(S) \ge \alpha|S|^3$, then $\Pr[D = \textsc{Fail}] \le \rho_{\mathrm{cmp}}$. Once D is fixed, membership in Y_D is deterministic and has controlled worst-case query cost. Moreover, for any fixed nonfailure record satisfying the retained-size guarantee $|Y_D| \ge c_{\mathrm{BSG}}\sqrt{\alpha}|S|$, any prescribed number t of downstream samples, and failure budget ζ, independent fixed-length rejection blocks succeed with probability at least $1-\zeta$; conditioned on their success, the returned sample vector has law exactly $U_{Y_D}^{\otimes t}$.

The construction separates finding a promising core from certifying the set that will be returned. In each of two multiplicity regimes, a counting argument assigns noticeable mass to useful candidates under an exactly sampleable weighted distribution.

For the high-multiplicity branch, put $r(d) = |S \cap (S+d)|$. Good fibres have mass at least $\Omega(\alpha)$ under $\mu_E(d) = r(d)^2/E(S)$. This law has a simple implementation: draw $U, V, W \sim U_S$, set $D = U + V$, and accept when $W + D \in S$. Then

$$\Pr[D = d \text{ and accept}] = \frac{r(d)^2}{|S|^3}, \qquad \Pr[D = d \mid \text{accept}] = \frac{r(d)^2}{E(S)}.$$

Thus the proposal samples by squared multiplicity without estimating that multiplicity first.

For low multiplicities, a visible-layer argument gives $O(1 + \log(1/\sqrt{\alpha}))$ thresholds to search. At a useful layer, sampling centres with probability proportional to squared relation degree puts constant mass on good candidates.

Shared sketches make the resulting predicates reliable throughout their later use. For example, define the translation density $p(d) = \Pr_{X \sim U_S}[X + d \in S]$. One block $X_1, \dots, X_M \sim U_S$ defines

$$\widehat{p}(d) = \frac{1}{M}\sum_{j=1}^{M} \mathbf{1}_S(X_j + d).$$

Hoeffding's inequality and a union bound over the 2^n translations give

$$\Pr\left[\sup_{d \in \mathbb{F}_2^n} |\widehat{p}(d) - p(d)| > \tau\right] \le \delta \quad \text{for} \quad M = O\big(\tau^{-2}\big(n + \log(1/\delta)\big)\big)\,.$$

The record stores the sample block and evaluates this estimate only where needed. Its accuracy already holds at every translation, including those chosen adaptively later. The relation sketches and sampled codegree graph use such simultaneous control and are fixed before candidate validation.

Fresh samples test candidate density and average bad-pair rate. An independent cleaning sample turns the latter into a pointwise condition on every accepted element, yielding a deterministic small-doubling criterion. The stored sketches, candidate, and cleaning data define Y_D; the eventual membership predicate uses no fresh randomness. This gives soundness on arbitrary inputs. Under high energy, the weighted proposal bounds ensure that a good candidate is found with high probability.

For a valid record, the retained-size bound gives acceptance probability at least $c_{\mathrm{BSG}}\sqrt{\alpha}$ when rejecting samples from S. Independent blocks of length $\lceil (c_{\mathrm{BSG}}\sqrt{\alpha})^{-1}\log(t/\zeta)\rceil$ supply the exact product law conditional on all blocks succeeding. The same compiled set therefore supports both downstream access modes. Appendix B gives the construction and worst-case bounds. The comparison with earlier implicit BSG access procedures (Tulsiani & Wolf, 2014) is detailed in Appendix M.

7 Exact Recovery at the Critical Radius

The common-cover objective also changes the cost of recovery when the latent support is unique. At doubling one, every support is affine, and exact reconstruction has the following boundary.

Proposition 6 (Exact recovery at doubling one)**.** *Suppose $K = 1$.*

1. *If $0 \le \eta < 1/3$ and $\mathcal{I}_{1,\eta}(B) \neq \varnothing$, then its unique affine support can be recovered from $\mathsf{SampMQ}(B)$ in $\mathrm{poly}(n, (1-3\eta)^{-1}, \log(1/\rho))$ time and oracle calls.*
2. *At $\eta = 1/3$, every nonempty identification set is still a singleton, but any randomized classical algorithm that exactly recovers its member with probability at least $2/3$ on every instance requires $\Omega(2^n)$ total sample and membership-query calls in the worst case when the latent cardinality is not supplied.*
3. *For every $\eta > 1/3$, there are observations with two distinct compatible affine supports, so exact recovery is not information-theoretically possible on every instance.*

In contrast, the coarse common structural output of Theorem 3 is polynomial-time computable at the fixed critical parameters $K = 1$, $\eta = 1/3$.

Below $1/3$, use the common learner at the fixed radius $1/3$ to obtain a constant-cost subspace. Fresh samples propose a containing affine flat W of size $O(|A|)$. Containment can be tested because the coverage gap between containing and noncontaining flats is

$$\frac{1}{1+\eta} - \frac{1}{2(1-\eta)} = \frac{1-3\eta}{2(1-\eta^2)} > 0.$$

Inside a certified flat, write $p = |A|/|W|$, $f = \mathbf{1}_{B\cap W}$, and $\mu = \mathbb{E}_W f$. The threshold $\mu/2$ separates Fourier characters defining the affine support from all others by margin $(1-3\eta)p/2$. Goldreich–Levin search (Goldreich & Levin, 1989) and coefficient estimation recover the affine equations.

At $\eta = 1/3$, two observed supports differing at a single point can have distinct unique latent explanations: nested subspaces of sizes a and $2a$. Distinguishing the resulting oracles requires exponentially many accesses in the worst case. Appendix J proves these claims and studies how candidate lists and knowledge of latent cardinality change the exact-recovery boundary.

Discussion. Even when the latent support is uniquely identified, efficient recovery can depend on how the learner accesses the observation. The statistical and query learners use the same majority-coverage rule. The statistical learner searches exhaustively over coset covers; the query learner constructs a short family to which the rule can be applied. The specified search-LPN assumption rules out an efficient passive learner at the matched covering budget.

The next questions are whether the general query learner can match the statistical covering bound and improve its XP dependence to FPT. Recursive residual predicates cause the current exponent, making their evaluation cost a concrete target. Appendix N discusses these questions and extensions beyond exact uniform samples and exact membership access.

AI USE STATEMENT

LLM-based tools were used during the development and preparation of this manuscript for mathematical discussion and checking, literature discovery, and writing assistance. In the technical development, they were used to discuss formulations and possible refinements of bounds and proof arguments, and to check algebraic manipulations, quantifier order, parameter dependence, failure-probability accounting, and assumptions in compositions of randomized algorithms. They were also used to improve proof exposition and theorem statements. All suggestions were independently assessed by the author, and all final mathematical claims and proofs were manually checked; Lean was additionally used during verification.

LLMs were also used to identify and compare potentially relevant literature and to assist with organization, drafting, wording, LaTeX formatting, and notation consistency. Retained references were checked against the original sources. The work contains no synthetic datasets or empirical experiments. The author takes responsibility for all mathematical statements, proofs, citations, and AI-assisted text in the final manuscript.

REFERENCES

Dana Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988. doi: 10.1023/A:1022821128753.

Srinivasan Arunachalam, Davi Castro-Silva, Arkopal Dutt, and Tom Gur. Classical and quantum Polynomial Freiman–Ruzsa algorithms. In *17th Innovations in Theoretical Computer Science Conference*, volume 362 of *Leibniz International Proceedings in Informatics*, pp. 11:1–11:8, 2026. doi: 10.4230/LIPIcs.ITCS.2026.11.

Pranjal Awasthi, Vitaly Feldman, and Varun Kanade. Learning using local membership queries. In *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pp. 398–431, 2013. URL https://proceedings.mlr.press/v30/Awasthi13.html.

Ainesh Bakshi and Pravesh K. Kothari. List-decodable subspace recovery: Dimension independent error in polynomial time. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms*, pp. 1279–1297, 2021. doi: 10.1137/1.9781611976465.78.

Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM*, 50(4):506–519, 2003. doi: 10.1145/792538.792543.

Nader H. Bshouty. On one-sided testing affine subspaces. *Theoretical Computer Science*, 1013:114745, 2024. doi: 10.1016/j.tcs.2024.114745.

Davi Castro-Silva, Jop Briët, Srinivasan Arunachalam, Arkopal Dutt, and Tom Gur. An algorithmic Polynomial Freiman–Ruzsa theorem. *arXiv preprint arXiv:2604.04547*, 2026. doi: 10.48550/arXiv.2604.04547.

Timothy M. Chan and Moshe Lewenstein. Clustered integer 3SUM via additive combinatorics. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*, pp. 31–40, 2015. doi: 10.1145/2746539.2746568.

Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 47–60, 2017. doi: 10.1145/3055399.3055491.

Aidao Chen, Anindya De, and Aravindan Vijayaraghavan. Learning a mixture of two subspaces over finite fields. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132 of *Proceedings of Machine Learning Research*, pp. 481–504. PMLR, 2021.

Anindya De, Shivam Nadimpalli, and Rocco A. Servedio. Approximating sumset size. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2339–2357, 2022. doi: 10.1137/1.9781611977073.94.

Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1047–1060, 2018. doi: 10.1145/3188745.3188758.

Vitaly Feldman. Attribute-efficient and non-adaptive learning of parities and dnf expressions. *Journal of Machine Learning Research*, 8(51):1431–1460, 2007.

Vitaly Feldman and Shrenik Shah. Separating models of learning with faulty teachers. *Theoretical Computer Science*, 410(19):1903–1912, 2009. doi: 10.1016/j.tcs.2009.01.017.

Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, pp. 25–32, 1989. doi: 10.1145/73007.73010.

Parikshit Gopalan, Adam R. Klivans, and David Zuckerman. List-decoding reed–muller codes over small fields. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pp. 265–274, 2008. doi: 10.1145/1374376.1374417.

W. T. Gowers, Ben Green, Freddie Manners, and Terence Tao. On a conjecture of marton. *Annals of Mathematics*, 201(2):515–549, 2025. doi: 10.4007/annals.2025.201.2.5.

Ishay Haviv and Oded Regev. The list-decoding size of fourier-sparse boolean functions. In *30th Conference on Computational Complexity*, volume 33 of *Leibniz International Proceedings in Informatics*, pp. 58–71, 2015. doi: 10.4230/LIPIcs.CCC.2015.58.

Aapo Hyvarinen, Hiroaki Sasaki, and Richard E. Turner. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 859–868, 2019. URL https://proceedings.mlr.press/v89/hyvarinen19a.html.

Jeffrey C. Jackson, Eli Shamir, and Clara Shwartzman. Learning with queries corrupted by classification noise. *Discrete Applied Mathematics*, 92(2–3):157–175, 1999. doi: 10.1016/S0166-218X(99)00045-1.

Ilyes Khemakhem, Diederik P. Kingma, Ricardo Pio Monti, and Aapo Hyvarinen. Variational autoencoders and nonlinear ICA: A unifying framework. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2207–2217, 2020. URL https://proceedings.mlr.press/v108/khemakhem20a.html.

Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4114–4124, 2019. URL https://proceedings.mlr.press/v97/locatello19a.html.

Prasad Raghavendra and Morris Yau. List decodable subspace recovery. In *Proceedings of the Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 3206–3226, 2020. URL https://proceedings.mlr.press/v125/raghavendra20a.html.

Christian Reiher and Tomasz Schoen. Note on the theorem of balog, szemerédi, and gowers. *Combinatorica*, 44(3):691–698, 2024. doi: 10.1007/s00493-024-00092-5.

Zhao Song. An output-size-optimal algorithmic balog-szemerédi-gowers theorem. *Electronic Colloquium on Computational Complexity, Report TR26-148*, 2026. URL https://eccc.weizmann.ac.il/report/2026/148/.

Stas Syrota, Yevgen Zainchkovskyy, Johnny Xi, Benjamin Bloem-Reddy, and Søren Hauberg. Identifying metric structures of deep latent variable models. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 58065–58087, 2025. URL https://proceedings.mlr.press/v267/syrota25a.html.

Madhur Tulsiani and Julia Wolf. Quadratic goldreich–levin theorems. *SIAM Journal on Computing*, 43(2):730–766, 2014. doi: 10.1137/12086827X.

Danru Xu, Dingling Yao, Sebastien Lachapelle, Perouz Taslakian, Julius von Kügelgen, Francesco Locatello, and Sara Magliacane. A sparsity principle for partially observable causal representation learning. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 55389–55433, 2024. URL https://proceedings.mlr.press/v235/xu24ac.html.

Dingling Yao, Danru Xu, Sébastien Lachapelle, Sara Magliacane, Perouz Taslakian, Georg Martius, Julius von Kügelgen, and Francesco Locatello. Multi-view causal representation learning with partial observability. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=OGtnhKQJms.

Appendix Contents

The detailed proofs and constructions underlying Theorems 1, 2, and 3 appear in Appendices H, I, and E–G, respectively.

Appendix	Purpose	Page
A. Probability and Exact-Sampling Tools	Conditional concentration, finite-family testing, adaptive failure accounting, and exact finite rejection-sampling primitives used throughout the proofs.	14
B. Persistent BSG Compiler	Full deterministic and randomized construction of the persistent BSG core, including searchable sifting, validation, exact sampling, and finite resource bounds.	16
C. Size-Oblivious Algorithmic PFR	Algorithmic PFR for an implicit structured set of unknown cardinality: localization, dense modeling, restricted-homomorphism search, validation, covering, and resource accounting.	30
D. Structural Lifting from an Observed Core	Lifting a core cover to a latent-support cover when their overlap is sufficiently large, with subspace-size control.	36
E. Intermediate List Construction	Arbitrary-corruption learner based on persistent-core peeling, PFR on each core, residual certificates, finite residual access, and terminal invariants.	38
F. From a List to One Common Hypothesis	Fresh-sample selection and majority-coverage arguments that convert the intermediate candidate list into one hypothesis valid for all compatible latent supports.	43
G. Common Learning under Moderate Corruption	Single-core analysis in the moderate-corruption regime and the proof that one persistent core contains enough clean mass for simultaneous structural lifting.	45
H. Sample-Only Common Learning	Sample-only learning by majority coverage and minimum-dimension selection, with uniform-convergence sample bounds.	47
I. Conditional Sample-Only Hardness	Search-LPN reduction for constant-cost coarse recovery, allowing polynomial-size output lists and known latent cardinality.	48
J. Exact Affine-Support Recovery at $K = 1$	Exact recovery below $\eta = 1/3$, exponential access cost at the boundary with unknown latent cardinality, and refinements using candidate lists or known cardinality.	50
K. Ambiguity Beyond Latent Cardinality	Sharp common-cover costs for a prescribed family of equally large latent subspaces with independent unseen directions.	56
L. One-Core Size-Only Mechanism Barrier	A limitation of the single-core size-based certification mechanism, explaining the intrinsic $K^{-1/2}$ scale of that proof architecture.	57
M. Detailed Relation to Prior Work	Extended comparisons with latent identifiability, robust and list-decodable learning, persistent-noise query models, algorithmic additive combinatorics, and exact affine recovery.	59
N. Algorithmic and Access-Model Questions	Open algorithmic questions on matching statistical covering quality, representing recursive residuals more efficiently, and weakening the required sample/membership access.	61

A Probability and Exact-Sampling Tools

The randomized constructions reveal information sequentially. At each stage, the object being tested may depend on the full past transcript, but the block used to test it is fresh. We therefore state all probability tools conditionally. Throughout, $\mathcal{F}$ denotes the sigma-field generated by the past; quantities described as fixed may be $\mathcal{F}$-measurable.

Arithmetic convention for public parameters. Public scalar parameters such as α, ρ, and the deterministic thresholds derived from them are available to the algorithm as exact public parameters, and arithmetic/comparisons involving only these public scalars are counted at unit cost. The displayed ordinary running-time bounds are stated under this exact-public-parameter arithmetic convention; we do not claim a separate bit-complexity bound for arbitrary real encodings of the public parameters. Data-dependent quantities are stored discretely: empirical probabilities are represented by their integer counts and sample sizes, and derived thresholds are reconstructed symbolically from those counts and the public parameters. In particular, the record-length bounds below count only data-dependent information stored in the record, not a second encoding of the public parameters.

Lemma 7 (Conditional concentration and finite-family control)**.** *Conditional on $\mathcal{F}$, let $X_1, \dots, X_m$ be independent Bernoulli variables with common $\mathcal{F}$-measurable mean p, and put $X = \sum_i X_i$. Then, almost surely,*

$$\Pr[X \le (1-\delta)mp \mid \mathcal{F}] \le \exp\left(-\frac{\delta^2 mp}{2}\right) \qquad (0 < \delta < 1), \tag{18}$$

$$\Pr[X \ge (1+\delta)mp \mid \mathcal{F}] \le \exp\left(-\frac{\delta^2 mp}{2+\delta}\right) \qquad (\delta > 0), \tag{19}$$

$$\Pr\left[\left|\frac{X}{m} - p\right| \ge \varepsilon \,\middle|\, \mathcal{F}\right] \le 2e^{-2m\varepsilon^2}. \tag{20}$$

Moreover, if $\mathcal{H}$ is an $\mathcal{F}$-measurable family of at most M functions $f : \Omega \to [0,1]$, and $Z_1, \dots, Z_m$ are fresh independent samples from a fixed conditional distribution μ, then

$$\Pr\left[\sup_{f \in \mathcal{H}} \left|\frac{1}{m}\sum_{i=1}^{m} f(Z_i) - \mathbb{E}_\mu f\right| > \varepsilon \,\middle|\, \mathcal{F}\right] \le 2Me^{-2m\varepsilon^2}. \tag{21}$$

Hence simultaneous error at most ε holds with conditional probability at least $1-\rho$ whenever

$$m \ge \frac{1}{2\varepsilon^2} \log \frac{2M}{\rho}. \tag{22}$$

Proof. Condition on $\mathcal{F}$. The ordinary Chernoff and Hoeffding bounds apply to the fresh block. The finite-family statement follows by applying Hoeffding to each $f \in \mathcal{H}$ and taking a union bound. Once the sample block is stored, every empirical value is a deterministic function of that same block; this is the basic persistence mechanism used below. □

Lemma 8 (Sequential conditional failure accounting)**.** *Let $\mathcal{F}_0 \subseteq \cdots \subseteq \mathcal{F}_R$ be a filtration. For each j, let $\mathcal{B}_j \in \mathcal{F}_j$ satisfy*

$$\Pr[\mathcal{B}_j \mid \mathcal{F}_{j-1}] \le \rho_j \quad \textit{almost surely}.$$

Then

$$\Pr\left[\bigcup_{j=1}^{R} \mathcal{B}_j\right] \le \sum_{j=1}^{R} \rho_j. \tag{23}$$

In particular, an algorithm with at most R adaptively reached test slots, each using a fresh block and having conditional error at most ρ, incurs total error at most $R\rho$.

Proof. Let $\mathcal{H}_j = \mathcal{B}_j \cap \bigcap_{i<j} \mathcal{B}_i^c$ be the event that the first failure occurs at stage j. The $\mathcal{H}_j$ are disjoint and the tower property gives

$$\Pr[\mathcal{H}_j] = \mathbb{E}\big[\mathbf{1}_{\cap_{i<j}\mathcal{B}_i^c} \Pr(\mathcal{B}_j \mid \mathcal{F}_{j-1})\big] \le \rho_j.$$

Summing proves the claim. Unreached slots may be padded with the empty bad event. □

Lemma 9 (Finite gap tests)**.** *The following two conditional tests are valid after any adaptively generated past.*

(a) Multiplicative gap. *Fix $0 < a < \theta < b$, let $q \in (0, 1]$ with $bq \le 1$, and accept a Bernoulli mean p when the empirical mean from m fresh samples is at least θq. Then*

$$p \ge bq \Longrightarrow \Pr[\text{reject} \mid \mathcal{F}] \le \exp\left(-\frac{(b-\theta)^2}{2b} mq\right), \tag{24}$$

$$p \le aq \Longrightarrow \Pr[\text{accept} \mid \mathcal{F}] \le \exp\left(-\frac{(\theta-a)^2}{a+\theta} mq\right). \tag{25}$$

Thus both errors are at most ρ if

$$m \ge C(a, \theta, b) q^{-1} \log \frac{2}{\rho}, \quad C(a, \theta, b) := \max\left\{\frac{2b}{(b-\theta)^2}, \frac{a+\theta}{(\theta-a)^2}\right\}. \tag{26}$$

(b) Additive gap. *Fix $0 \le a < \theta < b \le 1$, and accept when the empirical mean is at least θ. Then*

$$p \ge b \Longrightarrow \Pr[\text{reject} \mid \mathcal{F}] \le e^{-2m(b-\theta)^2}, \tag{27}$$

$$p \le a \Longrightarrow \Pr[\text{accept} \mid \mathcal{F}] \le e^{-2m(\theta-a)^2}. \tag{28}$$

It suffices that

$$m \ge \frac{1}{2\min\{(b-\theta)^2, (\theta-a)^2\}} \log \frac{2}{\rho}. \tag{29}$$

For the persistent BSG validation below it is enough to take

$$m_{\deg,\mathrm{H}} = \left\lceil 128 h^{-1} \log \frac{2}{\zeta} \right\rceil, \qquad m_{\deg,\mathrm{L}}(\delta_0) = \left\lceil 256 \delta_0^{-1} \log \frac{2}{\zeta} \right\rceil, \tag{30}$$

$$m_{\mathrm{pair}} = \left\lceil 512 \log \frac{2}{\zeta} \right\rceil. \tag{31}$$

Proof. Part (a) is the lower- and upper-tail Chernoff bound applied at the two gap endpoints; part (b) is Hoeffding. For the displayed BSG constants use $(a, \theta, b) = (3/4, 7/8, 1)$ in the high degree test, $(3/8, 7/16, 1/2)$ in the low degree test, and apply the additive test to $1 - \mathsf{Bad}$ with gaps of size $1/32$. □

Lemma 10 (Fixed-block exact rejection sampling)**.** *Conditional on $\mathcal{F}$, suppose a trial takes values in $\Omega \cup \{\perp\}$ and satisfies*

$$\Pr[V \in A \mid \mathcal{F}] = p\,\mu(A)$$

for an $\mathcal{F}$-measurable acceptance probability p and conditional law μ. In a block of T independent trials, return the first non-$\perp$ value, or FAIL *if all trials reject. Then*

$$\Pr[\mathsf{FAIL} \mid \mathcal{F}] = (1-p)^T \le e^{-pT}, \tag{32}$$

and conditioned on success the returned value has law μ. Independent blocks have the product of their conditional output laws when conditioned on their individual successes.

In particular, if an adaptively fixed $Y \subseteq S$ satisfies $|Y| \ge \beta|S|$, then t independent rejection blocks from U_S of length

$$T = \left\lceil \beta^{-1} \log \frac{t}{\rho} \right\rceil \tag{33}$$

all succeed with conditional probability at least $1 - \rho$, and, conditioned on success, their outputs have law

$$U_Y^{\otimes t}. \tag{34}$$

Proof. The probability that the first accepted value lies in A is

$$p\mu(A) \sum_{j=1}^{T} (1-p)^{j-1};$$

the geometric factor is independent of A, so conditioning on success leaves the law μ. For set rejection, one trial accepts with probability $|Y|/|S| \ge \beta$; the chosen T makes each block fail with probability at most ρ/t. A union bound and conditional independence give the product statement. □

B PERSISTENT BSG COMPILER: FULL CONSTRUCTION AND PROOF

The structural argument follows the high-/low-multiplicity strategy of Reiher & Schoen (2024). We formulate weighted certificates that can be sampled through the oracle and introduce margins for the sampled tests. We then construct a finite record defining one fixed subset and provide exact sampling from that set, proving Theorem 5.

B.1 DETERMINISTIC SEARCHABLE SIFTING

Let $S \subseteq \mathbb{F}_2^n$ be nonempty, write $N := |S|$, and define

$$r(d) := |S \cap (S+d)|, \qquad p(d) := \frac{r(d)}{N}.$$

Then

$$\sum_d p(d) = N, \qquad \sum_d p(d)^2 = \frac{E(S)}{N^2}. \tag{35}$$

Fix $\alpha \in (0,1]$ and put $h := \sqrt{\alpha}$. For every nonempty $Y \subseteq S$ and $s \in Y+Y$, fix the lexicographically first ordered representation $s = y_1(s) + y_2(s)$.

Lemma 11 (Two structural criteria)**.** *Each of the following deterministic conditions forces a large small-doubling subset.*

(H) Let $Q \subseteq \mathbb{F}_2^n$ and $Y \subseteq X \subseteq S$ satisfy

$$|Q| \le 2^{13} h^{-2} N, \qquad |X| \ge \frac{3}{4} hN, \qquad |Y| \ge \frac{1}{2}|X|, \tag{36}$$

and suppose every $y \in Y$ has at most $3|X|/8$ partners $w \in X$ with $y + w \notin Q$. Then

$$|Y| \ge \frac{3}{8} hN, \qquad |Y+Y| \le \frac{2^{33}}{9} h^{-6} |Y|.$$

(L) Let $Q \subseteq \mathbb{F}_2^n$, and suppose a symmetric graph $\mathcal{G} \subseteq S^2$, numbers $\delta_0, \tau, \kappa > 0$, and sets $Y \subseteq X \subseteq S$ satisfy

$$\delta_0 \ge \frac{7}{8} h, \qquad \delta_0 \tau^2 \ge 2^{-16} h^2, \qquad |Q| \le \frac{2\delta_0 N}{h\tau}, \tag{37}$$

$$|X| \ge \frac{3\delta_0 N}{8}, \qquad |Y| \ge \frac{1}{2}|X|, \qquad \kappa \ge 2^{-12}\delta_0^2. \tag{38}$$

Assume every graph edge has Q-codegree at least κN, and every $y \in Y$ has at most $3|X|/8$ non-neighbours in X. Then

$$|Y| \ge \frac{21}{128} hN, \qquad |Y+Y| \le \frac{2^{69}}{9} h^{-8} |Y|.$$

Proof. For (H), fix $s \in Y+Y$ with canonical endpoints y_1, y_2. At least $|X|/4$ points $w \in X$ satisfy both $y_1 + w \in Q$ and $y_2 + w \in Q$. The map $(s, w) \mapsto (y_1 + w, w + y_2) \in Q^2$ is injective, hence

$$|Y+Y| \frac{|X|}{4} \le |Q|^2.$$

Substitution gives the stated doubling bound and the size lower bound is immediate.

For (L), the same endpoint argument leaves at least $|X|/4$ common graph neighbours w. For each such w, each endpoint–w graph edge has at least κN witnesses in S. The resulting map $(s, w, z_1, z_2) \mapsto (q_1, q_2, q_3, q_4) \in Q^4$ is injective, so

$$|Y+Y| \frac{|X|}{4} (\kappa N)^2 \le |Q|^4. \tag{39}$$

Using equation 37–equation 38 yields the stated h^{-8} bound; the size estimate follows directly. □

For $Q \subseteq \mathbb{F}_2^n$, let

$$R_Q := \{(x,y) \in S^2 : x+y \in Q\}, \qquad \delta(Q) := \frac{|R_Q|}{N^2},$$

and define the Q-codegree

$$c_Q(x,y) := |\{z \in S : x+z \in Q,\ y+z \in Q\}|.$$

Set

$$\mathcal{H}_h := \{d : p(d) \geq h\}, \qquad \mathcal{L}_h := \{d : p(d) < h\}.$$

If $E(S) \geq h^2 N^3$, then

$$\sum_{d \in \mathcal{H}_h} p(d)^2 \geq \frac{h^2 N}{32} \quad \text{or} \quad \sum_{d \in \mathcal{L}_h} p(d)^2 \geq \frac{31 h^2 N}{32}. \tag{40}$$

Lemma 12 (Visible low-multiplicity layer)**.** *Assume the second alternative in equation 40. Define*

$$\mathcal{T}_h := \left\{ 2^{-j} : 0 \leq j \leq \left\lceil \log_2 \frac{2^8}{h} \right\rceil \right\}.$$

Then some $\tau \in \mathcal{T}_h$ has the following property. For

$$Q_\tau := \{d \in \mathcal{L}_h : p(d) \geq h\tau\}, \qquad \delta_\tau := \delta(Q_\tau),$$

one has

$$\delta_\tau \geq \frac{15}{16} h, \qquad \delta_\tau \tau^2 \geq \frac{h^2}{2^{15}}, \qquad h\tau \geq \frac{h^2}{2^8}, \qquad |Q_\tau| \leq \frac{\delta_\tau N}{h\tau}. \tag{41}$$

Every $d \in Q_\tau$ also satisfies $h\tau \leq p(d) < h$.

Proof. For $d \in \mathcal{L}_h$, put $x_d = p(d)/h$, and for $t \in (0,1]$ write $M(t) = \sum_{x_d \geq t} x_d$. Then

$$\sum_d x_d^2 \geq \frac{31}{32} N, \qquad \sum_d x_d \leq \frac{N}{h}, \qquad \int_0^1 M(t)\, dt = \sum_d x_d^2.$$

If no t simultaneously satisfied $M(t) \geq 15N/16$ and $M(t)t^2 \geq hN/2^{13}$, then

$$M(t) \leq \frac{15}{16} N + \min\left\{ \frac{N}{h}, \frac{hN/2^{13}}{t^2} \right\}.$$

Integrating and using $\int_0^1 \min\{A, Z/t^2\}\, dt \leq 2\sqrt{AZ}$ contradicts the $31N/32$ second-moment lower bound. Choose dyadic τ with $t/2 < \tau \leq t$. The quadratic bound loses at most a factor four. At this chosen t, the inequality $M(t)t^2 \geq hN/2^{13}$, together with $M(t) \leq N/h$, gives

$$t^2 \geq \frac{h^2}{2^{13}}, \qquad \tau > \frac{t}{2} \geq \frac{h}{2^{15/2}} > \frac{h}{2^8}.$$

Thus the finite dyadic grid reaches the needed scale and $h\tau > h^2/2^8$. Since

$$\delta_\tau = \frac{h}{N} M(\tau),$$

the first two displayed bounds follow, while $|Q_\tau| h\tau \leq \delta_\tau N$ gives the cardinality estimate. □

Lemma 13 (Searchable weighted certificates)**.** *Under the corresponding alternative in equation 40, the structural criteria above have noticeable mass under an exactly sampleable weighted law.*

(H) Set

$$\vartheta_{\mathrm{H}} := \frac{h^2}{2^{12}}, \qquad Q_{\mathrm{H}} := \{q : p(q) \geq \vartheta_{\mathrm{H}}\}, \qquad \mu_E(d) := \frac{r(d)^2}{E(S)}.$$

Then

$$|Q_{\mathrm{H}}| \leq 2^{12} h^{-2} N, \qquad \sum_{q \notin Q_{\mathrm{H}}} r(q)^2 \leq \vartheta_{\mathrm{H}} N^3. \tag{42}$$

If the first alternative in equation 40 holds, then μ_E assigns probability at least $7h^2/2^8$ to differences d for which $|X_d| \geq hN$ and the average fraction of pairs $(x,y) \in X_d^2$ with $x+y \notin Q_{\mathrm{H}}$ is at most $1/16$. Cleaning vertices of weak degree above $|X_d|/4$ leaves at least $3|X_d|/4$ points and satisfies criterion (H).

(L) Let $\tau, Q_\tau, \delta_\tau$ be supplied by Lemma 12, put $R_\tau = R_{Q_\tau}$, and let

$$\Gamma_{R_\tau}(a) := \{x \in S : (a,x) \in R_\tau\}, \quad d(a) := |\Gamma_{R_\tau}(a)|, \quad Z := \sum_a d(a)^2.$$

For

$$\mu_{R_\tau}(a) := \frac{d(a)^2}{Z}, \qquad \kappa_\tau := \frac{\delta_\tau^2}{2^{10}},$$

and any fixed symmetric graph $\mathcal{G}_\tau$ satisfying

$$c_{Q_\tau}(x,y) \geq 2\kappa_\tau N \Longrightarrow (x,y) \in \mathcal{G}_\tau \Longrightarrow c_{Q_\tau}(x,y) \geq \kappa_\tau N, \tag{43}$$

the law μ_{R_τ} assigns probability at least $15/32$ to centres a with

$$d(a) \geq \frac{\delta_\tau N}{2}$$

and at most a $1/16$ fraction of graph nonedges inside $\Gamma_{R_\tau}(a)^2$. Cleaning high-nonedge-degree vertices leaves at least three quarters of the neighbourhood and satisfies criterion (L).

Proof. For (H), define

$$B_Q(d) := |\{(x,y) \in X_d^2 : x + y \notin Q\}|.$$

Double counting ordered pairs of representations gives the identity

$$\sum_d B_Q(d) = \sum_{q \notin Q} r(q)^2.$$

The high-energy alternative places at least $h^2N^3/32$ squared- representation mass on $\{d : p(d) \geq h\}$. Differences with $B_{Q_{\mathrm{H}}}(d) > r(d)^2/16$ carry less than $16\sum_d B_{Q_{\mathrm{H}}}(d) \leq h^2N^3/2^8$. Thus at least $7h^2N^3/2^8$ good mass remains; because $E(S) \leq N^3$, its μ_E-probability is at least $7h^2/2^8$. Markov's inequality on weak degrees gives the cleaning claim.

For (L), Cauchy–Schwarz gives $Z \geq \delta_\tau^2 N^3$. Centres with degree below $\delta_\tau N/2$ carry at most half of Z. If $B_{\mathcal{G}}(a)$ denotes the ordered number of graph nonedges inside $\Gamma(a)^2$, then

$$\sum_a B_{\mathcal{G}}(a) = \sum_{(x,y) \notin \mathcal{G}} c_{Q_\tau}(x,y) < 2\kappa_\tau N^3.$$

Hence centres with $B_{\mathcal{G}}(a) > d(a)^2/16$ carry less than $32\kappa_\tau N^3 = \delta_\tau^2 N^3/32 \leq Z/32$. Removing the two bad classes leaves μ_{R_τ}-mass at least $1 - 1/2 - 1/32 = 15/32$, and Markov gives the cleaning claim. □

Theorem 14 (Weighted deterministic sifting)**.** *Let $S \subseteq \mathbb{F}_2^n$ be nonempty and suppose $E(S) \geq \alpha|S|^3$. Then at least one of the following holds.*

(H) Under μ_E, a set of differences of probability at least $7\alpha/2^8$ yields, after cleaning, a set $Y \subseteq S$ with

$$|Y| \geq \frac{3}{8}\sqrt{\alpha}\,|S|, \qquad |Y+Y| \leq \frac{2^{33}}{9}\alpha^{-3}|Y|.$$

(L) Some $\tau \in \mathcal{T}_h$ satisfies equation 41; for every graph obeying equation 43, the squared-degree law assigns mass at least $15/32$ to centres whose cleaning yields $Y \subseteq S$ with

$$|Y| \geq \frac{21}{128}\sqrt{\alpha}\,|S|, \qquad |Y+Y| \leq \frac{2^{69}}{9}\alpha^{-4}|Y|.$$

In particular one branch contains a set Y with

$$|Y| \geq \frac{21}{128}\sqrt{\alpha}\,|S|, \qquad |Y+Y| \leq \frac{2^{69}}{9}\alpha^{-4}|Y|.$$

Proof. Apply the dichotomy equation 40. In the high case use Lemma 13(H) and Lemma 11(H). In the low case use Lemmas 12, 13(L), and 11(L). Since $\alpha \leq 1$, the high- branch α^{-3} bound is also dominated by the displayed α^{-4} guarantee. □

The weighted bounds make the structural alternatives searchable: each places useful candidates under a distribution that the oracle algorithm can sample exactly.

B.2 Persistent randomized BSG: construction and proof

Let $S \subseteq \mathbb{F}_2^n$ be nonempty, write $N := |S|$, and fix $\alpha \in (0, 1]$. As in the preceding section, put $h := \sqrt{\alpha}$.

The randomized algorithm compiles the criteria of Section B.1 into one reusable predicate. All data defining the output set are sampled before it is returned. Fresh validation tests decide whether to accept a candidate; their randomness is separate from the data defining membership.

The algorithm estimates the representation function once, constructs one high-popularity set and finitely many low-multiplicity layers, proposes candidates from the weighted distributions of Section B.1, and tests them using fresh samples. An independent cleaning sample defines the final subset.

Set

$$c_{\mathrm{BSG}} := \frac{21}{128}, \qquad C_{\mathrm{BSG}} := \frac{2^{69}}{9}, \qquad C_{\mathrm{BSG,H}} := \frac{2^{33}}{9}.$$

Theorem 15 (Persistent randomized BSG, detailed form)**.** *There is a randomized algorithm with the following properties. It receives deterministic membership access to a nonempty set $S \subseteq \mathbb{F}_2^n$, independent exact uniform samples from S, and parameters $\alpha, \rho_{\mathrm{snd}}, \rho_{\mathrm{cmp}} \in (0, 1)$. It always terminates and returns either* FAIL *or a finite record D.*

Every nonfailure record D defines one deterministic set $Y_D \subseteq S$. For an arbitrary input set S, without an energy assumption,

$$\Pr\big[D \neq \mathsf{FAIL} \text{ and } \big(|Y_D| < c_{\mathrm{BSG}}\sqrt{\alpha}\,|S| \text{ or } |Y_D + Y_D| > C_{\mathrm{BSG}}\alpha^{-4}|Y_D|\big)\big] \le \rho_{\mathrm{snd}}. \tag{44}$$

If the output has a high-branch tag, then on the same soundness event it satisfies the stronger estimate

$$|Y_D + Y_D| \le C_{\mathrm{BSG,H}}\alpha^{-3}|Y_D|.$$

If, in addition,

$$E(S) \ge \alpha|S|^3, \tag{45}$$

then

$$\Pr[D = \mathsf{FAIL}] \le \rho_{\mathrm{cmp}}. \tag{46}$$

Once D is fixed, membership in Y_D is deterministic, and every later query refers to this same set. If

$$\Lambda := 1 + n + \log\frac{1}{\alpha} + \log\frac{1}{\rho_{\mathrm{snd}}} + \log\frac{1}{\rho_{\mathrm{cmp}}},$$

then the compilation uses

$$O(\alpha^{-2}\Lambda^4) \quad \text{samples from } S, \text{ and} \quad O(\alpha^{-4}\Lambda^5) \quad \text{membership queries to } S.$$

One membership query to Y_D uses $O(\alpha^{-4}\Lambda^3)$ membership queries to S, and the record length is $O(n\alpha^{-2}\Lambda)$ bits. All these bounds are worst-case over the random choices of the algorithm.

B.2.1 The construction

The exact block lengths and error allocations are given in Appendix B.3. The following description records the logical order of the algorithm.

All proposal loops and all rejection samplers have deterministic finite lengths. If a rejection block is exhausted, the current candidate is discarded. Exhaustion can cause the algorithm to miss a good candidate, but it cannot cause an invalid set to be returned.

B.2.2 Approximating the relevant relations

For $q \in \mathbb{F}_2^n$, recall that

$$p(q) = \Pr_{Z \sim U_S}[Z + q \in S].$$

Draw independent $Z_1, \dots, Z_{m_{\mathrm{rep}}} \sim U_S$ and define

$$\widehat{p}(q) := \frac{1}{m_{\mathrm{rep}}} \sum_{i=1}^{m_{\mathrm{rep}}} \mathbf{1}_S(Z_i + q).$$

Algorithm 1 PERSISTENTBSG

Require: membership access and uniform samples from S; $\alpha, \rho_{\mathrm{snd}}, \rho_{\mathrm{cmp}} \in (0,1)$

draw one representation sample block and define the high set $\widehat{Q}_{\mathrm{H}}$ and every dyadic layer $\widehat{Q}_\tau$
estimate the relation density of every dyadic layer and retain the layers passing the density test
for each retained layer, draw one codegree sample and define a fixed graph $\mathcal{G}_\tau$
generate prescribed finite lists of high and low candidates from the exact weighted proposal laws
for the candidates in a fixed deterministic order **do**
    test the density of the associated set X
    using fresh exact samples from X, test its average bad-pair rate
    if both tests pass **then**
        draw an independent cleaning sample from X
        if all required finite rejection blocks succeed **then**
            return the finite record defining the cleaned set
        end if
    end if
end for
return FAIL

The stored sample makes $q \mapsto \widehat{p}(q)$ a fixed deterministic function on the whole group.

Let

$$\vartheta_{\mathrm{H}} := \frac{h^2}{2^{12}}, \qquad \widehat{Q}_{\mathrm{H}} := \{q : \widehat{p}(q) \geq 3\vartheta_{\mathrm{H}}/4\}.$$

For the dyadic family

$$\mathcal{T}_h := \left\{ 2^{-j} : 0 \leq j \leq \left\lceil \log_2 \frac{2^8}{h} \right\rceil \right\},$$

put

$$\widehat{Q}_\tau := \{q : \widehat{p}(q) \geq 3h\tau/4\} \qquad (\tau \in \mathcal{T}_h).$$

For a fixed layer $\widehat{Q}_\tau$, write

$$\delta_\tau^{\mathrm{true}} := \Pr_{A,X \sim U_S}[A + X \in \widehat{Q}_\tau].$$

A fresh sample gives an estimate $\widetilde{\delta}_\tau$, and we use the lower confidence value

$$\delta_{\tau,0} := \max\{0, \widetilde{\delta}_\tau - h/2^7\}.$$

We retain the layer when

$$\delta_{\tau,0} \geq \frac{7}{8}h \quad \text{and} \quad \delta_{\tau,0}\tau^2 \geq \frac{h^2}{2^{16}}. \tag{47}$$

Lemma 16 (Representation and layer approximation)**.** *With the block sizes of Appendix B.3.3, the following event has probability at least* $1 - \rho_{\mathrm{rep}} - \rho_{\mathrm{den}}$.

First,

$$|\widehat{p}(q) - p(q)| \leq \frac{h^2}{2^{16}} \qquad \textit{for every } q \in \mathbb{F}_2^n.$$

On this event,

$$\{q : p(q) \geq \vartheta_{\mathrm{H}}\} \subseteq \widehat{Q}_{\mathrm{H}} \subseteq \left\{q : p(q) \geq \frac{2}{3}\vartheta_{\mathrm{H}}\right\}, \qquad |\widehat{Q}_{\mathrm{H}}| \leq 2^{13}h^{-2}N, \tag{48}$$

$$\{q : p(q) \geq h\tau\} \subseteq \widehat{Q}_\tau \subseteq \left\{q : p(q) \geq \frac{2}{3}h\tau\right\} \qquad (\tau \in \mathcal{T}_h). \tag{49}$$

Moreover,

$$\sum_{q \notin \widehat{Q}_{\mathrm{H}}} r(q)^2 \leq \vartheta_{\mathrm{H}} N^3. \tag{50}$$

Second, all relation-density estimates are simultaneously accurate:

$$|\widetilde{\delta}_\tau - \delta_\tau^{\mathrm{true}}| \le \frac{h}{2^7} \qquad (\tau \in \mathcal{T}_h).$$

Every retained layer satisfies

$$\delta_{\tau,0} \le \delta_\tau^{\mathrm{true}} \le \frac{57}{56}\delta_{\tau,0}, \qquad |\widehat{Q}_\tau| \le \frac{2\delta_{\tau,0}N}{h\tau}. \tag{51}$$

Finally, under the low-energy alternative of Theorem 14, the threshold supplied by Lemma 12 is retained by the density test.

Proof. The uniform estimate for $\widehat{p}$ follows from Lemma 7, applied to the 2^n translation predicates. Since the error is $\vartheta_{\mathrm{H}}/16$, the high inclusions follow immediately. The cardinality estimate follows from

$$\frac{2}{3}\vartheta_{\mathrm{H}}|\widehat{Q}_{\mathrm{H}}| \le \sum_{q\in\widehat{Q}_{\mathrm{H}}} p(q) \le N.$$

Outside $\widehat{Q}_{\mathrm{H}}$ one has $p(q) < \vartheta_{\mathrm{H}}$, which gives equation 50 after summing $p(q)^2 \le \vartheta_{\mathrm{H}}p(q)$.

For an arbitrary searched low threshold, the ceiling in the dyadic grid only guarantees

$$h\tau > \frac{h^2}{2^9}.$$

Hence the same representation error is strictly smaller than $2^{-7}h\tau$, which is still well inside the gap between the thresholds $h\tau$, $3h\tau/4$, and $2h\tau/3$. This proves equation 49. The density estimates follow from the same lemma, now applied to the J_h searched layers.

If a layer is retained, then $\delta_{\tau,0} \le \delta_\tau^{\mathrm{true}} \le \delta_{\tau,0} + h/64$. Since $\delta_{\tau,0} \ge 7h/8$, this is at most $(57/56)\delta_{\tau,0}$. Every element of $\widehat{Q}_\tau$ has $p(q) \ge (2/3)h\tau$, and hence

$$|\widehat{Q}_\tau| \le \frac{3\delta_\tau^{\mathrm{true}}N}{2h\tau} < \frac{2\delta_{\tau,0}N}{h\tau}.$$

For the visible threshold of Lemma 12, the ideal layer is contained in $\widehat{Q}_\tau$, so

$$\delta_\tau^{\mathrm{true}} \ge \frac{15}{16}h, \qquad \delta_\tau^{\mathrm{true}}\tau^2 \ge \frac{h^2}{2^{15}}.$$

Subtracting the two-sided density error leaves $\delta_{\tau,0} > 7h/8$. The relative loss is at most $1/60$, so the quadratic condition in equation 47 also holds. □

Fix a retained layer and abbreviate

$$Q := \widehat{Q}_\tau, \qquad \delta_0 := \delta_{\tau,0}, \qquad \kappa_0 := \frac{\delta_0^2}{2^{12}}.$$

Draw independent $W_1, \dots, W_{m_{\mathrm{code}}} \sim U_S$ and set

$$\widehat{\gamma}_Q(x,y) := \frac{1}{m_{\mathrm{code}}} \sum_{i=1}^{m_{\mathrm{code}}} \mathbf{1}_Q(x+W_i)\mathbf{1}_Q(y+W_i).$$

The corresponding fixed graph is

$$\mathcal{G}_\tau := \{(x,y) \in S^2 : \widehat{\gamma}_Q(x,y) \ge 3\kappa_0/2\}.$$

Lemma 17 (Codegree graph)**.** *With the block size of Appendix B.3.3, except with probability ρ_{code}, simultaneously for every $x, y \in S$,*

$$\left|\widehat{\gamma}_Q(x,y) - \frac{c_Q(x,y)}{N}\right| \le \frac{\kappa_0}{4}.$$

On this event,

$$c_Q(x,y) \ge 2\kappa_0 N \implies (x,y) \in \mathcal{G}_\tau \implies c_Q(x,y) \ge \kappa_0 N. \tag{52}$$

Proof. Condition on the previously sampled data, which fixes Q, δ_0, κ_0. The codegree sample is fresh, and each summand is Bernoulli with mean $c_Q(x,y)/N$. Lemma 7, applied to at most 2^{2n} ordered pairs, gives the uniform estimate. The two implications follow by comparing the thresholds $2\kappa_0$, $3\kappa_0/2$, and the error $\kappa_0/4$. □

The weighted laws in Lemma 13 admit simple exact sampling rules.

Proposition 18 (Exact weighted proposals)**.** *The following one-trial procedures have the stated laws.*

1. *Draw $U, V, W \sim U_S$, put $D := U + V$, and accept D when $W + D \in S$. The trial accepts with probability $E(S)/N^3$, and, conditioned on acceptance,*

$$\Pr[D = d] = \frac{r(d)^2}{E(S)}.$$

2. *Fix $Q \subseteq \mathbb{F}_2^n$. Draw $A, X, Y \sim U_S$ and accept A when $A + X, A + Y \in Q$. If $d_Q(a) := |\{x \in S : a + x \in Q\}|$ and $Z_Q := \sum_a d_Q(a)^2$, then the trial accepts with probability Z_Q/N^3, and, conditioned on acceptance,*

$$\Pr[A = a] = \frac{d_Q(a)^2}{Z_Q}.$$

If the relation density of Q is δ, then $Z_Q/N^3 \geq \delta^2$.

Proof. For the first sampler,

$$\Pr[D = d \text{ and accept}] = \frac{r(d)}{N^2}\frac{r(d)}{N} = \frac{r(d)^2}{N^3}.$$

Summing over d gives the acceptance probability, and normalizing gives the conditional law.

For the second sampler, the probability of drawing $A = a$ and accepting is $d_Q(a)^2/N^3$. The final lower bound follows from Cauchy–Schwarz:

$$Z_Q \geq \frac{1}{N}\left(\sum_a d_Q(a)\right)^2 = \delta^2 N^3.$$

□

Each candidate request uses a fixed number of independent trials and returns the first accepted value. By Lemma 10, conditioned on success the output has exactly the law in Proposition 18. The precise block lengths are recorded in Appendix B.3.4.

B.2.3 Proposing, testing, and cleaning candidates

The two branches use the same validation pattern. Their candidate sets and bad-pair predicates are:

branch	candidate set X	$\mathsf{Bad}(x,w)$
high, candidate d	$X_d := S \cap (S + d)$	$\mathbf{1}[x + w \notin \widehat{Q}_{\mathrm{H}}]$
low, candidate a in layer τ	$X_a := \{x \in S : a + x \in \widehat{Q}_\tau\}$	$\mathbf{1}[(x,w) \notin \mathcal{G}_\tau]$.

We first test the density of X inside S. If this test passes, we obtain exact uniform samples from X by fixed-length rejection sampling from S, and estimate the average bad-pair rate

$$\omega_X := \Pr_{x,w \sim U_X}[\mathsf{Bad}(x,w) = 1].$$

The candidate is accepted when the empirical bad-pair rate is at most $3/32$.

Lemma 19 (Candidate validation)**.** *Condition on all randomness revealed before a candidate is fixed. With the fresh validation block sizes of Appendix B.3.5, the following statements hold, each with conditional error at most ζ.*

For a high candidate d, the degree test accepts at threshold $7h/8$. It accepts with high probability when $|X_d| \geq hN$, while, except with probability ζ, every candidate that passes satisfies

$$|X_d| > \frac{3}{4}hN.$$

For a low candidate a associated with a retained layer of certified density δ_0, the degree test accepts at threshold $7\delta_0/16$. It accepts with high probability when $|X_a| \geq \delta_0 N/2$, while, except with probability ζ, every candidate that passes satisfies

$$|X_a| > \frac{3\delta_0 N}{8}.$$

In either branch, conditional on the success of the exact X-sampling blocks, the bad-pair test has the following gap:

$$\omega_X \leq \frac{1}{16} \quad \Longrightarrow \quad \Pr[\textit{reject}] \leq \zeta,$$

whereas, except with probability ζ, every candidate that passes satisfies $\omega_X < 1/8$.

Proof. After conditioning, all validation samples are fresh. Apply Lemma 9 to the two degree tests and Lemma 9 to the complementary good-pair indicators in the bad-pair test. The resulting constants are recorded in Lemma 9 and Appendix B.3.5. □

Validation controls an average bad-pair rate. The final output needs a pointwise guarantee for every point that it accepts.

Lemma 20 (Independent cleaning)**.** *Let $X \subseteq S$ be nonempty, and let $\mathsf{Bad} : X^2 \to \{0,1\}$ be fixed with*

$$\mathbb{E}_{x,w\sim U_X} \mathsf{Bad}(x,w) \leq \frac{1}{8}.$$

Draw fresh independent $W_1, \ldots, W_{m_{\mathrm{clean}}} \sim U_X$ and define

$$Y := \left\{ x \in X : \frac{1}{m_{\mathrm{clean}}} \sum_{i=1}^{m_{\mathrm{clean}}} \mathsf{Bad}(x, W_i) \leq \frac{5}{16} \right\}.$$

With the block size of Appendix B.3.5, except with probability ρ_{clean},

$$|Y| \geq \frac{1}{2}|X|, \qquad \Pr_{W\sim U_X}[\mathsf{Bad}(y,W)=1] \leq \frac{3}{8} \quad (y \in Y). \tag{53}$$

Proof. Lemma 7, applied to the family of bad degree predicates indexed by $x \in X$, gives a simultaneous additive error at most $1/16$. Thus every point of true bad degree at most $1/4$ enters Y, while every point in Y has true bad degree at most $3/8$. By Markov's inequality, at least half the points of X have true bad degree at most $1/4$. □

The output record stores only the data used by the final membership predicate. Validation samples are not stored.

For a high candidate d, the record contains the representation sample, d, and the cleaning sample. It defines

$$Y_D := \left\{ x \in S : x + d \in S, \quad \frac{1}{m_{\mathrm{clean}}} \sum_i \mathbf{1}[x + W_i \notin \widehat{Q}_{\mathrm{H}}] \leq \frac{5}{16} \right\}.$$

For a low candidate a in layer τ, the record contains the representation sample, τ, the empirical density count from which δ_0 is reconstructed using the public parameters, the codegree sample, a, and the cleaning sample. It defines

$$Y_D := \left\{ x \in S : a + x \in \widehat{Q}_\tau, \quad \frac{1}{m_{\mathrm{clean}}} \sum_i \mathbf{1}[(x, W_i) \notin \mathcal{G}_\tau] \leq \frac{5}{16} \right\}.$$

In both cases, once the record is fixed, this is a deterministic membership predicate.

Deterministic assembly of a returned record. Assume that the representation, density, codegree, validation, and cleaning conclusions used along the realized execution are all correct. If Algorithm 1 returns a nonfailure record D, then its fixed set Y_D satisfies one of the two structural criteria in Lemma 11. In particular,

$$|Y_D| \geq c_{\mathrm{BSG}} hN, \qquad |Y_D + Y_D| \leq C_{\mathrm{BSG}} h^{-8} |Y_D|,$$

and a high-branch output satisfies the stronger

$$|Y_D + Y_D| \leq C_{\mathrm{BSG,H}} h^{-6} |Y_D|.$$

Indeed, in the high branch the approximation event gives $|\widehat{Q}_{\mathrm{H}}| \leq 2^{13} h^{-2} N$; validation gives $|X_d| > 3hN/4$ and average bad-pair rate below $1/8$; and independent cleaning gives $|Y_D| \geq |X_d|/2$ with pointwise bad degree at most $3|X_d|/8$. These are exactly criterion (H). In the low branch, the retained-layer certificate gives

$$\delta_0 \geq \frac{7}{8} h, \qquad \delta_0 \tau^2 \geq \frac{h^2}{2^{16}}, \qquad |\widehat{Q}_\tau| \leq \frac{2\delta_0 N}{h\tau},$$

while validation and cleaning give

$$|X_a| > \frac{3\delta_0 N}{8}, \qquad |Y_D| \geq \frac{1}{2} |X_a|.$$

The sampled codegree sandwich ensures that every graph edge has true $\widehat{Q}_\tau$-codegree at least $2^{-12}\delta_0^2 N$, so criterion (L) applies.

B.2.4 Proof of the main theorem

Useful candidates retain noticeable proposal mass after the ideal relations are replaced by the sampled approximations.

Lemma 21 (Good candidates have noticeable proposal mass)**.** *Assume the representation event of Lemma 16.*

If the high-energy alternative of Theorem 14 holds, then the energy-biased distribution assigns probability at least $7\alpha/2^8$ to differences d satisfying

$$|X_d| \geq hN, \qquad \Pr_{x,w \sim U_{X_d}} [x + w \notin \widehat{Q}_{\mathrm{H}}] \leq \frac{1}{16}.$$

If the low-energy alternative holds, let τ be the visible threshold. On the corresponding codegree event, the squared-degree distribution for $\widehat{Q}_\tau$ assigns probability at least $63/128$ to centres a satisfying

$$|X_a| \geq \frac{\delta_\tau^{\mathrm{true}} N}{2} \geq \frac{\delta_{\tau,0} N}{2},$$

and whose average nonedge rate in $\mathcal{G}_\tau$ is at most $1/16$.

Proof. For the high branch, the weak-pair identity from Lemma 13(H) and equation 50 show that high fibres with bad-pair rate above $1/16$ carry squared-representation mass at most $h^2 N^3/2^8$. The high-energy alternative carries at least $h^2 N^3/32$, leaving at least $7h^2 N^3/2^8$ on the stated good fibres. Since $E(S) \leq N^3$, their energy-biased probability is at least $7\alpha/2^8$.

For the low branch, write $\delta := \delta_\tau^{\mathrm{true}}$ and $\delta_0 := \delta_{\tau,0}$. Centres of degree below $\delta N/2$ carry at most half of the squared-degree mass. By the left-hand implication in equation 52, every graph nonedge has codegree below $2\kappa_0 N$. The double-counting argument of Lemma 13(L) shows that centres with more than a $1/16$ fraction of nonedges carry at most $32\kappa_0 N^3$ squared-degree mass. Since $\kappa_0 = \delta_0^2/2^{12}$ and $\delta_0 \leq \delta$, this is at most $1/128$ of the total squared-degree mass. The remaining mass is at least $1 - 1/2 - 1/128 = 63/128$. □

Proof of Theorem 15. The exact parameter choices are stated in Appendix B.3.2. Let

$$J_h := |\mathcal{T}_h|, \qquad R_\star := R_{\mathrm{H}} + J_h R_{\mathrm{L}}, \qquad \rho_\wedge := \min\{\rho_{\mathrm{snd}}, \rho_{\mathrm{cmp}}\}.$$

Soundness. No energy assumption is made. Let $\mathcal{E}_{\mathrm{snd}}$ be the intersection of the representation and density events, all codegree events, all validation-soundness events, and all cleaning events for the finitely many possible slots. Conditional on the complete past, each validation or cleaning sample is fresh. Lemma 8 therefore gives

$$\begin{aligned}\Pr[\mathcal{E}_{\mathrm{snd}}^{c}] &\le \rho_{\mathrm{rep}} + \rho_{\mathrm{den}} + J_h\rho_{\mathrm{code}} + 2R_\star\zeta + R_\star\rho_{\mathrm{clean}} \\ &= \frac{21}{200}\rho_\wedge < \rho_{\mathrm{snd}}. \end{aligned} \tag{54}$$

Proposal or rejection-block exhaustion is not included here: it discards a candidate and cannot produce an output.

On $\mathcal{E}_{\mathrm{snd}}$, every nonfailure output satisfies the hypotheses of the deterministic assembly above. This proves equation 44.

Completeness. Assume equation 45. By Theorem 14, the high- or low-energy alternative holds.

In the high case, each successful proposal request returns an exact energy-biased candidate, and Lemma 21 gives good-candidate probability at least $7\alpha/2^8$. The choice

$$R_{\mathrm{H}} := \left\lceil 64\alpha^{-1}\log\frac{40}{\rho_{\mathrm{cmp}}} \right\rceil$$

makes the probability of missing every good high candidate at most $\rho_{\mathrm{cmp}}/40$.

In the low case, the visible layer is retained on the global approximation event. Each successful proposal request for that layer has the exact squared-degree law, and Lemma 21 gives good-centre probability at least $63/128$. The choice

$$R_{\mathrm{L}} := \left\lceil 3\log\frac{40J_h}{\rho_{\mathrm{cmp}}} \right\rceil$$

makes the probability of missing every good centre in the visible layer at most $\rho_{\mathrm{cmp}}/(40J_h)$, and hence at most $\rho_{\mathrm{cmp}}/40$.

Every good candidate satisfies the completeness premises of Lemma 19. If its finite rejection blocks succeed, it passes validation and produces a cleaned record except on the allocated validation and cleaning events. The complete failure accounting is

failure source	upper bound
global approximation, validation-soundness, and cleaning events	$21\rho_{\mathrm{cmp}}/200$
no good candidate proposed	$\rho_{\mathrm{cmp}}/40$
proposal-block exhaustion	$\rho_{\mathrm{cmp}}/100$
candidate-sampling exhaustion	$\rho_{\mathrm{cmp}}/100$
cleaning-sampling exhaustion	$\rho_{\mathrm{cmp}}/100$
validation-completeness failures	$\rho_{\mathrm{cmp}}/50$

because

$$2R_\star\zeta = \frac{\rho_\wedge}{50} \le \frac{\rho_{\mathrm{cmp}}}{50}.$$

The sum is at most

$$\frac{9}{50}\rho_{\mathrm{cmp}} < \rho_{\mathrm{cmp}},$$

which proves equation 46.

Complexity. Appendix B.3.7 proves the displayed sample, query, record-length, and per-query bounds. Every loop and rejection block has a deterministic length, so these are worst-case rather than expected bounds. □

B.2.5 Exact sampling from the output set

The record provides deterministic membership in Y_D, but later algorithms also need a prescribed finite block of exact uniform samples from this same set.

Proposition 22 (Exact finite sampling from the output)**.** *Fix a nonfailure record satisfying*

$$|Y_D| \geq c_{\mathrm{BSG}}\sqrt{\alpha}\,|S|.$$

Let $t \geq 1$ be fixed before sampling begins, and let $\rho_{\mathrm{samp}} \in (0,1)$. For each requested sample, use an independent block of

$$T := \left\lceil \frac{1}{c_{\mathrm{BSG}}\sqrt{\alpha}} \log \frac{t}{\rho_{\mathrm{samp}}} \right\rceil$$

uniform samples from S, returning the first point accepted by the membership predicate of Y_D. If any block is exhausted, return FAIL.

The procedure fails with probability at most ρ_{samp}. Conditioned on success, the returned tuple has law

$$U_{Y_D}^{\otimes t}. \tag{55}$$

It uses at most tT samples from S and

$$O\left(t\alpha^{-9/2}\Lambda^3 \log \frac{et}{\rho_{\mathrm{samp}}}\right)$$

membership queries to S.

Proof. Condition on the fixed record D. Each draw from S lands in Y_D with probability at least $c_{\mathrm{BSG}}\sqrt{\alpha}$. Thus one block is exhausted with probability at most

$$(1 - c_{\mathrm{BSG}}\sqrt{\alpha})^T \leq e^{-c_{\mathrm{BSG}}\sqrt{\alpha}T} \leq \frac{\rho_{\mathrm{samp}}}{t}.$$

A union bound proves the failure estimate. In a successful block, the first accepted point is exactly uniform on Y_D. Independence of the blocks and their block-local success events gives equation 55. The query bound follows from the per-query cost in Theorem 15. □

All samples and membership queries in the proposition refer to the same realized set Y_D. Later compositions invoke it on the structural-good event and charge the soundness error of Theorem 15 separately.

B.3 Finite implementation and resource bounds

The following block lengths implement Section B.2 with deterministic loop bounds. We account separately for approximation, proposals, validation, and cleaning.

Throughout, $N := |S|$, $h := \sqrt{\alpha}$, and

$$\Lambda := 1 + n + \log \frac{1}{\alpha} + \log \frac{1}{\rho_{\mathrm{snd}}} + \log \frac{1}{\rho_{\mathrm{cmp}}}.$$

B.3.1 Fixed-length rejection sampling

All rejection procedures use Lemma 10. In particular, when a fixed set $X \subseteq S$ has density at least β, assigning each of m requested samples an independent block of length

$$T_X(\beta, m, \varepsilon) := \left\lceil \beta^{-1} \log \frac{m}{\varepsilon} \right\rceil \tag{56}$$

makes the probability that any request fails at most ε. Conditioned on all blocks succeeding, the returned tuple has the exact product law $U_X^{\otimes m}$, by Lemma 10.

B.3.2 Parameter and error allocation

Let

$$J_h := |\mathcal{T}_h| = 1 + \left\lceil \log_2 \frac{2^8}{h} \right\rceil.$$

We use

$$R_{\mathrm{H}} := \left\lceil 64\alpha^{-1} \log \frac{40}{\rho_{\mathrm{cmp}}} \right\rceil, \qquad R_{\mathrm{L}} := \left\lceil 3 \log \frac{40 J_h}{\rho_{\mathrm{cmp}}} \right\rceil, \qquad R_\star := R_{\mathrm{H}} + J_h R_{\mathrm{L}}. \tag{57}$$

The constants 64 and 3 dominate respectively $2^8/7$ and $128/63$, the reciprocals of the proposal masses in Lemma 21 after removing the factors α and 1.

Put

$$\rho_\wedge := \min\{\rho_{\mathrm{snd}}, \rho_{\mathrm{cmp}}\}.$$

The global and slotwise error parameters are

$$\rho_{\mathrm{rep}} := \frac{\rho_\wedge}{40}, \qquad \rho_{\mathrm{den}} := \frac{\rho_\wedge}{40}, \qquad \rho_{\mathrm{code}} := \frac{\rho_\wedge}{40 J_h}, \tag{58}$$

$$\zeta := \frac{\rho_\wedge}{100 R_\star}, \qquad \rho_{\mathrm{clean}} := \frac{\rho_\wedge}{100 R_\star}, \tag{59}$$

$$\varepsilon_{\mathrm{prop}} := \frac{\rho_{\mathrm{cmp}}}{100 R_\star}, \qquad \varepsilon_{\mathrm{pair}} := \frac{\rho_{\mathrm{cmp}}}{100 R_\star}, \qquad \varepsilon_{\mathrm{clean}} := \frac{\rho_{\mathrm{cmp}}}{100 R_\star}. \tag{60}$$

These choices give the soundness sum

$$\rho_{\mathrm{rep}} + \rho_{\mathrm{den}} + J_h \rho_{\mathrm{code}} + 2R_\star \zeta + R_\star \rho_{\mathrm{clean}} = \frac{21}{200} \rho_\wedge.$$

They also give the validation-completeness contribution $2R_\star \zeta = \rho_\wedge/50$, which is displayed separately in the completeness table of the main proof.

B.3.3 Global representation, density, and codegree blocks

Set

$$\xi_{\mathrm{rep}} := \frac{h^2}{2^{16}}, \qquad \xi_{\mathrm{den}} := \frac{h}{2^7}.$$

It is enough to take

$$m_{\mathrm{rep}} := \left\lceil \frac{1}{2\xi_{\mathrm{rep}}^2} \left(n \log 2 + \log \frac{2}{\rho_{\mathrm{rep}}} \right) \right\rceil, \tag{61}$$

$$m_{\mathrm{den}} := \left\lceil \frac{1}{2\xi_{\mathrm{den}}^2} \log \frac{2 J_h}{\rho_{\mathrm{den}}} \right\rceil. \tag{62}$$

For every retained low layer, with $\kappa_0 := \delta_0^2/2^{12}$, take the layer-specific length

$$m_{\mathrm{code}} := \left\lceil \frac{8}{\kappa_0^2} \left(2n \log 2 + \log \frac{2}{\rho_{\mathrm{code}}} \right) \right\rceil. \tag{63}$$

Hoeffding's inequality gives respectively the uniform errors ξ_{rep}, ξ_{den}, and $\kappa_0/4$. In the resource bounds, m_{code} denotes the maximum of these layer-specific lengths over retained layers; this maximum is deterministically bounded because every retained layer has $\delta_0 \geq 7h/8$.

Every retained layer has $\delta_0 \geq 7h/8$, so $\kappa_0 \geq (49/64)h^2/2^{12}$. Consequently,

$$m_{\mathrm{rep}} = O(\alpha^{-2}\Lambda), \qquad m_{\mathrm{den}} = O(\alpha^{-1}\Lambda), \qquad m_{\mathrm{code}} = O(\alpha^{-2}\Lambda). \tag{64}$$

A membership query to any compiled difference set evaluates $\widehat{p}(q)$ and therefore uses

$$q_Q := m_{\mathrm{rep}} = O(\alpha^{-2}\Lambda) \tag{65}$$

membership queries to S. One adjacency query to a low-branch graph evaluates two difference predicates for every codegree sample, and hence uses

$$q_{\mathcal{G}} := 2 m_{\mathrm{code}} q_Q = O(\alpha^{-4}\Lambda^2) \tag{66}$$

membership queries to S.

B.3.4 CANDIDATE PROPOSAL AND LOCAL SAMPLING BLOCKS

Under the energy promise, a high proposal trial accepts with probability at least α. For every retained low layer, $\delta_\tau^{\mathrm{true}} \geq \delta_0 \geq 7h/8$, so a low proposal trial accepts with probability at least $(7h/8)^2 = 49\alpha/64$. We therefore use the common lower bound

$$\beta_{\mathrm{prop}} := \frac{49}{64}\alpha$$

and assign every proposal request a block of length

$$T_{\mathrm{prop}} := \left\lceil \beta_{\mathrm{prop}}^{-1} \log \frac{1}{\varepsilon_{\mathrm{prop}}} \right\rceil. \tag{67}$$

On an arbitrary input this block still has deterministic length and, when successful, returns the exact conditional proposal law. The lower bound on its success probability is used only in the completeness proof.

A candidate that reaches pair validation or cleaning is sampled from its associated set X through independent rejection blocks. We use the common nominal density

$$\beta_X := \frac{h}{4}.$$

For a bad-pair test requesting $2m_{\mathrm{pair}}$ points, give each request an independent block of length

$$T_{\mathrm{pair}} := \left\lceil \beta_X^{-1} \log \frac{2m_{\mathrm{pair}}}{\varepsilon_{\mathrm{pair}}} \right\rceil. \tag{68}$$

For the m_{clean} cleaning points, use

$$T_{\mathrm{clean}} := \left\lceil \beta_X^{-1} \log \frac{m_{\mathrm{clean}}}{\varepsilon_{\mathrm{clean}}} \right\rceil. \tag{69}$$

For an ideal high candidate, $|X| \geq hN$; for an ideal low candidate, $|X| \geq \delta_0 N/2 \geq 7hN/16$. Thus $\beta_X = h/4$ is valid in the completeness proof. In the soundness proof, block exhaustion merely discards a candidate. On the degree-validation soundness event, every candidate that proceeds beyond the degree test also has density greater than β_X.

The resulting asymptotic bounds are

$$T_{\mathrm{prop}} = O(\alpha^{-1}\Lambda), \qquad T_{\mathrm{pair}}, T_{\mathrm{clean}} = O(\alpha^{-1/2}\Lambda). \tag{70}$$

B.3.5 VALIDATION AND CLEANING BLOCK SIZES

Lemma 9 gives the explicit choices

$$m_{\mathrm{deg,H}} := \left\lceil 128 h^{-1} \log \frac{2}{\zeta} \right\rceil, \tag{71}$$

$$m_{\mathrm{deg,L}}(\delta_0) := \left\lceil 256 \delta_0^{-1} \log \frac{2}{\zeta} \right\rceil, \tag{72}$$

$$m_{\mathrm{pair}} := \left\lceil 512 \log \frac{2}{\zeta} \right\rceil. \tag{73}$$

For cleaning, a uniform additive error $1/16$ over the 2^n possible query points is obtained with

$$m_{\mathrm{clean}} := \left\lceil 128 \left(n \log 2 + \log \frac{2}{\rho_{\mathrm{clean}}} \right) \right\rceil. \tag{74}$$

Since every retained layer has $\delta_0 \geq 7h/8$,

$$m_{\mathrm{deg,H}}, m_{\mathrm{deg,L}} = O(\alpha^{-1/2}\Lambda), \qquad m_{\mathrm{pair}}, m_{\mathrm{clean}} = O(\Lambda). \tag{75}$$

In the resource table below, $m_{\mathrm{deg,L}}$ denotes the maximum of $m_{\mathrm{deg,L}}(\delta_0)$ over retained layers. Since $\delta_0 \geq 7h/8$, this maximum is deterministic.

All validation statements are conditional on the randomness revealed before the current candidate is fixed. The conditional error bound is uniform over that past, so summing over the deterministic number $R_\star$ of possible slots is valid.

Table 1: Worst-case resources for persistent BSG compilation.

stage	number of slots	raw S-samples	membership queries to S
representation sample	1	m_{rep}	0
layer-density estimates	J_h	$2J_h m_{\mathrm{den}}$	$J_h m_{\mathrm{den}} q_Q$
codegree samples	at most J_h	$J_h m_{\mathrm{code}}$	0
high proposals	R_{H}	$3R_{\mathrm{H}} T_{\mathrm{prop}}$	$R_{\mathrm{H}} T_{\mathrm{prop}}$
low proposals	$J_h R_{\mathrm{L}}$	$3J_h R_{\mathrm{L}} T_{\mathrm{prop}}$	$2J_h R_{\mathrm{L}} T_{\mathrm{prop}} q_Q$
degree tests	at most $R_\star$	$R_{\mathrm{H}} m_{\mathrm{deg,H}} + J_h R_{\mathrm{L}} m_{\mathrm{deg,L}}$	$R_{\mathrm{H}} m_{\mathrm{deg,H}} + J_h R_{\mathrm{L}} m_{\mathrm{deg,L}} q_Q$
candidate pair blocks	at most $R_\star$	$2R_\star m_{\mathrm{pair}} T_{\mathrm{pair}}$	$2m_{\mathrm{pair}} T_{\mathrm{pair}}(R_{\mathrm{H}} + J_h R_{\mathrm{L}} q_Q)$
bad-pair evaluation	at most $R_\star$	0	$R_{\mathrm{H}} m_{\mathrm{pair}} q_Q + J_h R_{\mathrm{L}} m_{\mathrm{pair}} q_{\mathcal{G}}$
cleaning blocks	at most $R_\star$	$R_\star m_{\mathrm{clean}} T_{\mathrm{clean}}$	$m_{\mathrm{clean}} T_{\mathrm{clean}}(R_{\mathrm{H}} + J_h R_{\mathrm{L}} q_Q)$

B.3.6 Failure arithmetic

The soundness contribution is

$$\begin{aligned}&\rho_{\mathrm{rep}} + \rho_{\mathrm{den}} + J_h \rho_{\mathrm{code}} + 2R_\star \zeta + R_\star \rho_{\mathrm{clean}} \\ &\quad = \frac{\rho_\wedge}{40} + \frac{\rho_\wedge}{40} + \frac{\rho_\wedge}{40} + \frac{\rho_\wedge}{50} + \frac{\rho_\wedge}{100} = \frac{21}{200}\rho_\wedge.\end{aligned}$$

For completeness, the candidate counts in equation 57 give

$$\Pr[\text{miss every good high candidate}] \le \exp\left(-\frac{7\alpha}{2^8} R_{\mathrm{H}}\right) \le \frac{\rho_{\mathrm{cmp}}}{40},$$

$$\Pr[\text{miss every good centre in the visible layer}] \le \exp\left(-\frac{63}{128} R_{\mathrm{L}}\right) \le \frac{\rho_{\mathrm{cmp}}}{40 J_h}.$$

The exhaustion contributions follow from Lemma 10 and equation 60:

$$R_\star \varepsilon_{\mathrm{prop}} = R_\star \varepsilon_{\mathrm{pair}} = R_\star \varepsilon_{\mathrm{clean}} = \frac{\rho_{\mathrm{cmp}}}{100}.$$

Finally,

$$2R_\star \zeta = \frac{\rho_\wedge}{50} \le \frac{\rho_{\mathrm{cmp}}}{50}.$$

Thus the full completeness sum is bounded by

$$\frac{21}{200} + \frac{1}{40} + \frac{1}{100} + \frac{1}{100} + \frac{1}{100} + \frac{1}{50} = \frac{9}{50}.$$

B.3.7 Worst-case resource accounting

The candidate and layer counts satisfy

$$J_h = O(\Lambda), \qquad R_{\mathrm{H}} = O(\alpha^{-1}\Lambda), \qquad R_{\mathrm{L}} = O(\Lambda), \qquad R_\star = O(\alpha^{-1}\Lambda^2). \tag{76}$$

The following table charges every possible slot, even when the algorithm halts earlier. A raw sample means one call to the uniform sampler for S.

Summing the raw-sample column and substituting equation 64, equation 70, equation 75, and equation 76 gives

$$S_{\mathrm{compile}} = O(\alpha^{-2}\Lambda^4). \tag{77}$$

The membership-query column is dominated by low-branch bad-pair evaluation,

$$J_h R_{\mathrm{L}} m_{\mathrm{pair}} q_{\mathcal{G}} = O(\alpha^{-4}\Lambda^5),$$

and every other term is no larger. Hence

$$Q_{\mathrm{compile}} = O(\alpha^{-4}\Lambda^5). \tag{78}$$

A high-output membership query evaluates at most m_{clean} compiled difference predicates, with cost

$$O(m_{\mathrm{clean}} q_Q) = O(\alpha^{-2}\Lambda^2).$$

A low-output query evaluates one compiled difference predicate and at most m_{clean} graph predicates, giving

$$q_{Y_D} = O(q_Q + m_{\mathrm{clean}} q_{\mathcal{G}}) = O(\alpha^{-4}\Lambda^3). \tag{79}$$

A high record stores one representation sample, one difference, and one cleaning sample. A low record additionally stores one threshold, one empirical density count, one codegree sample, and one centre. If that count is $k_\tau \in \{0, \dots, m_{\mathrm{den}}\}$, then the derived value

$$\delta_{\tau,0} = \max\left\{0, \frac{k_\tau}{m_{\mathrm{den}}} - \frac{h}{2^7}\right\}$$

is reconstructed from k_τ and the public parameters; it is not stored as a separate rational number. Thus the count requires $O(\log m_{\mathrm{den}})$ bits, while the threshold τ is stored by its dyadic index. All later threshold comparisons use this reconstructed expression under the exact-public-parameter arithmetic convention stated at the start of Appendix A. The representation and codegree samples are the dominant components, and

$$|D|_{\mathrm{bits}} = O(n m_{\mathrm{rep}}) = O(n\alpha^{-2}\Lambda). \tag{80}$$

Accounting for vector arithmetic and predicate evaluation, the ordinary running time is

$$O\big(\alpha^{-4}\Lambda^5 \operatorname{poly}(n)\big).$$

All bounds in this appendix are deterministic upper bounds over every random execution.

B.4 Deduction of the main-text compiler theorem

The detailed theorem above gives promise-free soundness, high-energy completeness, a deterministic membership predicate once the record is fixed, and the displayed worst-case compilation and per-query bounds. Proposition 22 states that, for any fixed record satisfying the retained-size conclusion, independent rejection blocks of length

$$\left\lceil (c_{\mathrm{BSG}}\sqrt{\alpha})^{-1} \log \frac{t}{\zeta} \right\rceil$$

fail with total probability at most ζ, and conditioned on all blocks succeeding return exactly $U_{Y_D}^{\otimes t}$. These are precisely the clauses of Theorem 5.

C Size-Oblivious Algorithmic PFR

The persistent compiler of Appendix B produces a fixed implicit set together with membership access and a prescribed finite block of exact samples, but it does not reveal the cardinality of that set. This appendix records the clean-input PFR interface needed by the outer learners. The mathematical PFR step is based on the algorithmic framework of Castro–Silva et al. (Castro-Silva et al., 2026); the wrapper below makes the input size unnecessary, keeps every randomized stage within a deterministic finite budget, and exposes an adaptive-invocation interface compatible with Appendix B.

The construction is promise based: throughout this appendix the input $A \subseteq \mathbb{F}_2^n$ is nonempty and satisfies

$$|A + A| \le K|A|. \tag{81}$$

The routine does not try to certify this promise. Its role in the full paper is downstream of a persistent BSG core whose small-doubling guarantee has already been certified on the global good event.

Theorem 23 (Size-oblivious algorithmic PFR)**.** *There is a nondecreasing polynomial P_{PFR} and a nondecreasing computable function F_{PFR} with the following property. Given $K \ge 1$, $\rho \in (0,1)$, deterministic membership access to a nonempty $A \subseteq \mathbb{F}_2^n$ satisfying equation 81, and independent exact uniform samples from A, a randomized algorithm that is not given $|A|$ outputs, with probability at least $1-\rho$, a basis for a subspace $V \le \mathbb{F}_2^n$ such that*

$$|V| \le |A|, \qquad \mathcal{N}_V(A) \le P_{\mathrm{PFR}}(K). \tag{82}$$

The number of samples requested from A is fixed before they are drawn and obeys

$$s_{\mathrm{PFR}}(n,\rho) = O\left(n + \log\frac{1}{\rho}\right). \tag{83}$$

The membership-query complexity and ordinary running time are at most

$$F_{\mathrm{PFR}}(K)\,\mathrm{poly}\left(n, \log\frac{1}{\rho}\right), \qquad F_{\mathrm{PFR}}(K) = 2^{O(K)}K^{O(\log K)}. \tag{84}$$

The same statement holds conditionally after an arbitrary adaptive past: if a past transcript fixes a nonempty set A, its deterministic membership predicate, and the promise equation 81, then supplying a fresh block with law $U_A^{\otimes s_{\mathrm{PFR}}(n,\rho)}$ and fresh internal randomness gives conditional failure probability at most ρ.

The polynomial in the theorem may be enlarged, if necessary, to the fixed nondecreasing dominating polynomial P_{PFR} used in the main text. No later argument depends on its exact coefficients.

C.1 Localization without knowing the input size

A short sample block first replaces the possibly exponentially sparse ambient instance by a dense instance inside an explicitly known subspace. Draw t_{loc} independent samples $X_1,\dots,X_{t_{\mathrm{loc}}} \sim U_A$, let

$$U := \mathrm{span}(X_1,\dots,X_{t_{\mathrm{loc}}}), \qquad A_0 := A \cap U,$$

and write $r := \dim U$.

Lemma 24 (Localization and safe size bracketing)**.** *For every $\rho_{\mathrm{loc}}, \rho_{\mathrm{size}} \in (0,1)$, one can choose*

$$t_{\mathrm{loc}} = \left\lceil 8\left(n + \log\frac{1}{\rho_{\mathrm{loc}}}\right)\right\rceil \tag{85}$$

so that, except with probability ρ_{loc},

$$|A_0| \ge \frac{1}{2}|A|, \qquad \mathrm{span}(A_0) = U, \qquad |A_0 + A_0| \le L|A_0|, \quad L := 2K. \tag{86}$$

On this event, using $2^{O(L)}\log(1/\rho_{\mathrm{size}})$ uniform samples from the explicit subspace U and the same number of membership queries to A, the algorithm computes integers $\underline{N}, \widehat{N}$ such that, except with an additional probability ρ_{size},

$$1 \le \underline{N} \le |A_0| \le \widehat{N} < 4|A_0|. \tag{87}$$

If $|A_0| \ge 4$, then also $\underline{N} \ge |A_0|/4$.

Proof. Let $U_j = \mathrm{span}(X_1,\dots,X_j)$. As long as $|A \cap U_j| < |A|/2$, the next exact sample lies outside U_j with conditional probability greater than $1/2$, and every such event increases the dimension by one. Coupling these growth indicators from below by independent Bernoulli$(1/2)$ variables and applying the conditional Chernoff bound from Appendix A proves the first claim for the choice equation 85. Since the sampled points lie in A_0 and span U, we also have $\mathrm{span}(A_0) = U$; the doubling estimate follows from $|A_0 + A_0| \le |A + A| \le K|A| \le 2K|A_0|$.

For the size bracket we use the finite-field spanning estimate from the clean algorithmic PFR framework (Castro-Silva et al., 2026, Theorem 5.1): if $B \subseteq W$, $\mathrm{span}(B) = W$, and $|B+B| \le L|B|$, then

$$|W| \le \frac{2^{2L}}{2L}|B|. \tag{88}$$

Applied to $A_0 \subseteq U$, this gives the density lower bound $\theta_0 := |A_0|/|U| \ge 2^{-2L}$. The learner can sample uniformly from U using its basis, so a multiplicative estimate $\widehat{\theta}$ of θ_0 with relative error at most $1/4$ requires only $2^{O(L)}\log(1/\rho_{\mathrm{size}})$ membership tests. Set

$$N_- := \frac{2}{3}\widehat{\theta}|U|, \qquad N_+ := \frac{4}{3}\widehat{\theta}|U|,$$

$$\underline{N} := \max\{1, \lfloor N_- \rfloor\}, \qquad \widehat{N} := \begin{cases} 1, & N_+ = 0, \\ 2^{\lceil \log_2 N_+ \rceil}, & N_+ > 0. \end{cases}$$

On the concentration event,

$$\frac{1}{2}|A_0| \le N_- \le |A_0| \le N_+ \le \frac{5}{3}|A_0|.$$

Integer and dyadic rounding give equation 87; if $|A_0| \ge 4$, rounding the lower endpoint loses at most another factor two. □

The upper scale $\widehat{N}$ is used only to choose a dense linear model; the lower scale $\underline{N}$ is used only to enforce the final nontriviality condition $|V| \le |A|$.

C.2 A BOUNDED DENSE MODEL WITH DETERMINISTIC IMAGE ORACLES

Fix a coordinate isomorphism $\chi_U : U \to \mathbb{F}_2^r$. Let

$$m_0 := \left\lceil \log_2 \widehat{N} + 4 \log_2 L + c_{\mathrm{mod}} \right\rceil, \qquad b_L := \lceil 2L \rceil + 4, \qquad m := \max\{m_0, r - b_L\}, \quad (89)$$

where c_{mod} is a sufficiently large absolute constant. If $m \ge r$, use the fixed injection $u \mapsto (\chi_U(u), 0^{m-r})$. If $m < r$, draw five independent uniform $m \times r$ binary matrices and use the first full-row-rank matrix; if none exists, the current model trial fails.

Lemma 25 (Safe dense linear model and image access)**.** *Assume the events in Lemma 24. The construction above has the following properties.*

(a) If $m < r$, model generation fails with probability below $1/32$, and conditioned on success the selected map $\pi : U \to \mathbb{F}_2^m$ is exactly uniform among surjective linear maps.

(b) The safeguard in equation 89 is inactive on the good size event, so $m = m_0$ and

$$2^m \le C_{\mathrm{mod}} L^4 |A_0| \quad (90)$$

for an absolute constant C_{mod}. Conditional on successful model generation, π is Freiman-good,

$$\ker \pi \cap (4A_0) = \{0\}, \quad (91)$$

with probability at least $3/4$.

(c) For every generated model, whether or not it is Freiman-good, the image set $S_\pi := \pi(A_0)$ has a deterministic membership oracle and a deterministic canonical inverse oracle $g_\pi : S_\pi \to \mathbb{F}_2^r$, each using at most $2^{b_L} = 2^{O(K)}$ membership queries to A and $2^{O(K)} \operatorname{poly}(n)$ time.

Proof. A uniform $m \times r$ matrix with $m < r$ has full row rank with probability

$$\prod_{i=0}^{m-1} (1 - 2^{i-r}) > \frac{1}{2}.$$

Thus all five draws fail with probability below 2^{-5}. Every full-rank matrix has the same probability of being the first successful draw, so the conditional distribution is uniform over surjections.

The spanning bound equation 88 gives $r \le \log_2 |A_0| + 2L$. Since $\widehat{N} \ge |A_0|$, the term $r - b_L$ is strictly below m_0; hence $m = m_0$. The upper bracket $\widehat{N} < 4|A_0|$ then gives equation 90. For a uniform surjection and any fixed nonzero $z \in \mathbb{F}_2^r$,

$$\Pr[\pi(z) = 0] = \frac{2^{r-m} - 1}{2^r - 1} \le 2^{-m}.$$

Plünnecke's inequality gives $|4A_0| \le L^4 |A_0|$, so a union bound and the choice of c_{mod} imply equation 91 with probability at least $3/4$. If $m \ge r$, the fixed injection is Freiman-good deterministically.

Finally, $m \ge r - b_L$ ensures $\dim \ker \pi \le b_L$ on every transcript. To answer a query at $x \in \mathbb{F}_2^m$, solve $\pi(u) = x$ in U. If the fiber is empty then $x \notin S_\pi$; otherwise enumerate its at most 2^{b_L} points, test membership in A, and return the first accepted point in a fixed total order. Composing that point with χ_U gives the canonical inverse value $g_\pi(x)$. □

The safeguard $m \geq r - b_L$ is important only off the main good event: it prevents a bad size estimate from creating exponentially large fibers and therefore makes every trial a finite worst-case computation.

C.3 The imported restricted-homomorphism primitive

The wrapper uses one nontrivial external algorithmic primitive: the restricted homomorphism routine from the clean algorithmic PFR framework. The theorem below is a normalized restatement of the restricted-homomorphism lemma of Castro–Silva et al. (Castro-Silva et al., 2026, Lemma 5.10); we spell out the interface because both the hypothesis and the conclusion are measured on the ambient scale 2^m.

Theorem 26 (Algorithmic restricted homomorphism; imported interface)**.** *There are a polynomial P_{RH} and a nondecreasing computable function $F_{\mathrm{RH}}(\mathcal{K}) = \mathcal{K}^{O(\log(2\mathcal{K}))}$ such that the following holds. Let $S' \subseteq \mathbb{F}_2^m$ and $f : S' \to \mathbb{F}_2^r$. If*

$$\left|\left\{(x_1,x_2,x_3,x_4) \in (S')^4 : \begin{array}{l} x_1+x_2 = x_3+x_4, \\ f(x_1)+f(x_2) = f(x_3)+f(x_4) \end{array}\right\}\right| \geq \frac{2^{3m}}{\mathcal{K}}, \tag{92}$$

then, given membership-query access to S' and value-query access to f, the routine returns with probability at least 0.7 a linear map $M : \mathbb{F}_2^m \to \mathbb{F}_2^r$ and $v \in \mathbb{F}_2^r$ satisfying

$$|\{x \in S' : f(x) = Mx + v\}| \geq \frac{2^m}{P_{\mathrm{RH}}(\mathcal{K})}. \tag{93}$$

Whenever the premise holds, its query and time bounds are $F_{\mathrm{RH}}(\mathcal{K})\,\mathrm{poly}(m+r)$. In the wrapper below, we fix a deterministic resource cap $B_{\mathrm{RH}}(\mathcal{K}, m, r)$ for the imported implementation; after enlarging F_{RH} if necessary, this cap is at most $F_{\mathrm{RH}}(\mathcal{K})\,\mathrm{poly}(m+r)$ on the displayed asymptotic scale. Every invocation is forcibly stopped at this cap; a timeout, an explicit failure output, or a malformed output is treated as a failed trial. Hence every wrapper execution has a finite worst-case budget even when the premise is false, while on promised inputs the cap does not alter the imported guarantee.

On a Freiman-good model, g_π is the true inverse of $\pi|_{A_0}$. Hence every additive quadruple in S_π is respected by g_π, and the number of respected quadruples equals $E(A_0)$. Since

$$E(A_0) \geq \frac{|A_0|^4}{|A_0 + A_0|} \geq \frac{|A_0|^3}{L},$$

while equation 90 gives $2^m/|A_0| \leq C_{\mathrm{mod}} L^4$, the premise equation 92 holds with

$$\mathcal{K}_\star(L) := C_{\mathrm{RH}} L^{13}, \qquad C_{\mathrm{RH}} := C_{\mathrm{mod}}^3. \tag{94}$$

Thus the imported primitive is invoked only at an L-dependent parameter.

C.4 Fresh validation, covering, and safe truncation

Replace P_{RH} by a nondecreasing polynomial upper envelope with $P_{\mathrm{RH}} \geq 1$ and set

$$\beta_L := \frac{1}{2P_{\mathrm{RH}}(\mathcal{K}_\star(L))}. \tag{95}$$

For a returned affine candidate (M, v), define its absolute agreement

$$p_\pi(M, v) := \Pr_{X \sim U_{\mathbb{F}_2^m}}[X \in S_\pi \text{ and } g_\pi(X) = MX + v]. \tag{96}$$

A successful restricted-homomorphism call on a Freiman-good model has $p_\pi(M, v) \geq 2\beta_L$. The wrapper validates (M, v) on a fresh uniform block from $\mathbb{F}_2^m$ and accepts when the empirical agreement is at least $3\beta_L/2$.

Lemma 27 (Validated affine agreement yields a safe PFR output)**.** *Assume the localization and size-bracketing events and the model-size bound equation 90.*

(a) For any conditional error budget $\zeta \in (0,1)$, a fresh validation block of

$$m_{\mathrm{agr}} = \left\lceil 16\beta_L^{-1} \log \frac{2}{\zeta} \right\rceil \tag{97}$$

accepts every candidate with $p_\pi(M,v) \geq 2\beta_L$ except with probability ζ, while a candidate with $p_\pi(M,v) \leq \beta_L$ is accepted with probability at most ζ.

(b) Every candidate with $p_\pi(M,v) > \beta_L$ determines a subspace $H \leq U$ satisfying

$$\mathcal{N}_H(A) < \frac{L}{\beta_L}. \tag{98}$$

This implication does not require the realized model to be Freiman-good.

(c) Let

$$d_V := \min\{\dim H, \lfloor \log_2 \underline{N} \rfloor\}$$

and choose any d_V-dimensional $V \leq H$. Then

$$|V| \leq |A|, \qquad [H:V] \leq C_{\mathrm{tr}} L^4, \tag{99}$$

for an absolute constant C_{tr}, and consequently

$$\mathcal{N}_V(A) \leq 2C_{\mathrm{tr}} L^5 P_{\mathrm{RH}}(C_{\mathrm{RH}} L^{13}). \tag{100}$$

Proof. Part (a) is the multiplicative gap test of Lemma 9 with $(a,\theta,b) = (1,3/2,2)$ and scale $q = \beta_L$; the constant in equation 97 is sufficient for both tails.

For part (b), let

$$C := \{f_\pi(x) : x \in S_\pi,\ g_\pi(x) = Mx + v\},$$

where $f_\pi(x)$ denotes the canonical support point represented by $g_\pi(x)$. Distinct image points have distinct canonical representatives, so $|C| = p_\pi(M,v)2^m$. Moreover, C lies in one affine coset of $H := \chi_U^{-1}(\mathrm{Im}\, M)$. Since $m \geq m_0 \geq \log_2 \widehat{N}$,

$$|C| > \beta_L 2^m \geq \beta_L \widehat{N} \geq \beta_L |A_0| \geq \frac{\beta_L}{2}|A|.$$

Also $C + C \subseteq H$ and $C \subseteq A$, hence

$$|A + C| \leq |A + A| \leq K|A| < \frac{2K}{\beta_L}|C| = \frac{L}{\beta_L}|C|.$$

Ruzsa's covering lemma therefore covers A by fewer than L/β_L translates of $C + C \subseteq H$, proving equation 98. Notice that this argument uses only validated agreement and the original small-doubling promise; Freiman-goodness is needed for completeness, not for soundness of an accepted candidate.

For part (c), the definition gives $|V| \leq \underline{N} \leq |A_0| \leq |A|$. If $V \neq H$, then $|V| > \underline{N}/2$, so

$$[H:V] < \frac{2|H|}{\underline{N}} \leq \frac{2 \cdot 2^m}{\underline{N}}.$$

When $|A_0| \geq 4$, Lemma 24 gives $\underline{N} \geq |A_0|/4$, and equation 90 yields $[H:V] = O(L^4)$. When $|A_0| < 4$, the trivial lower bound $|V| \geq 1$ and the same model-size estimate give the same conclusion after increasing the absolute constant. Refining each H-coset into $[H:V]$ many V-cosets and substituting $\beta_L^{-1} = 2P_{\mathrm{RH}}(C_{\mathrm{RH}} L^{13})$ proves equation 100. □

C.5 Algorithm, correctness, and deterministic budgets

The wrapper has one preparation stage followed by a fixed number of model trials. It returns the first validated subspace; different successful trials need not produce comparable subspaces.

Algorithm 2 SIZEOBLIVIOUSPFR

Require: membership access and exact uniform samples from A; $K \geq 1$; failure parameter ρ

draw the prescribed localization sample block, form $U = \mathrm{span}(X_1, \ldots, X_{t_{\mathrm{loc}}})$, and set $A_0 = A \cap U$
estimate the density of A_0 in the explicit subspace U and compute $\underline{N}, \widehat{N}$
for a prescribed number R_{trial} of independent trials **do**
    generate a bounded-time linear model $\pi : U \to \mathbb{F}_2^m$
    expose deterministic membership/value oracles for (S_π, g_π)
    invoke the restricted-homomorphism routine with resource cap $B_{\mathrm{RH}}(\mathcal{K}_\star, m, r)$
    if the call times out, returns failure, or does not return a well-formed pair (M, v) of the prescribed dimensions **then**
        continue to the next model trial
    end if
    validate the returned affine map using a fresh ambient block
    if validation passes **then**
        construct H, truncate to $V \leq H$ using $\underline{N}$, and **return** V
    end if
end for
return FAIL

Proof of Theorem 23. Allocate

$$\rho_{\mathrm{loc}} := \frac{\rho}{4}, \qquad \rho_{\mathrm{size}} := \frac{\rho}{4}, \qquad R_{\mathrm{trial}} := \left\lceil \frac{\log(4/\rho)}{\log(5/3)} \right\rceil, \qquad \zeta := \frac{\rho}{8R_{\mathrm{trial}}}.$$

Condition on successful localization and size bracketing. In one fresh model trial, bounded-time model generation succeeds with probability at least $31/32$; conditional on success, the model is Freiman-good with probability at least $3/4$. On a Freiman-good model, the discussion after Theorem 26 verifies the restricted- homomorphism premise with parameter $C_{\mathrm{RH}}L^{13}$. The prescribed resource cap therefore does not truncate the promised execution, and the imported routine returns a well-formed good affine candidate with probability at least $7/10$. A good affine candidate has true agreement at least $2\beta_L$, and fresh validation accepts it with probability at least $1 - \zeta \geq 7/8$. Therefore one trial produces a good accepted candidate with probability at least

$$\frac{31}{32} \cdot \frac{3}{4} \cdot \frac{7}{10} \cdot \frac{7}{8} = \frac{4557}{10240} > \frac{2}{5}. \tag{101}$$

The probability that all R_{trial} trials miss is at most $(3/5)^{R_{\mathrm{trial}}} \leq \rho/4$.

Soundness of accepted candidates uses a different event. A candidate may be chosen adaptively from the entire preceding execution, but its validation block is fresh. By Lemma 9 and sequential conditional accounting from Lemma 8, the probability that any accepted candidate has true agreement at most β_L is at most $R_{\mathrm{trial}}\zeta = \rho/8$. On the localization, size-bracketing, and validation-soundness events, every returned subspace therefore satisfies Lemma 27. The total failure probability is at most

$$\frac{\rho}{4} + \frac{\rho}{4} + \frac{\rho}{4} + \frac{\rho}{8} = \frac{7\rho}{8} < \rho.$$

Taking

$$P_{\mathrm{PFR}}(K) := 2C_{\mathrm{tr}}(2K)^5 P_{\mathrm{RH}}\big(C_{\mathrm{RH}}(2K)^{13}\big) \tag{102}$$

proves equation 82; enlarging this polynomial if necessary makes it agree with the dominating convention of the main text.

Only the localization stage requests samples from the input set A, so equation 83 is fixed before the block is drawn. The size estimate samples from the explicitly known subspace U and therefore requires only membership queries to A. Every image-model operation costs at most $2^{O(K)}$ membership queries to A. There are $O(\log(1/\rho))$ model trials; every imported call is capped at $B_{\mathrm{RH}}(\mathcal{K}, m, r)$, with $\mathcal{K} = C_{\mathrm{RH}}(2K)^{13}$, and this cap is bounded by $F_{\mathrm{RH}}(\mathcal{K})\,\mathrm{poly}(m + r)$ for the fixed imported implementation. Thus bad models cannot create unbounded executions. Validation has polynomial dependence on K. Thus the query and time bound has the form equation 84. More explicitly, for a sufficiently large absolute constant C, one may take

$$F_{\mathrm{PFR}}(K) = 2^{CK}\left(1 + F_{\mathrm{RH}}(C_{\mathrm{RH}}(2K)^{13}) + P_{\mathrm{RH}}(C_{\mathrm{RH}}(2K)^{13})\right) = 2^{O(K)}K^{O(\log K)}.$$

Finally condition on an arbitrary adaptive past transcript that fixes the input predicate and supplies a fresh product block $U_A^{\otimes s_{\mathrm{PFR}}(n,\rho)}$. After conditioning, every random choice made above has exactly the standalone law, so the same proof gives conditional failure probability at most ρ. This is precisely the form needed when the input A is a persistent core generated earlier in the outer learner. Proposition 22 supplies such a fresh exact block from the fixed core using finite rejection budgets. □

C.6 Resource ledger and interface summary

For completeness, the deterministic budgets can be read directly from the proof. With $L = 2K$, let

$$\mathcal{K}_\star = C_{\mathrm{RH}} L^{13}, \qquad \beta_L = \frac{1}{2P_{\mathrm{RH}}(\mathcal{K}_\star)}.$$

The input-sample block has size

$$t_{\mathrm{loc}} = \left\lceil 8\left(n + \log\frac{4}{\rho}\right)\right\rceil.$$

The localized density is at least 2^{-2L}, so

$$m_{\mathrm{size}} = O\left(2^{2L}\log\frac{1}{\rho}\right)$$

uniform points of U suffice for size bracketing. Every model trial uses exactly five random matrices when $m < r$, at most one resource-capped restricted-homomorphism call, and (only when that call returns a well-formed affine candidate) a validation block of length

$$m_{\mathrm{agr}} = O\left(P_{\mathrm{RH}}(C_{\mathrm{RH}}L^{13})\log\frac{R_{\mathrm{trial}}}{\rho}\right).$$

The number of trials is $R_{\mathrm{trial}} = O(\log(1/\rho))$. Every query to the image model expands to at most $2^{b_L} = 2^{O(K)}$ membership queries to the original input. All these budgets are fixed before the corresponding fresh random blocks are drawn.

The interface used later can therefore be summarized as

$$\boxed{\begin{array}{c}\text{fixed membership predicate for a nonempty } A\\ + U_A^{\otimes s_{\mathrm{PFR}}(n,\rho)} \text{ supplied exactly}\\ + |A+A| \le K|A|\end{array}} \quad\Longrightarrow\quad \boxed{\begin{array}{c}|V| \le |A|,\\ \mathcal{N}_V(A) \le P_{\mathrm{PFR}}(K)\end{array}}$$

with conditional failure probability at most ρ even after an arbitrary adaptive past. This is the only PFR interface required by the moderate- and arbitrary-corruption learners.

D Structural Lifting from an Observed Core

A cover learned on a core $Y \subseteq B$ transfers to a compatible support A whenever the core contains enough clean mass. The two deterministic bounds below provide this transfer and enforce the size cap. They apply to every compatible support after the random core is fixed.

For $\gamma \in [0,1]$ and $\eta \in [0,1)$, define

$$\Delta(\gamma,\eta) := \gamma - \max\{\gamma, 1-\gamma\}\eta. \tag{103}$$

Lemma 28 (Clean mass forced by observed size)**.** *Suppose $Y \subseteq B \subseteq \mathbb{F}_2^n$, $|Y| \ge \gamma|B|$, and $|A \triangle B| \le \eta|A|$. Then*

$$|A \cap Y| \ge [\Delta(\gamma,\eta)]_+|A|, \qquad [x]_+ := \max\{x,0\}. \tag{104}$$

In particular, when $0 < \gamma \le 1/2$,

$$\Delta(\gamma,\eta) = \gamma - (1-\gamma)\eta. \tag{105}$$

Proof. Write

$$u := \frac{|A \setminus B|}{|A|}, \qquad v := \frac{|B \setminus A|}{|A|}.$$

Then $u, v \geq 0$, $u + v \leq \eta$, and $|B| = (1 - u + v)|A|$. Because $Y \subseteq B$,

$$\begin{aligned} |A \cap Y| &\geq |Y| - |B \setminus A| \\ &\geq \gamma|B| - v|A| \\ &= \big(\gamma - \gamma u - (1-\gamma)v\big)|A| \\ &\geq \Delta(\gamma, \eta)|A|. \end{aligned}$$

Combining with the trivial lower bound zero proves equation 104. If $\gamma \leq 1/2$, then $1 - \gamma \geq \gamma$, giving equation 105. □

The next lemma is the common structural transfer used in both algorithmic regimes. Its hypotheses are phrased directly in terms of clean mass; the moderate-corruption learner obtains that mass from Lemma 28, while the arbitrary-corruption learner obtains it by averaging across the peeled cores.

Lemma 29 (Clean-core structural lifting with size-oblivious truncation)**.** *Let $A, Y, B \subseteq \mathbb{F}_2^n$, with $A \neq \varnothing$ and $Y \subseteq B$, and let $H \leq \mathbb{F}_2^n$. Suppose*

$$|A + A| \leq K|A|, \qquad |A \cap Y| \geq \delta|A|, \qquad |H| \leq |Y|, \qquad \mathcal{N}_H(Y) \leq P, \tag{106}$$

for some $\delta > 0$, and assume only the observable-size comparison

$$|B| < 2|A|. \tag{107}$$

Define $V \leq H$ by the fixed rule

$$V := H \quad \text{if } H = \{0\}, \qquad V := \operatorname{span}(h_1, \ldots, h_{r-1}) \quad \text{otherwise}, \tag{108}$$

where $h_1, \ldots, h_r$ is the basis of H returned by the PFR routine. Then

$$|V| \leq |A|, \qquad \mathcal{N}_V(A) \leq \frac{2KP}{\delta}. \tag{109}$$

The construction of V uses neither A, $|A|$, η, nor δ.

Proof. Set $C := A \cap Y$, so $|C| \geq \delta|A|$. Cover Y by $m := \mathcal{N}_H(Y) \leq P$ distinct H-cosets. One coset contains a subset $C' \subseteq C$ of size at least $|C|/m \geq \delta|A|/m$. Since $C' \subseteq A$,

$$|A + C'| \leq |A + A| \leq K|A| \leq \frac{Km}{\delta}|C'|.$$

Ruzsa's covering lemma therefore yields

$$\mathcal{N}_H(A) \leq \frac{Km}{\delta} \leq \frac{KP}{\delta}. \tag{110}$$

It remains only to enforce the nontriviality condition in the cost q_A. From $|H| \leq |Y| \leq |B| < 2|A|$, if $H = \{0\}$ then $|H| \leq |A|$ already. Otherwise the rule equation 108 produces an index-two subspace $V \leq H$, so

$$|V| = \frac{|H|}{2} < |A|.$$

Every H-coset splits into two V-cosets, hence

$$\mathcal{N}_V(A) \leq 2\mathcal{N}_H(A) \leq \frac{2KP}{\delta}.$$

The rule depends only on the returned basis of H, proving the final claim. □

Two specializations will be used repeatedly. First, if $|A \triangle B| \leq \eta|A|$, then $|B| \leq (1 + \eta)|A| < 2|A|$, so equation 107 is automatic. Second, if a core satisfies $|Y| \geq \gamma|B|$ and $\Delta(\gamma, \eta) > 0$, then Lemma 28 and Lemma 29 give

$$q_A(V) = \mathcal{N}_V(A) \leq \frac{2KP}{\Delta(\gamma, \eta)}. \tag{111}$$

This is the one-core transfer used in the moderate-corruption theorem. In the list theorem, the analysis instead proves $|A \cap Y_j| \geq \delta_A|A|$ for a serving index $j = j(A)$ and applies equation 109 directly. No size-only purity argument is needed there.

E Intermediate List Construction under Arbitrary Corruption

We construct candidates for the arbitrary-corruption learner by peeling persistent cores, applying PFR to each core, and lifting its cover to latent supports. The compiler, PFR, and lifting guarantees are supplied by Appendices B–D. One observation-dependent success event will serve the entire identification set.

Throughout, fix $K \geq 1$, $\eta \in [0, 1)$, and

$$\varepsilon := 1 - \eta, \qquad \mathcal{I} := \mathcal{I}_{K,\eta}(B) \neq \varnothing.$$

For every $A \in \mathcal{I}$, write $C_A := A \cap B$. The shared insertion–deletion budget implies

$$|C_A| \geq \varepsilon|A|, \qquad |B| \leq (1+\eta)|C_A|, \qquad |C_A| > \frac{|B|}{2}. \tag{112}$$

The last inequality is the strict-majority property from the main text.

Set

$$\alpha_\star := \frac{\varepsilon}{128K}, \qquad \gamma_\star := c_{\mathrm{BSG}}\sqrt{\alpha_\star}, \qquad L_\star := C_{\mathrm{BSG}}\alpha_\star^{-4}, \qquad P_\star := P_{\mathrm{PFR}}(L_\star). \tag{113}$$

Thus

$$\gamma_\star = \Theta\Big(\sqrt{\varepsilon/K}\Big), \qquad L_\star = O(K^4\varepsilon^{-4}).$$

The deterministic peel cap is

$$M_\star := \left\lceil \frac{\ln 4}{-\ln(1-\gamma_\star)} \right\rceil = O\Big(\sqrt{K/\varepsilon}\Big). \tag{114}$$

Only K, η, and the absolute compiler constants enter these parameters.

Proposition 30 (Intermediate universal structural list)**.** *There exist absolute constants $C_{\mathrm{list}}, C_{\mathrm{list,L}} > 0$ and a computable function G_{list} such that the following holds. Let $\varepsilon = 1-\eta$, suppose $\mathcal{I}_{K,\eta}(B) \neq \varnothing$, and let $\rho \in (0,1)$. Given K, η, ρ and $\mathsf{SampMQ}(B)$, a randomized algorithm always terminates, returning either FAIL or a finite list $\mathcal{L}(B) = \{V_1, \ldots, V_m\}$. With probability at least $1-\rho$, the returned list is nonempty and simultaneously satisfies*

$$1 \leq m \leq C_{\mathrm{list}}\sqrt{K/\varepsilon},$$

and, for every $A \in \mathcal{I}_{K,\eta}(B)$,

$$\min_{V \in \mathcal{L}(B)} q_A(V) \leq C_{\mathrm{list}} K^{3/2}\varepsilon^{-3/2} P_{\mathrm{PFR}}\big(C_{\mathrm{list,L}} K^4 \varepsilon^{-4}\big).$$

The number of exact samples from B is polynomial in $K, \varepsilon^{-1}, n, \log(1/\rho)$, while membership queries and ordinary running time are at most

$$G_{\mathrm{list}}(K, \varepsilon^{-1}) \left(1 + n + \log\frac{1}{\rho}\right)^{O(\sqrt{K/\varepsilon})}.$$

Appendix F selects one common output from this list.

E.1 The relative residual certificate

Measure the clean mass remaining in a residual relative to the *initial observed clean mass* $C_A = A \cap B$. This gives a constant-fraction stopping target even when $1 - \eta$ is small.

Lemma 31 (Retained observed mass forces residual energy)**.** *Let $R \subseteq B$. If some $A \in \mathcal{I}$ satisfies*

$$|A \cap R| \geq \frac{1}{2}|C_A|, \tag{115}$$

then

$$\frac{E(R)}{|R|^3} \geq \frac{\varepsilon}{16K(1+\eta)^3} \geq \frac{\varepsilon}{128K} = \alpha_\star. \tag{116}$$

Consequently, if $E(R) < \alpha_\star|R|^3$, then simultaneously for every $A \in \mathcal{I}$,

$$|A \cap R| < \frac{1}{2}|A \cap B|. \tag{117}$$

Proof. Fix a witness A to equation 115. Since $A \cap R \subseteq A$,

$$E(R) \geq E(A \cap R) \geq \frac{|A \cap R|^4}{|(A \cap R) + (A \cap R)|} \geq \frac{|C_A|^4}{16K|A|}.$$

Also $R \subseteq B$, so by equation 112,

$$|R|^3 \leq |B|^3 \leq (1+\eta)^3 |C_A|^3.$$

Therefore

$$\frac{E(R)}{|R|^3} \geq \frac{|C_A|}{16K(1+\eta)^3|A|} \geq \frac{\varepsilon}{16K(1+\eta)^3}.$$

Since $1 + \eta < 2$, the latter quantity is at least $\varepsilon/(128K)$. The final statement is the contrapositive, applied separately to every $A \in \mathcal{I}$. □

The contrapositive is uniform over the identification set: a low-energy residual retains less than half of the initial clean overlap for every compatible explanation.

E.2 Finite residual access

At round j, previously compiled cores $Y_0, \ldots, Y_{j-1}$ determine the residual

$$R_0 := B, \qquad R_j := B \setminus \bigcup_{i<j} Y_i. \tag{118}$$

Conditional on the complete past, R_j is one fixed deterministic set. Its predicate is

$$\mathbf{1}_{R_j}(x) = \mathbf{1}_B(x) \prod_{i<j} \big(1 - \mathbf{1}_{Y_i}(x)\big). \tag{119}$$

We use the fixed constants

$$\xi_{\mathrm{res}} := \frac{1}{32}, \qquad \theta_{\mathrm{res}} := \frac{3}{16}, \qquad \beta_{\mathrm{res}} := \frac{5}{32}.$$

A residual-density estimate with additive error ξ_{res} therefore creates a strict gap between stopping and continuation.

Proposition 32 (Residual test and exact compiler input)**.** *Fix a past transcript that determines $R \subseteq B$, a requested sample count $s \geq 1$, and a conditional failure budget $\rho_0 \in (0,1)$. There is a finite procedure using fresh samples from B and the fixed predicate of R with the following properties.*

Draw

$$m_{\mathrm{res}} := \left\lceil \frac{1}{2\xi_{\mathrm{res}}^2} \log \frac{2}{\rho_0} \right\rceil \tag{120}$$

samples from B, let $\widehat{p}_R$ be the empirical fraction lying in R, and return STOP *when $\widehat{p}_R \leq \theta_{\mathrm{res}}$. Except with conditional probability ρ_0,*

$$\mathsf{STOP} \Longrightarrow \frac{|R|}{|B|} \leq \frac{7}{32}, \qquad \mathsf{CONTINUE} \Longrightarrow \frac{|R|}{|B|} > \frac{5}{32}. \tag{121}$$

In particular, on this event a STOP *decision implies*

$$|A \cap R| < \frac{1}{2}|A \cap B| \qquad \textit{for every } A \in \mathcal{I}. \tag{122}$$

On the density-estimation good event in equation 121, and conditional on CONTINUE*, generate each requested point by an independent rejection block from U_B of length*

$$T_{\mathrm{res}} := \left\lceil \beta_{\mathrm{res}}^{-1} \log \frac{s}{\rho_0} \right\rceil. \tag{123}$$

The probability that any of the s blocks is exhausted is at most ρ_0; conditioned on all blocks succeeding, their accepted values have the exact product law

$$U_R^{\otimes s}. \tag{124}$$

Proof. Conditional on the past, the indicators $\mathbf{1}_R(X)$ for fresh $X \sim U_B$ are i.i.d. Bernoulli with mean $p_R = |R|/|B|$. Hoeffding's inequality and equation 120 give $|\widehat{p}_R - p_R| \le 1/32$ except with probability ρ_0. The threshold $3/16 = 6/32$ then yields equation 121.

If STOP, then

$$|A \cap R| \le |R| \le \frac{7}{32}|B| < \frac{1}{4}|B| < \frac{1}{2}|A \cap B|$$

for every compatible A, where the last inequality uses equation 112.

On the density-estimation good event, a continue branch has $p_R > \beta_{\mathrm{res}}$. Conditional on that event and the realized CONTINUE decision, the fixed-block exact rejection lemma, Lemma 10, with acceptance set R, gives block-exhaustion probability at most ρ_0/s and exact U_R law for the first accepted point conditional on success. Independence of the blocks gives equation 124. □

E.3 Peeling and the uniform terminal invariant

Let

$$\rho_0 := \frac{\rho}{10M_\star}. \tag{125}$$

Let $s^\star_{\mathrm{BSG}}$ be a deterministic upper bound on the number of uniform samples requested by Theorem 15 at parameters $(\alpha_\star, \rho_0, \rho_0)$, and let

$$s^\star_{\mathrm{PFR}} := s_{\mathrm{PFR}}(n, \rho_0)$$

be the fixed sample count in Theorem 23.

For a reached round j, condition on the complete past and on a successful residual sample bridge, so that the compiler receives exact product samples from the now-fixed residual R_j. Let D_j denote its output and define the two one-sided compiler failure events

$$\mathcal{F}_{\mathrm{snd},j} := \big\{D_j \neq \mathsf{FAIL} \text{ and the structural conclusion of Theorem 15 fails}\big\}, \tag{126}$$

$$\mathcal{F}_{\mathrm{cmp},j} := \big\{E(R_j) \ge \alpha_\star |R_j|^3 \text{ and } D_j = \mathsf{FAIL}\big\}. \tag{127}$$

By the promise-free soundness and high-energy completeness clauses of Theorem 15, each event has conditional probability at most ρ_0, uniformly over the realized past. Thus, outside $\mathcal{F}_{\mathrm{snd},j}$, every nonfailure record is structurally valid; outside $\mathcal{F}_{\mathrm{cmp},j}$, an internal compiler FAIL implies $E(R_j) < \alpha_\star |R_j|^3$.

The learner starts with $R_0 = B$. In a reached round $j < M_\star$, it performs the following operations using fresh random blocks.

1. Apply Proposition 32 to R_j with $s = s^\star_{\mathrm{BSG}}$. On STOP, terminate normally. Exhaustion of a residual rejection block is an algorithmic FAIL.
2. Run the persistent compiler on the exact R_j-sample block with parameter $\alpha_\star$ and soundness/completeness budgets (ρ_0, ρ_0). If the compiler returns its internal FAIL, terminate normally. Otherwise let $Y_j \subseteq R_j$ be the fixed output set.
3. Outside $\mathcal{F}_{\mathrm{snd},j}$,
$$|Y_j| \ge \gamma_\star |R_j|, \qquad |Y_j + Y_j| \le L_\star |Y_j|. \tag{128}$$
Since a reached continue branch has $|R_j| > (5/32)|B|$, such a core obeys
$$\frac{|Y_j|}{|B|} > \beta_Y, \qquad \beta_Y := \frac{5\gamma_\star}{32}. \tag{129}$$
Generate the $s^\star_{\mathrm{PFR}}$ samples required by the PFR routine directly from B, using independent Y_j-rejection blocks of length
$$T_Y := \left\lceil \beta_Y^{-1} \log \frac{s^\star_{\mathrm{PFR}}}{\rho_0} \right\rceil. \tag{130}$$
Exhaustion is an algorithmic FAIL; conditional on success, the block has exact law $U_{Y_j}^{\otimes s^\star_{\mathrm{PFR}}}$.
4. Invoke Theorem 23 on the fixed set Y_j, with doubling parameter $L_\star$ and error budget ρ_0. A PFR failure is an algorithmic FAIL. Otherwise let H_j be the returned subspace and define V_j from H_j by the fixed truncation rule equation 108.

5. Store Y_j, V_j and set $R_{j+1} := R_j \setminus Y_j$.

The distinction between the two kinds of failure is important. A compiler FAIL is a *normal stop*: on its completeness good event it certifies that the current residual does not meet the energy threshold. Rejection-block exhaustion and PFR failure carry no such structural meaning and cause the outer learner to return FAIL.

Lemma 33 (Uniform terminal invariant and serving core)**.** *Assume the following observation-dependent good events hold in every reached round: the density conclusion of Proposition 32; success of the residual and core sample bridges; neither $\mathcal{F}_{\mathrm{snd},j}$ nor $\mathcal{F}_{\mathrm{cmp},j}$ occurs; and the PFR output is valid. Suppose the outer learner does not return algorithmic* FAIL.

If it produces $m \geq 1$ cores, then $m \leq M_\star$, the cores are pairwise disjoint, and its final residual R_{fin} satisfies simultaneously for every $A \in \mathcal{I}$,

$$|A \cap R_{\mathrm{fin}}| < \frac{1}{2}|A \cap B|. \tag{131}$$

Consequently,

$$\sum_{j<m} |A \cap Y_j| > \frac{1}{2}|A \cap B| \geq \frac{\varepsilon}{2}|A| \qquad (A \in \mathcal{I}), \tag{132}$$

and for every $A \in \mathcal{I}$ there exists an index $j = j(A) < m$ such that

$$|A \cap Y_j| \geq \frac{\varepsilon}{2M_\star}|A|. \tag{133}$$

Moreover, on the same good events the process cannot terminate normally before producing its first core.

Proof. Because $Y_j \subseteq R_j$ and $|Y_j| \geq \gamma_\star |R_j|$,

$$|R_{j+1}| = |R_j| - |Y_j| \leq (1 - \gamma_\star)|R_j|.$$

Thus the stored cores are pairwise disjoint and, after $M_\star$ valid peels,

$$|R_{M_\star}| \leq (1 - \gamma_\star)^{M_\star}|B| \leq \frac{1}{4}|B|.$$

By equation 112,

$$\frac{1}{4}|B| < \frac{1}{2}|A \cap B|$$

for every $A \in \mathcal{I}$. Hence reaching the deterministic peel cap implies equation 131.

If instead the residual-density test returns STOP, then equation 122 gives the same conclusion. Finally, suppose the persistent compiler returns its internal FAIL. Outside $\mathcal{F}_{\mathrm{cmp},j}$, an internal compiler FAIL at the final reached round can happen only when

$$E(R_{\mathrm{fin}}) < \alpha_\star |R_{\mathrm{fin}}|^3.$$

Lemma 31 then gives equation 131. These exhaust the normal stopping mechanisms.

Since

$$B \setminus R_{\mathrm{fin}} = \bigsqcup_{j<m} Y_j,$$

the terminal invariant gives

$$\sum_{j<m} |A \cap Y_j| = |A \cap B| - |A \cap R_{\mathrm{fin}}| > \frac{1}{2}|A \cap B| \geq \frac{\varepsilon}{2}|A|,$$

which is equation 132. Averaging over $m \leq M_\star$ cores yields equation 133.

It remains to rule out a normal stop at $R_0 = B$. For any $A \in \mathcal{I}$,

$$|A \cap R_0| = |A \cap B| > \frac{1}{2}|A \cap B|,$$

so a density stop is impossible on its good event, and Lemma 31 gives $E(B) \geq \alpha_\star |B|^3$. Hence a compiler FAIL is also impossible outside the corresponding one-sided event $\mathcal{F}_{\mathrm{cmp},0}$. Therefore at least one valid core is produced. □

For the serving index in equation 133, the PFR output satisfies

$$|H_j| \le |Y_j|, \qquad \mathcal{N}_{H_j}(Y_j) \le P_\star.$$

Applying Lemma 29 with

$$\delta = \frac{\varepsilon}{2M_\star}$$

gives

$$q_A(V_j) = \mathcal{N}_{V_j}(A) \le \frac{4KM_\star}{\varepsilon} P_\star. \tag{134}$$

Substituting equation 113 and equation 114 yields

$$q_A(V_j) \le O(K^{3/2}\varepsilon^{-3/2}) P_{\mathrm{PFR}}\big(O(K^4\varepsilon^{-4})\big), \tag{135}$$

uniformly over $A \in \mathcal{I}$, with the serving index allowed to depend on A.

E.4 Failure accounting and resources

All probability statements are organized chronologically. At the beginning of a reached round, the past transcript fixes the residual predicate. Conditional on the density-test good event and successful residual bridge, the compiler receives an exact product block from that fixed residual. Conditional on a valid compiled core and successful direct core bridge, the PFR routine receives an exact product block from that fixed core. Thus every invoked guarantee can be applied conditionally on the entire past.

There are at most six charged events per reached round: residual-density estimation, residual-bridge exhaustion, the one-sided events $\mathcal{F}_{\mathrm{snd},j}$ and $\mathcal{F}_{\mathrm{cmp},j}$, core-bridge exhaustion, and PFR failure. With the slot budget equation 125, sequential conditional failure accounting, Lemma 8, gives

$$\Pr[\mathcal{E}_{\mathrm{global}}^c] \le 6M_\star\rho_0 = \frac{3}{5}\rho < \rho. \tag{136}$$

On $\mathcal{E}_{\mathrm{global}}$, all hypotheses of Lemma 33 hold simultaneously.

For explicit complexity bounds, put

$$\Lambda_E := 1 + n + \log\frac{K}{\varepsilon} + \log\frac{M_\star}{\rho}. \tag{137}$$

Proposition 34 (Finite sample bound and recursive oracle cost)**.** *The outer list learner can be implemented with*

$$O\left(\left(\frac{K}{\varepsilon}\right)^{5/2}\Lambda_E^5\right) \tag{138}$$

uniform samples from B.

There is a nondecreasing computable function $G_{\mathrm{list}}(K,\varepsilon^{-1})$ such that its membership-query complexity and ordinary running time are both at most

$$G_{\mathrm{list}}(K,\varepsilon^{-1})\left(1+n+\log\frac{1}{\rho}\right)^{O(M_\star)} = G_{\mathrm{list}}(K,\varepsilon^{-1})\left(1+n+\log\frac{1}{\rho}\right)^{O(\sqrt{K/\varepsilon})}. \tag{139}$$

All rejection blocks and all sample counts are deterministic once K, η, ρ, n are fixed.

Proof. At one reached round, the density test uses $O(\log(1/\rho_0))$ samples from B. By Theorem 15, the compiler requests

$$s^\star_{\mathrm{BSG}} = O(\alpha_\star^{-2}\Lambda_E^4) = O\big((K/\varepsilon)^2\Lambda_E^4\big)$$

uniform samples from its residual input. Proposition 32 realizes those requests from B with blocks of length

$$T_{\mathrm{res}} = O\left(\log\frac{s^\star_{\mathrm{BSG}}}{\rho_0}\right) = O(\Lambda_E).$$

Hence the residual bridge uses

$$O\big((K/\varepsilon)^2\Lambda_E^5\big)$$

samples from B per round.

The PFR routine requests

$$s^{\star}_{\mathrm{PFR}} = O\left(n + \log\frac{1}{\rho_0}\right) = O(\Lambda_E)$$

samples from Y_j. Since $\beta_Y = \Theta(\gamma_\star)$, equation 130 gives

$$T_Y = O\left(\sqrt{K/\varepsilon}\,\Lambda_E\right).$$

Thus the direct core bridge costs at most $O(\sqrt{K/\varepsilon}\,\Lambda_E^2)$ samples from B per round. Multiplying all per-round costs by $M_\star = O(\sqrt{K/\varepsilon})$ proves the conservative bound equation 138.

For queries, let Q_j be the worst-case number of MQ_B calls needed to evaluate the fixed residual predicate $\mathbf{1}_{R_j}(x)$ once. The compiler theorem gives a bound

$$C_\star = O(\alpha_\star^{-4}\Lambda_E^3)$$

on the number of membership queries to R_i needed for one evaluation of a stored Y_i-predicate. Therefore

$$Q_0 = 1, \qquad Q_j \le 1 + C_\star \sum_{i<j} Q_i \le (1 + C_\star)^j. \tag{140}$$

The density tests and residual rejection bridges evaluate $\mathbf{1}_{R_j}$; compiler compilation makes only a parameter-bounded number of such calls; evaluating Y_j costs at most $C_\star Q_j$; and the PFR invocation makes $F_{\mathrm{PFR}}(L_\star)\,\mathrm{poly}(\Lambda_E)$ calls to that fixed Y_j-predicate. Since $\alpha_\star^{-1}$, $L_\star$, and all nonpolynomial dependence of the PFR routine are functions only of K, ε^{-1}, summing the reached rounds and absorbing these parameter-only factors into G_{list} gives

$$G_{\mathrm{list}}(K, \varepsilon^{-1})\Lambda_E^{O(M_\star)}.$$

Replacing Λ_E by a parameter-only factor times $1 + n + \log(1/\rho)$ proves equation 139. The same recurrence bounds ordinary running time because each oracle evaluation and every linear-algebra operation used by the compiler/PFR routines has the corresponding finite worst-case cost. □

E.5 Assembly of the intermediate list guarantee

Proof of Proposition 30. Run the peeling procedure above. It always terminates: there are at most $M_\star$ reached rounds, and every randomized subroutine and rejection bridge has a predetermined finite budget. Outside the event $\mathcal{E}_{\mathrm{global}}$, the algorithm may return FAIL.

Work now on $\mathcal{E}_{\mathrm{global}}$, which has probability at least $1-\rho$ by equation 136. Because $\mathcal{I} \neq \varnothing$, Lemma 33 shows that the returned list is nonempty and has

$$1 \le m \le M_\star = O\left(\sqrt{K/\varepsilon}\right).$$

For every $A \in \mathcal{I}$, the same lemma supplies a serving core $Y_{j(A)}$, and equation 135 together with the deterministic truncation in Lemma 29 gives

$$q_A(V_{j(A)}) \le O(K^{3/2}\varepsilon^{-3/2})P_{\mathrm{PFR}}\left(O(K^4\varepsilon^{-4})\right).$$

The event $\mathcal{E}_{\mathrm{global}}$, the realized list, and all stored predicates depend only on B and the learner's randomness. The argument choosing $j(A)$ is deterministic after that event is fixed. Hence the same realized list satisfies the guarantee simultaneously for every $A \in \mathcal{I}$, which is the required $\forall A\,\exists j(A)$ quantifier order.

Finally, Proposition 34 gives the polynomial sample bound and the XP membership-query/time bound stated in Proposition 30. □

F From an Intermediate List to One Common Hypothesis

The list from Appendix E supplies feasible descriptions for the majority-coverage rule. Fresh samples let us apply that rule to the fixed list and prove Lemma 4.

Proof of Lemma 4. Condition on an arbitrary adaptive past that fixes an explicit list $V_1, \dots, V_m$ satisfying the serving property in the lemma. Put $b = \varepsilon/8$. Draw a fresh sample block from U_B. For each V_i, group sampled points by their V_i-coset, let Z_i be the union of the at most Q most frequent observed cosets, and write $\widehat{P}(Z_i)$ for their empirical mass. Keep indices with

$$\widehat{P}(Z_i) \geq \frac{1}{2} + b$$

and return, with fixed tie-breaking, a kept entry of minimum dimension. If no entry is kept, return failure.

For an elementary finite-class implementation, consider all unions of at most Q cosets of any listed subspace. This class has at most $m2^{nQ}$ nonempty descriptions: use Q ambient representatives, repeating representatives when fewer cosets are needed. Hence

$$s = \left\lceil \frac{32}{\varepsilon^2} \left(nQ \ln 2 + \ln \frac{2m}{\zeta} \right) \right\rceil$$

samples suffice, by Hoeffding and a union bound, for one event of probability at least $1 - \zeta$ on which

$$|\widehat{P}(Z) - U_B(Z)| \leq b \tag{141}$$

for every such union. The event includes the data-dependent top-frequency choices Z_i.

Fix this event. For every compatible A, choose a serving entry $V_{i(A)}$. Since $\mathcal{N}_{V_{i(A)}}(A) \leq Q$, some union Z_A of at most Q $V_{i(A)}$-cosets contains A. The shared budget gives

$$U_B(Z_A) \geq \frac{|A \cap B|}{|B|} \geq \frac{1}{1+\eta} = \frac{1}{2-\varepsilon} \geq \frac{1}{2} + \frac{\varepsilon}{4}.$$

The empirically most frequent Q cosets maximize empirical coverage for that fixed subspace, so

$$\widehat{P}(Z_{i(A)}) \geq \widehat{P}(Z_A) \geq \frac{1}{2} + \frac{\varepsilon}{8}.$$

Thus every compatible support has at least one retained serving entry. The minimum-dimension choice therefore yields

$$|V_\star| \leq |V_{i(A)}| \leq |A| \qquad \text{simultaneously for all } A \in \mathcal{I}_{K,\eta}(B).$$

The selected union satisfies $U_B(Z_\star) \geq 1/2$ by equation 141. For every compatible A, the majority-capture calculation from Section 3 gives

$$|A \cap B \cap Z_\star| \geq \frac{\varepsilon}{2}|A|.$$

One of the at most Q cosets in $Z_\star$ therefore contains a set $C \subseteq A \cap B$ of size at least $\varepsilon|A|/(2Q)$. Because $C + C \subseteq V_\star$ and $|A + C| \leq K|A|$, the maximal-disjoint-translate covering argument yields

$$\mathcal{N}_{V_\star}(A) \leq \frac{|A + C|}{|C|} \leq \frac{2KQ}{\varepsilon}.$$

The same concentration event works for all compatible supports; no union bound over A is taken. This proves the lemma. □

Dimension-free sample refinement. For a fixed V_i, unions of at most Q of its cosets have VC dimension at most Q: a shattered set must occupy distinct cosets, and more than Q points cannot be shattered. Standard uniform convergence followed by a union bound over the m fixed subspaces therefore permits

$$s = O\big(\varepsilon^{-2}\big[Q \log(2/\varepsilon) + \log(2m/\zeta)\big]\big).$$

The finite-description bound above remains available when one wants a fully elementary count.

Proof of the arbitrary-corruption clause of Theorem 3. Run Proposition 30 with failure budget $\rho/2$. Abort on its failure or if its list exceeds the prescribed parameter-only cap, which makes every downstream execution finite. On its good event, apply Lemma 4 with

$$Q = \left\lceil C_{\mathrm{list}} K^{3/2} \varepsilon^{-3/2} P_{\mathrm{PFR}}\big(C_{\mathrm{list,L}} K^4 \varepsilon^{-4}\big) \right\rceil$$

and $\zeta = \rho/2$. The two failure probabilities sum to at most ρ, and the lemma gives

$$q_A(V_\star) \le O\Big(K^{5/2}\varepsilon^{-5/2} P_{\mathrm{PFR}}(O(K^4\varepsilon^{-4}))\Big) \qquad \forall A \in \mathcal{I}_{K,\eta}(B).$$

The postprocessor makes no membership queries. Its additional sample and linear-algebra costs are polynomial in the structural parameters and the list size, so the query/time recurrence from Proposition 34 retains the XP form equation 9. This proves the full-range clause. □

G Common Learning under Moderate Corruption

Under moderate corruption, a single persistent core is large enough to contain clean mass from every compatible support. We prove this size guarantee and then apply deterministic lifting on one observation-level success event, establishing the first part of Theorem 3.

Fix

$$\alpha_K := \frac{1}{2K}, \qquad \gamma_K := c_{\mathrm{BSG}}\sqrt{\alpha_K} = \frac{c_{\mathrm{BSG}}}{\sqrt{2K}}, \qquad L_K := C_{\mathrm{BSG}}\alpha_K^{-4} = 16C_{\mathrm{BSG}}K^4. \tag{142}$$

Choose once and for all an absolute constant

$$0 < c_{\mathrm{rob}} \le \min\left\{\frac{1}{32}, \frac{c_{\mathrm{BSG}}}{4\sqrt{2}}\right\}. \tag{143}$$

The numerical values are not important; only the two constant margins below are used.

Lemma 35 (Observation energy and uniform clean mass)**.** *Let*

$$\bar\eta := \frac{c_{\mathrm{rob}}}{\sqrt{K}}, \qquad \mathcal{I}_{\mathrm{rob}} := \mathcal{I}_{K,\bar\eta}(B),$$

and suppose $\mathcal{I}_{\mathrm{rob}} \neq \varnothing$. Then $B \neq \varnothing$ and

$$E(B) \ge \alpha_K |B|^3. \tag{144}$$

Moreover, every set $Y \subseteq B$ satisfying $|Y| \ge \gamma_K|B|$ obeys, simultaneously for all $A \in \mathcal{I}_{\mathrm{rob}}$,

$$|A \cap Y| \ge \delta_K|A|, \qquad \delta_K \ge \frac{3c_{\mathrm{BSG}}}{4\sqrt{2K}} = \Omega(K^{-1/2}). \tag{145}$$

Proof. Fix $A \in \mathcal{I}_{\mathrm{rob}}$ and put $C := A \cap B$. Since $|A \triangle B| \le \bar\eta|A|$,

$$|C| \ge (1-\bar\eta)|A|, \qquad |B| \le (1+\bar\eta)|A|.$$

Also $C + C \subseteq A + A$, hence

$$E(B) \ge E(C) \ge \frac{|C|^4}{|C+C|} \ge \frac{(1-\bar\eta)^4}{K}|A|^3.$$

After normalizing by $|B|^3$,

$$\frac{E(B)}{|B|^3} \ge \frac{(1-\bar\eta)^4}{(1+\bar\eta)^3 K}. \tag{146}$$

By equation 143, $\bar\eta \le 1/32$, and the constant factor on the right is larger than $1/2$. This proves equation 144.

For the second claim, apply Lemma 28 with $\gamma = \gamma_K$ and $\eta = \bar\eta$. Since $\gamma_K < 1/2$,

$$\Delta(\gamma_K, \bar\eta) = \gamma_K - (1-\gamma_K)\bar\eta \ge \gamma_K - \bar\eta.$$

Using equation 143,

$$\gamma_K - \bar\eta \ge \frac{c_{\mathrm{BSG}}}{\sqrt{2K}} - \frac{c_{\mathrm{BSG}}}{4\sqrt{2K}} = \frac{3c_{\mathrm{BSG}}}{4\sqrt{2K}}.$$

The same lower bound holds for each compatible A, proving equation 145. □

G.1 One observation-level event

Run Theorem 15 on $S = B$ with energy parameter α_K and soundness/completeness budgets $\rho/4$. By Lemma 35, the high-energy premise holds. On the intersection of the compiler soundness and completeness events, the compiler returns a fixed record D defining $Y_D \subseteq B$ with

$$|Y_D| \geq \gamma_K |B|, \qquad |Y_D + Y_D| \leq L_K |Y_D|. \tag{147}$$

Let

$$s^\star_{\mathrm{PFR}} := s_{\mathrm{PFR}}(n, \rho/4).$$

Use Proposition 22 to generate this prescribed number of exact uniform samples from the same fixed Y_D, with bridge-failure budget $\rho/4$. Conditional on bridge success, the block has law

$$U_{Y_D}^{\otimes s^\star_{\mathrm{PFR}}}.$$

Invoke Theorem 23 on this fixed implicit set with doubling parameter L_K and failure budget $\rho/4$. On its good event the returned subspace H satisfies

$$|H| \leq |Y_D|, \qquad \mathcal{N}_H(Y_D) \leq P_{\mathrm{PFR}}(L_K). \tag{148}$$

Finally define V from the returned basis of H by the deterministic truncation rule equation 108.

All random objects in this construction depend only on B, K, the failure budget, and the learner's randomness. In particular, no distinguished $A \in \mathcal{I}_{\mathrm{rob}}$ is selected.

G.2 Proof of the moderate-corruption theorem

Moderate-corruption clause of Theorem 3. Let $\mathcal{E}_{\mathrm{rob}}$ be the intersection of the compiler soundness/completeness events, success of the finite Y_D-sample bridge, and the PFR success event. Chronological conditioning and the allocated budgets give

$$\Pr[\mathcal{E}_{\mathrm{rob}}] \geq 1 - \rho.$$

Fix a realization in $\mathcal{E}_{\mathrm{rob}}$. Equations equation 147 and equation 148 are now deterministic statements about one fixed Y_D, H, V. Lemma 35 gives, for every $A \in \mathcal{I}_{\mathrm{rob}}$,

$$|A \cap Y_D| \geq \delta_K |A|.$$

Because $|A \triangle B| \leq \bar{\eta}|A|$ with $\bar{\eta} < 1$, the size condition $|B| < 2|A|$ required by Lemma 29 is automatic. Apply that lemma with

$$\delta = \delta_K, \qquad P = P_{\mathrm{PFR}}(L_K).$$

The same realized V then satisfies, simultaneously for all $A \in \mathcal{I}_{\mathrm{rob}}$,

$$\begin{aligned} q_A(V) &= \mathcal{N}_V(A) && (149)\\ &\leq \frac{2K}{\delta_K} P_{\mathrm{PFR}}(L_K) && (150)\\ &\leq C K^{3/2} P_{\mathrm{PFR}}(16 C_{\mathrm{BSG}} K^4) =: P_{\mathrm{rob}}(K) && (151) \end{aligned}$$

for an absolute constant C. Enlarging P_{rob}, if necessary, makes it a nondecreasing polynomial. Since the good event was fixed before quantifying over A, this proves exactly

$$\Pr[\forall A \in \mathcal{I}_{\mathrm{rob}} : q_A(V) \leq P_{\mathrm{rob}}(K)] \geq 1 - \rho,$$

with no union bound over the identification set.

It remains to record complexity. At $\alpha_K = 1/(2K)$, the persistent compiler uses

$$F_1(K)\,\mathrm{poly}\left(n, \log \frac{1}{\rho}\right)$$

samples, queries, and time for a computable F_1. The retained fraction $\gamma_K = \Theta(K^{-1/2})$ makes the finite output-sampling bridge cost only an additional

$$F_2(K)\,\mathrm{poly}\left(n, \log \frac{1}{\rho}\right)$$

samples and predicate evaluations. The PFR call has doubling parameter $L_K = O(K^4)$, so Theorem 23 contributes

$$F_{\mathrm{PFR}}(O(K^4))\,\mathrm{poly}\left(n, \log\frac{1}{\rho}\right).$$

One evaluation of the fixed Y_D-predicate has parameter-dependent worst-case cost by Theorem 15. Absorbing all K-only factors into a computable $F_{\mathrm{rob}}(K)$ proves the FPT sample, query, and time bound in the theorem. □

H Sample-Only Common Learning: Full Proof

This appendix supplies the detailed proof of Theorem 1. Let $p = p(K)$ be the integer-valued clean PFR covering bound fixed in the main text.

Lemma 36 (Shared budget and majority capture)**.** *For every $A \in \mathcal{I}_{K,\eta}(B)$,*

$$|B| \le (1+\eta)|A \cap B|.$$

If $Z \subseteq G$ satisfies $|Z \cap B| \ge |B|/2$, then

$$|A \cap B \cap Z| \ge \frac{\varepsilon}{2}|A|.$$

Proof. Set $a = |A|$, $c = |A \cap B|$, $d = |A \setminus B|$, and $o = |B \setminus A|$. From $d + o \le \eta(c + d)$ we get $o \le \eta c - \varepsilon d \le \eta c$, proving the first assertion. For the second,

$$|A \cap B \cap Z| \ge |B \cap Z| - o \ge \frac{a - d + o}{2} - o = \frac{a - d - o}{2} \ge \frac{\varepsilon a}{2}.$$

□

Lemma 37 (VC dimension of parallel affine unions)**.** *The family of unions of at most p parallel affine subspaces of $\mathbb{F}_2^n$ has VC dimension $O((n+1)p\log(2p))$.*

Proof. First consider all affine subspaces, allowing the empty set. Any $n+2$ distinct points are affinely dependent, so some point lies in the affine span of the others. No affine subspace can contain all the other points while excluding that point. Thus the affine-subspace class has VC dimension at most $v = n + 1$. By Sauer's lemma its number of traces on $t \ge v$ points is at most $(et/v)^v$. Unions of p arbitrary affine subspaces have at most $(et/v)^{vp}$ traces, and parallel unions form a subclass. If t points are shattered, then $2^t \le (et/v)^{vp}$. Writing $x = t/(vp)$ gives $2^x \le epx$, which forces $x = O(\log(2p))$. □

Proof of Theorem 1. Let

$$\mathcal{Z}_p = \left\{ \bigcup_{j=1}^{p} (z_j + V) : V \le G,\ z_j \in G \right\},$$

where repeated representatives allow fewer than p cosets. Set $\delta = \varepsilon/8$. Draw the sample block prescribed in equation 7; standard VC uniform convergence and Lemma 37 give, with probability at least $1 - \rho$,

$$\sup_{Z \in \mathcal{Z}_p} |\widehat{P}(Z) - U_B(Z)| \le \delta. \tag{152}$$

On this event, among all candidate descriptions with $\widehat{P}(Z) \ge 1/2 + \delta$, choose one whose underlying subspace has minimum dimension, with fixed tie-breaking.

For every compatible A, clean PFR supplies a subspace V_A and a union Z_A of at most p of its cosets such that

$$A \subseteq Z_A, \qquad |V_A| \le |A|.$$

By Lemma 36,

$$U_B(Z_A) \ge \frac{|A \cap B|}{|B|} \ge \frac{1}{1+\eta} = \frac{1}{2-\varepsilon} \ge \frac{1}{2} + \frac{\varepsilon}{4}.$$

Hence every such witness is empirically feasible on equation 152. The minimum-dimension rule gives

$$|\widehat{V}| \le |V_A| \le |A| \qquad \forall A \in \mathcal{I}_{K,\eta}(B).$$

For the selected union $\widehat{Z}$, the same uniform event gives $U_B(\widehat{Z}) \ge 1/2$. Lemma 36 therefore yields

$$|A \cap B \cap \widehat{Z}| \ge \frac{\varepsilon}{2}|A| \qquad \forall A \in \mathcal{I}_{K,\eta}(B).$$

One of the at most p $\widehat{V}$-cosets contains $C \subseteq A \cap B$ with $|C| \ge \varepsilon|A|/(2p)$. Since $C + C \subseteq \widehat{V}$ and $|A + C| \le K|A|$, a maximal family of disjoint translates $t + C$, $t \in A$, has size at most $|A + C|/|C|$ and maximality implies that the corresponding translates of $C + C$ cover A. Thus

$$\mathcal{N}_{\widehat{V}}(A) \le \frac{|A + C|}{|C|} \le \frac{2Kp}{\varepsilon}.$$

The same sample event works for every compatible support, establishing the joint guarantee.

For an elementary alternative to the VC bound, every subspace is the row span of some $n \times n$ binary matrix and the p representatives use np bits. Thus there are at most 2^{n^2+np} descriptions, and Hoeffding plus a union bound gives the sufficient count

$$O\left(\frac{n^2 + np + \log(1/\rho)}{\varepsilon^2}\right).$$

Exhaustive enumeration over these descriptions gives the finite exponential-time implementation.

For sharpness, take $G = \mathbb{F}_2^t$, $B = \{0\}$, and $\varepsilon = 2^{-t}$. Both $A_0 = \{0\}$ and $A_1 = G$ are compatible at $K = 1$. Any common output must satisfy $|V| \le |A_0| = 1$, so $V = \{0\}$, and its cost on A_1 is $|G| = \varepsilon^{-1}$. Conversely $V = \{0\}$ costs at most $|G|$ on every subset of G. This proves the displayed equality in the theorem. □

Population interpretation. The proof also gives a deterministic population rule over candidate descriptions: among all $(V, z_1, \ldots, z_p)$ whose induced set $Z = \bigcup_j (z_j + V)$ satisfies $U_B(Z) \ge 1/2$, select a description minimizing $\dim V$. The same argument gives cost $2Kp/\varepsilon$. The sample-only learner is exactly a uniform-convergence implementation of this population structural ERM.

I Conditional Sample-Only Hardness via Learning Parities with Noise

We prove Theorem 2 by reducing search-LPN at the stated noise rate to coarse quotient-cover learning. Blocking forces a bounded-cost output to reveal a candidate secret, and a frozen-table coupling realizes the required fixed-support observation model.

Search-LPN convention. For a dimension d and fixed $0 < \theta < 1/2$, a secret $s \in \mathbb{F}_2^d$ is uniform, and the learner receives independent samples

$$(X, \langle s, X\rangle + E), \qquad X \sim U_{\mathbb{F}_2^d}, \quad E \sim \mathrm{Bernoulli}(\theta).$$

This is the standard noisy-parity problem studied, for example, by Blum, Kalai, and Wasserman (Blum et al., 2003). The assumption used here is that no randomized polynomial-time algorithm recovers the uniform secret with constant success probability for all sufficiently large d. The reduction below in fact succeeds uniformly for every fixed secret, so this average-case convention is sufficient.

Fix the theorem's constant quality budget Q, let

$$\ell = \lfloor \log_2 Q \rfloor, \qquad k = \ell + 1, \qquad D = kd, \qquad n = k(d + 1),$$

and define, for $s \in \mathbb{F}_2^d$,

$$A_s = \{(x_1, \ldots, x_k, \langle s, x_1\rangle, \ldots, \langle s, x_k\rangle) : x_i \in \mathbb{F}_2^d\}. \tag{153}$$

It is a D-dimensional linear subspace and hence has doubling one.

Lemma 38 (Decoding a coarse subspace output)**.** *From a basis of any $V \le \mathbb{F}_2^n$ satisfying $q_{A_s}(V) \le Q$, one can construct in $\mathrm{poly}(n, k, 2^{k+\ell})$ time a list of at most $k2^{k+\ell} \le 2kQ^2$ vectors of $\mathbb{F}_2^d$ containing s. On arbitrary inputs, discard any V with $\dim V \notin [D-\ell, D]$ before enumeration.*

Proof. Because A_s is a subspace,

$$\mathcal{N}_V(A_s) = 2^{D-\dim(A_s \cap V)}.$$

The assumed cost gives $\dim(A_s \cap V) \ge D - \ell$, while the size cap in q_{A_s} gives $\dim V \le D$. Hence $\dim V^\perp \le k + \ell$ and

$$\dim(A_s^\perp \cap V^\perp) = n - \dim(A_s + V) \ge k - \ell = 1.$$

Directly from equation 153,

$$A_s^\perp = \{(a_1 s, \dots, a_k s, a_1, \dots, a_k) : a \in \mathbb{F}_2^k\}.$$

Enumerate the at most $2^{k+\ell}$ vectors $w = (b_1, \dots, b_k, a_1, \dots, a_k) \in V^\perp$, and insert b_j in the candidate list whenever $a_j = 1$. A nonzero vector in the displayed intersection has some $a_j = 1$, and then $b_j = s$. □

Lemma 39 (Frozen-table coupling)**.** *Set $\theta = \eta_0/(8k)$. For every complete input tuple $x = (x_1, \dots, x_k) \in \mathbb{F}_2^D$, independently draw and freeze $E(x) \in \mathbb{F}_2^k$ with independent Bernoulli(θ) coordinates, and let B_E be the resulting noisy graph. Then*

$$|B_E| = |A_s| = 2^D,$$

$$\Pr_E[|A_s \triangle B_E| > \eta_0 |A_s|] \le \exp(-9\eta_0^2 2^D/32),$$

and for any adaptive sample-only algorithm making at most T requests, the transcript law under fresh grouped LPN samples and the law obtained by first drawing E and then sampling independently from U_{B_E} have total variation distance at most

$$\binom{T}{2} 2^{-D}.$$

The frozen table is not conditioned on satisfying the support promise.

Proof. Different complete inputs produce different graph points. The two graphs differ at input x exactly when $E(x) \neq 0$, and each changed input contributes two points to the symmetric difference. The changed-input indicators are independent with mean

$$a = 1 - (1-\theta)^k \le k\theta = \eta_0/8.$$

The support promise holds whenever their empirical mean over 2^D inputs is at most $\eta_0/2$. Hoeffding with gap $3\eta_0/8$ gives the stated invalid-table probability.

For the coupling, use the same independent uniform complete input tuples in both experiments. On the first visit to a tuple, use its fresh LPN noise to reveal the previously unexposed frozen-table entry. Until a complete tuple repeats, the two streams coincide. A repetition among T tuples has probability at most $\binom{T}{2}/2^D$. Unexposed entries can then be filled independently, giving exactly the advance-sampled table distribution. Pad adaptively stopped executions to T requests and couple internal randomness as well. □

Lemma 40 (Uniqueness at doubling one through the critical radius)**.** *If $0 \le \eta \le 1/3$, then $\mathcal{I}_{1,\eta}(B)$ contains at most one support.*

Proof. Every nonempty doubling-one set in $\mathbb{F}_2^n$ is an affine subspace. For distinct affine subspaces A, A' of the same size a, their intersection has size at most $a/2$, so $|A \triangle A'| \ge a = (|A| + |A'|)/2$. If their sizes differ, assume $|A'| = b \ge 2a = 2|A|$; then

$$|A \triangle A'| \ge b - a \ge (a+b)/3.$$

Thus simultaneous compatibility is impossible for $\eta < 1/3$. At $\eta = 1/3$, equality could only occur for a nested pair of sizes a and $2a$. Their distance is a, but the two integer discrepancy budgets sum to $\lfloor a/3 \rfloor + \lfloor 2a/3 \rfloor = a - 1$, since a is a power of two. The triangle inequality again gives a contradiction. □

Proof of Theorem 2. Assume for contradiction a polynomial-time sample-only learner with the stated guarantee. Given search-LPN samples for secret $s \in \mathbb{F}_2^d$, group them in blocks of k, thereby producing samples in the ambient space of equation 153. If size information is allowed, supply $\dim A_s = kd$ and $|A_s| = 2^{kd}$. Let $T(n)$ be a polynomial upper bound on the learner's sample requests.

In the frozen-table experiment, every table satisfying the support budget is a legitimate fixed instance for the learner. Average the learner's constant success guarantee over such tables without conditioning the simulator on validity. Lemma 39 shows that the LPN-driven simulation returns a serving subspace with probability at least

$$c_0 - \exp(-9\eta_0^2 2^{kd}/32) - \binom{T(n)}{2} 2^{-kd},$$

where $c_0 > 1/2$ can be obtained, if necessary, by running the learner an absolute constant number of independent times and merging the returned structural lists. If one run succeeds with probability $c > 0$, then after $r = O(1)$ runs the merged list contains a serving entry with probability $1 - (1-c)^r$; its length remains polynomial, and this amplification does not require identifying which entry serves the latent support.

For every output subspace, run Lemma 38. A polynomial-size structural list therefore yields a polynomial-size union of candidate secrets because Q and k are fixed. Validate all candidates on one fresh LPN block. The true secret has mismatch probability θ, whereas every false candidate has mismatch probability exactly $1/2$. The constant gap $1/2 - \theta$, Hoeffding's inequality, and a union bound over the polynomial candidate set recover the true secret with constant probability for all sufficiently large d. This contradicts the stated search-LPN assumption.

Finally, $\eta_0 < 1/3$ and Lemma 40 show that every valid observation used in the theorem has at most one compatible K=1 support. The reduction therefore does not rely on latent non-identifiability. □

Assumption boundary. The reduction invokes search-LPN at $\theta = \eta_0/(8k)$, as specified in Theorem 2. Using an assumption at another rate would require a noise-rate reduction.

J Exact Affine-Support Recovery at $K = 1$

The following postprocessing result uses Fourier learning after finding an explicit affine flat in which the latent support has noticeable density. The output is the entire affine subspace, specified by an offset and a basis.

Lemma 41 (Heavy-Fourier primitive)**.** *For a fixed $f : \mathbb{F}_2^r \to \{0,1\}$, define $\widehat{f}(\xi) = \mathbb{E}_z f(z)(-1)^{\langle \xi, z\rangle}$. Given value-query access to f, a threshold $0 < \tau < 1$, and a failure budget $0 < \zeta < 1$, one can in*

$$\mathrm{poly}(r, \tau^{-1}, \log(1/\zeta))$$

time and queries produce a list containing every ξ with $|\widehat{f}(\xi)| \geq \tau$ and containing at most $O(\tau^{-2})$ entries, except with probability ζ. The routine can have deterministic finite budgets on every execution.

This is the standard Goldreich–Levin interface (Goldreich & Levin, 1989). For finite budgets, estimate squared Fourier prefix masses to error $\tau^2/4$, retain estimates above $\tau^2/2$, and cap each level at $\lceil 4/\tau^2 \rceil$. Parseval controls the retained count and heavy prefixes survive. For a length-j prefix a, the bounded unbiased estimator is $f(u,w)f(v,w)(-1)^{\langle a, u+v\rangle}$, for independent uniform $u, v \in \mathbb{F}_2^j$, $w \in \mathbb{F}_2^{r-j}$. Fixed estimation slots and the cap bound all executions. Dimension zero is evaluated directly.

Theorem 42 (From a coarse list to exact recovery)**.** *Let $A \subseteq \mathbb{F}_2^n$ be a nonempty affine subspace and let $|A \triangle B| \leq \eta |A|$ with a supplied $0 \leq \eta < 1/3$. Suppose one is given $m \geq 1$ explicitly based subspaces $V_1, \ldots, V_m$, at least one of which satisfies $q_A(V_i) \leq Q$, where $Q \geq 1$ is supplied. The other entries need not satisfy any guarantee. Given exact uniform samples and membership queries to B, a randomized postprocessor recovers A exactly with probability at least $1 - \rho$, in*

$$\mathrm{poly}\left(n, m, Q, (1-3\eta)^{-1}, \log(1/\rho)\right)$$

samples, membership queries, and time. The same guarantee holds after an adaptive past fixing the list, using fresh randomness thereafter.

Proof. Write $a = |A|$, $C = A \cap B$, and $H = A - A$. All stages below have prescribed finite budgets. Set $\ell = \lfloor \log_2 Q \rfloor$, $\kappa = 1/2 - \eta > 1/6$, and

$$t = \left\lceil \frac{8(\ell + \ln 4)}{\kappa} \right\rceil, R = \left\lceil 2(1+\eta)^{t+1} \ln \frac{4}{\rho} \right\rceil.$$

For each V_i, make R independent proposals using $t+1$ fresh samples $X_0, \ldots, X_t$ from B, and form the affine flat

$$W = X_0 + \mathrm{span}(V_i, X_1 - X_0, \ldots, X_t - X_0).$$

All proposals are explicit affine flats, so membership and uniform sampling within them are elementary linear-algebra operations.

A useful proposal occurs with noticeable probability. Fix a serving entry V. The image of A in G/V is an affine subspace of dimension at most ℓ. Conditional on all proposal samples lying in C, they are independent uniform samples from C. Any proper affine subspace of this image has a preimage in A of size at most $a/2$. Since $|C| \geq (1-\eta)a$, a new conditional sample lies outside a currently proper affine span with probability at least

$$\frac{|C| - a/2}{|C|} \geq \frac{1}{2} - \eta = \kappa.$$

There are at most ℓ affine-dimension increases to make after the first sample. Until the image is spanned, each new conditional sample increases its affine dimension with conditional probability at least κ. Couple each not-yet-spanned step with an independent $\mathrm{Bernoulli}(\kappa)$ success that is dominated by the actual increase event. Then failure to accumulate the required ℓ increases is contained in the event $\{\mathrm{Binomial}(t, \kappa) < \ell\}$. Since

$$t\kappa \geq 8(\ell + \ln 4),$$

a lower-tail Chernoff bound gives

$$\Pr[\text{fewer than } \ell \text{ increases}] \leq \Pr[\mathrm{Binomial}(t, \kappa) < \ell] < \frac{1}{2}$$

(with the case $\ell = 0$ immediate). Thus the projected samples span the entire image with probability at least $1/2$. The probability all samples are clean is at least $(1+\eta)^{-(t+1)}$. Thus a single proposal equals $A + V$ with probability at least $\frac{1}{2}(1+\eta)^{-(t+1)}$. Such a proposal contains A and has

$$|W| = |A + V| = \mathcal{N}_V(A)|V| \leq Qa.$$

The probability of missing all such proposals for the serving entry is at most $\rho/4$. Since $\eta < 1/3$ and $\kappa > 1/6$, R is polynomial in Q and $\log(1/\rho)$ with absolute exponents.

A sample test certifies containment. For any proposed affine flat W, let $b_W = \mathbb{P}_{U_B}[X \in W]$. If $A \subseteq W$, then

$$b_W \geq u := \frac{1}{1+\eta}.$$

If $A \not\subseteq W$, then $|A \cap W| \leq a/2$. Writing $d = |A \setminus B|$, $o = |B \setminus A|$, one obtains

$$b_W \leq \frac{a/2 + o}{a - d + o} \leq v := \frac{1}{2(1-\eta)}.$$

The second inequality follows from $d + o \leq \eta a$; maximizing the ratio sets $d = \eta a - o$, after which the ratio is nonincreasing in o. The gap is

$$u - v = \frac{1 - 3\eta}{2(1-\eta^2)} > 0.$$

Estimate b_W with fresh samples from B to accuracy $(u-v)/4$, accepting only if the estimate is at least $(u+v)/2$. On the simultaneous test event, every accepted flat contains A, and every flat containing A is accepted.

A second test certifies density. Put $\beta = 1/Q$. Using uniform points of the explicit flat W and membership queries to B, estimate $\mu = |B \cap W|/|W|$ to accuracy $\beta/16$. Retain a flat passing the containment test only if this estimate is at least $\beta/2$. A useful proposal with $A \subseteq W$ and $|W| \le Qa$ has $\mu \ge (1-\eta)\beta > 2\beta/3$, so it passes. Every retained flat has $\mu \ge 7\beta/16$; since it contains A, its latent density $p = a/|W|$ obeys

$$p \ge \frac{\mu}{1+\eta} > \frac{1}{4Q} =: p_{\min}.$$

This test handles oversized nonserving entries in the input list. Allocate a total failure budget $\rho/4$ to all containment and density estimates over the at most mR proposals. Take the first flat passing both tests; if there is none, return failure.

A relative Fourier threshold recovers all affine equations. Parameterize the retained flat by an affine bijection $\psi : \mathbb{F}_2^r \to W$, and let $A' = \psi^{-1}(A) = b + H'$ and $f = \mathbf{1}_B \circ \psi$. Then

$$\widehat{\mathbf{1}_{A'}}(\xi) = \begin{cases} p(-1)^{\langle \xi, b\rangle}, & \xi \in (H')^{\perp}, \\ 0, & \xi \notin (H')^{\perp}. \end{cases}$$

The support discrepancy gives $|\widehat{f}(\xi) - \widehat{\mathbf{1}_{A'}}(\xi)| \le \eta p$ for every ξ, while $(1-\eta)p \le \mu \le (1+\eta)p$. Consequently

$$\begin{aligned} \xi \in (H')^{\perp} &\implies |\widehat{f}(\xi)| - \mu/2 \ge (1-3\eta)p/2, \\ \xi \notin (H')^{\perp} &\implies \mu/2 - |\widehat{f}(\xi)| \ge (1-3\eta)p/2. \end{aligned} \tag{154}$$

Run Lemma 41 with threshold $\tau = 1/(8Q)$ and error budget $\rho/4$. Every frequency in $(H')^{\perp}$ has magnitude at least $(1-\eta)p > 1/(6Q) > \tau$, so all are listed. Add the zero frequency explicitly. On a fresh block, estimate μ and all listed coefficients to accuracy $e = (1-3\eta)/(64Q)$, with total error budget $\rho/4$. Retain exactly those whose estimated absolute coefficient is greater than half the estimated μ. The gap in equation 154 is at least $8e$, whereas estimation moves this comparison by at most $3e/2$. The retained set is therefore exactly $(H')^{\perp}$. For a retained nonzero coefficient, the latent coefficient has magnitude p and the support discrepancy perturbs it by at most $\eta p < p$; hence the observed coefficient already has the same sign as $\widehat{\mathbf{1}_{A'}}(\xi)$. Its magnitude is at least $(1-\eta)p > 1/(6Q)$, whereas the final estimation error is at most $e < 1/(64Q)$, so the empirical coefficient has that same sign as well. Choose a basis $\xi_1, \dots, \xi_c$ of this set and solve

$$\langle \xi_j, z\rangle = b_j, \qquad (-1)^{b_j} = \operatorname{sign}\big(\widetilde{\widehat{f}}(\xi_j)\big),$$

where the tilde denotes the empirical coefficient estimate. The solution affine space is exactly A', giving $A = \psi(A')$. Inconsistent systems on unsuccessful executions are discarded.

The four failure contributions are useful-proposal absence, proposal-test errors, Fourier-list failure, and coefficient-estimation errors, each at most $\rho/4$. The containment and density tests use fixed error margins; Fourier search uses the fixed lower density $p_{\min}$ rather than an unverified estimated reciprocal. All stages therefore have the displayed polynomial worst-case budgets on every transcript. Conditioning on an earlier fixed list changes none of these arguments. □

Corollary 43 (Exact recovery for every fixed radius below $1/3$)**.** *For $K = 1$, $0 \le \eta < 1/3$, and nonempty $\mathcal{I}_{1,\eta}(B)$, exact recovery of its unique affine support is possible from* $\mathsf{SampMQ}(B)$ *with probability $1-\rho$, using*

$$\operatorname{poly}\big(n, (1-3\eta)^{-1}, \log(1/\rho)\big)$$

samples, queries, and time.

Proof. Invoke the arbitrary-corruption clause of Theorem 3 at fixed parameters $K = 1$ and radius $1/3$, with error budget $\rho/2$. The true A is compatible at that larger radius. At these fixed parameters the returned common subspace has an absolute serving-quality bound, and the XP resource bound is polynomial with an absolute exponent. Feed the singleton list containing this subspace to Theorem 42, using the supplied radius $\eta < 1/3$ and remaining error budget $\rho/2$. □

J.1 THE CRITICAL RADIUS: UNIQUENESS WITH EXPONENTIAL ACCESS COST

Theorem 44 (Unconditional critical-radius access lower bound)**.** *At $K = 1$ and $\eta = 1/3$, every nonempty identification set is a singleton, but a randomized classical algorithm that exactly recovers its member with probability at least $2/3$ on every instance requires $\Omega(2^n)$ total uniform-sample and membership-query calls in the worst case. This bound holds without any restriction on internal computation. Enumerating B with 2^n membership queries suffices with unbounded subsequent computation.*

Proof. Uniqueness is Lemma 40. Let $a = 2^{2t} = 3r + 1$, choose nested linear subspaces $U \subset H$ of sizes $a, 2a$, and fix any $T \subset H \setminus U$ of size $r = (a-1)/3$. Put

$$B_0 = U \cup T, \qquad Z = H \setminus B_0, \qquad B_z = B_0 \cup \{z\} \quad (z \in Z).$$

For B_0, the unique latent support is U, because $|B_0 \triangle U| = r \le a/3$. For each B_z, it is H, because $|B_z \triangle H| = a - r - 1 \le 2a/3$.

Even disclose U, H, T to the learner. Compare the B_0 experiment with the experiment that first chooses z uniformly in Z and then provides access to B_z. Suppose there are at most q membership queries and s sample requests. Couple internal randomness and all replies until either a membership query hits z or a sample returns z. Under the B_0 transcript, queried points are independent of the initially uniform hidden z, so the former event has probability at most $q/|Z|$. A sample from B_z equals z with probability $1/(|B_0| + 1)$ and otherwise is uniform in B_0, so the latter event has probability at most $s/(|B_0| + 1)$. Consequently the transcript total variation distance is at most

$$\frac{q}{|Z|} + \frac{s}{|B_0| + 1}.$$

Exact recovery with probability $2/3$ on both experiments distinguishes them with a probability gap at least $1/3$. Hence this displayed quantity is at least $1/3$. Since $|Z| = (2a + 1)/3$ and $|B_0| + 1 = (4a + 2)/3$, the total access budget is $\Omega(a)$. For $n = 2t + 1$, $a = 2^{n-1}$. Padding this known construction by at most one zero coordinate gives the same $\Omega(2^n)$ worst-case bound in the remaining dimensions. The argument applies to adaptive requests by using the hypothetical B_0 execution to define the queried set and padding to deterministic call caps. A cap on the total number of calls also gives such individual caps and the same asymptotic lower bound.

For the upper bound, query every ambient point to determine B, enumerate affine subspaces, and select the unique compatible one. Computation is unrestricted in this upper bound. □

Remark 45 (Role of latent cardinality). The critical construction uses different latent sizes a and $2a$. It applies to the present unknown-cardinality model, and does not prove the same lower bound when the true cardinality is supplied. This restriction differs from the LPN reduction, which remains valid even when latent cardinality is known.

Theorem 46 (Non-identifiability above $1/3$)**.** *For every $\eta > 1/3$, there are two different doubling-one supports compatible with the same observation at radius η. Thus exact recovery cannot succeed with probability greater than $1/2$ on every instance, even with full knowledge of B.*

Proof. Take nested $U \subset H$ of sizes $a, 2a$, with a a sufficiently large power of two. Choose an integer j nearest to $a/3$, a set $T \subset H \setminus U$ of size j, and $B = U \cup T$. The two relative discrepancies are j/a and $(a - j)/(2a)$, both tending to $1/3$ as $a \to \infty$. For large enough a they are below the supplied η. The two instances give identical observations but require different exact outputs, so their two success probabilities sum to at most one. □

Proof of Proposition 6. Item 1 is Corollary 43. Item 2 is Theorem 44, with uniqueness supplied by Lemma 40. Item 3 is Theorem 46. At the fixed critical parameters, Theorem 3 has constant structural parameters and therefore polynomial running time, giving the coarse common-output contrast stated in the proposition. □

Exact versus coarse recovery at the same critical radius. At $K = 1$ and $\eta = 1/3$, Theorem 3 computes one common bounded-cost subspace in polynomial time, whereas Theorem 44 requires exponentially many accesses for exact recovery when latent cardinality is unknown. This contrast does not rely on the LPN assumption.

J.2 EXACT-SUPPORT LISTS AND THE VALUE OF LATENT CARDINALITY

Lemma 47 (At most two affine explanations below one half)**.** *For $0 \le \eta < 1/2$, the class $\mathcal{I}_{1,\eta}(B)$ contains at most two members, and at most one member of any given cardinality.*

Proof. Two different affine subspaces of the same size a intersect in at most $a/2$ points. Their symmetric difference is at least a, whereas compatibility with B would bound it by $2\eta a < a$. Every compatible size lies in

$$\frac{|B|}{1+\eta} \le |A| \le \frac{|B|}{1-\eta}.$$

The ratio of the endpoints is less than 3. As affine sizes are powers of two, this interval contains at most two possible sizes. Combining the two observations proves the assertion. □

Theorem 48 (Polynomial-time exact-support lists of size two)**.** *Let $0 \le \eta < 1/2$, $\delta = 1/2 - \eta$, and $\bar{\eta} = \eta + \delta/2$. Given $\mathsf{SampMQ}(B)$ with $\mathcal{I}_{1,\eta}(B) \neq \varnothing$, one can in*

$$\mathrm{poly}\big(n, \delta^{-1}, \log(1/\rho)\big)$$

time and oracle calls return a list $\mathcal{C}$ of at most two explicitly represented affine subspaces such that, with probability $1-\rho$,

$$\mathcal{I}_{1,\eta}(B) \subseteq \mathcal{C} \subseteq \mathcal{I}_{1,\bar{\eta}}(B).$$

The algorithm has deterministic finite budgets, returning failure on exceptional executions if necessary. The polynomial has absolute exponents, although they can depend on the absolute quality constant from the fixed-parameter invocation of Theorem 3.

Proof. Invoke the arbitrary-corruption clause of Theorem 3 at $K = 1$, radius $1/2$, and failure budget $\rho/4$. At these fixed parameters it returns one common subspace V_0 with an absolute integer serving-quality bound Q_0 for every support in $\mathcal{I}_{1,1/2}(B)$, and its XP resource bound is polynomial with an absolute exponent. Treat $\{V_0\}$ as a singleton list. Reject a failed invocation before any downstream work, so the prescribed slot bounds apply on every execution. All subsequent randomness is fresh, conditional on this fixed subspace.

Propose a containing flat for every true explanation. Put $\ell = \lfloor \log_2 Q_0 \rfloor$ and

$$p_\star = (1+\eta)^{-(\ell+1)}\delta^\ell, \qquad R = \left\lceil p_\star^{-1} \ln \frac{8}{\rho} \right\rceil.$$

For the common entry V_0, make R proposals from $\ell + 1$ fresh samples:

$$W = X_0 + \mathrm{span}(V_0, X_1 - X_0, \ldots, X_\ell - X_0).$$

For a fixed true A, the common V_0 serves it, and all these samples are in $C = A \cap B$ with probability at least $(1+\eta)^{-(\ell+1)}$. Conditional on this event they are independent and uniform in C. Any proper affine subspace of the image of A modulo V_0 contains at most $|A|/2$ preimage points. Until that image is spanned, a fresh clean sample increases the affine dimension with probability at least $1/2 - \eta = \delta$. Its dimension is at most ℓ, so the chance of spanning it with these samples is at least δ^ℓ. Consequently a proposal equals $A + V_0$, contains A, and has size at most $Q_0|A|$ with probability at least $p_\star$. By Lemma 47, there are at most two true explanations. The chance of missing a useful proposal for either is at most $2\exp(-Rp_\star) \le \rho/4$.

Guess the latent size and decode within each flat. For every proposed W, try all $a = 2^d$, $0 \le d \le n$, satisfying $a \le |W| \le Q_0 a$. Here a is a guess, not supplied side information. Let $p = a/|W| \ge 1/Q_0$, and parameterize W by an affine bijection ψ. Use the heavy-Fourier primitive on $f = \mathbf{1}_B \circ \psi$ with threshold $1/(4Q_0)$, and estimate its listed coefficients to accuracy $\delta/(4Q_0)$. Retain estimated magnitudes above $p/2$. Take a basis of the retained frequencies, solve the affine equations specified by their estimated signs, and check all retained equations. Discard inconsistent outputs and outputs whose cardinality differs from the guess a. Map any remaining affine solution back to the ambient space. This produces at most one candidate per flat and size guess.

For a useful flat and the true guess, writing $\psi^{-1}(A) = b + H'$, one has

$$\xi \in (H')^\perp \implies |\widehat{f}(\xi)| \ge (1-\eta)p, \qquad \xi \notin (H')^\perp \implies |\widehat{f}(\xi)| \le \eta p.$$

The threshold $p/2$ has margin at least $\delta p \geq \delta/Q_0$, and every required frequency has magnitude at least $p/2 > 1/(4Q_0)$. Thus all required frequencies are listed and correctly selected, their signs are correct, and this branch produces exactly A. Allocate total failure probability $\rho/4$ to all Fourier calls and coefficient estimates over their prescribed finite slots. Incorrect size guesses and noncontaining flats need not have any Fourier guarantee beyond producing an explicit candidate of the guessed size.

Validate with a small radius slack. For each produced candidate D, independently estimate

$$u_D = \mathbb{P}_{U_D}[X \in B], \qquad v_D = \mathbb{P}_{U_B}[X \in D]$$

to additive accuracy $e = \delta/128$. The first estimate uses membership queries to B and uniform points of the explicit affine space D; the second uses samples from B. Allocate total error probability $\rho/4$ across these estimates. Reject if $\widehat{v}_D < 1/2$. Otherwise compute

$$\widehat{r}_D = 1 + \widehat{u}_D/\widehat{v}_D - 2\widehat{u}_D,$$

and accept only if $\widehat{r}_D \leq \eta + \delta/4$. When $v_D > 0$, the exact identity is

$$\frac{|D \triangle B|}{|D|} = 1 + \frac{u_D}{v_D} - 2u_D.$$

On the simultaneous estimation event, the gate guarantees $v_D \geq 1/2 - e$. Since $\delta \leq 1/2$, elementary bounds on the ratio give

$$\left| \widehat{r}_D - \frac{|D \triangle B|}{|D|} \right| \leq 10e < \delta/8.$$

Every true A passes: its $v_A \geq 1/(1+\eta) > 2/3$, and its discrepancy ratio is at most η. Every accepted candidate has discrepancy ratio at most $\eta + 3\delta/8 < \bar{\eta} < 1/2$. Deduplicate the accepted affine subspaces. Lemma 47 at radius $\bar{\eta}$ bounds their number by two. If more than two survive on an exceptional execution, return failure rather than an oversized list. The four error allocations sum to ρ.

The constants Q_0, ℓ are absolute. Therefore $R = O(\delta^{-\ell} \log(1/\rho))$, all Fourier thresholds are constant, and all estimation margins are inverse polynomial in δ^{-1}. The number of size guesses is at most $n + 1$. Together with the original fixed-parameter invocation, this proves the claimed worst-case budgets. □

Corollary 49 (Known cardinality moves exact recovery to one half)**.** *Under the same structural algorithm input, if the true cardinality $a = |A|$ is supplied, then exact recovery is possible for every $0 \leq \eta < 1/2$ in $\mathrm{poly}(n, (1-2\eta)^{-1}, \log(1/\rho))$ time and oracle calls. Restrict the algorithm's guesses and outputs to size a; Lemma 47 ensures that every accepted candidate, including those valid only at the relaxed radius, equals the true A.*

Theorem 50 (Exponential exact-support ambiguity at one half)**.** *At $K = 1$, $\eta = 1/2$, even when the latent cardinality is known to be two, one observation has $2^n - 1$ compatible latent supports. Hence a list covering all explanations must have exponential length. More generally, a randomized list of length at most L that includes each possible true support with probability at least $1 - \rho$ must satisfy $L \geq (1-\rho)(2^n - 1)$.*

Proof. Take $B = \{0\}$. For every nonzero $v \in \mathbb{F}_2^n$, the linear subspace $A_v = \{0, v\}$ has size two and discrepancy $|A_v \triangle B| = 1$. These $2^n - 1$ supports are distinct. Summing the inclusion probabilities gives the randomized-list lower bound. The same construction applies at any larger radius. □

What the two-output result says at the critical point. At $\eta = 1/3$ there is a unique exactly compatible support, yet the algorithm can return two candidates because its validity radius is slightly larger. It includes the exact support but need not identify which entry has discrepancy at most the original $1/3$ budget. This is consistent with Theorem 44: constructing two exact-support candidates is polynomial time, whereas selecting the unique correct entry can require exponential oracle access. For any fixed $\eta \in (1/3, 1/2)$, the nested construction from Theorem 46 also shows that two candidates are sometimes information-theoretically necessary.

K AMBIGUITY BEYOND LATENT CARDINALITY

The sharp example in Theorem 1 uses the compatible supports $\{0\}$ and G, whose cardinalities are very different. This appendix shows that the linear corruption cost is not merely a consequence of uncertainty about latent size. A prescribed family of equally large subspaces can force the same scale through independent unseen directions. The qualifier "prescribed family" is important: for the observation below, B itself also belongs to the full identification set and can impose a smaller admissible subspace size than the equal-size family alone.

Theorem 51 (Equal-size directional ambiguity)**.** *Let*

$$G = U \oplus X_1 \oplus \cdots \oplus X_r$$

be an internal direct sum over $\mathbb{F}_2$, where $d := \dim U$, $\dim X_i = t \geq 1$, and $r \geq 1$. Define

$$H_i := U \oplus X_i, \qquad B := U.$$

Then every H_i has doubling one, all H_i have the same cardinality, and

$$|H_i \triangle B| = (1 - 2^{-t})|H_i|.$$

Moreover,

$$\min_{\substack{V \leq G \\ \dim V \leq d+t}} \max_{1 \leq i \leq r} \mathcal{N}_V(H_i) = 2^{\,t - \lfloor t/r \rfloor}. \tag{155}$$

In particular, if $r > t$, then the optimum equals

$$2^t = (1 - \eta)^{-1}, \qquad \eta := 1 - 2^{-t},$$

even though all prescribed latent explanations have identical size.

Proof. The first assertions are immediate: each H_i is a subspace of cardinality 2^{d+t}, while $B = U \subseteq H_i$ has cardinality 2^d, so

$$\frac{|H_i \triangle B|}{|H_i|} = 1 - \frac{|U|}{|H_i|} = 1 - 2^{-t}.$$

Fix a subspace $V \leq G$ with $\dim V \leq d + t$. Put

$$J_i := V \cap H_i, \qquad u := \dim(V \cap U).$$

The quotients $J_i/(V \cap U)$ form an internal direct sum. Indeed, if $x_i \in J_i$ and $\sum_i x_i \in U$, projection onto the independent coordinate X_i forces the X_i-component of each x_i to vanish, so every $x_i \in U$. Consequently,

$$\sum_{i=1}^{r} \dim J_i - (r-1)u = \dim\left(\sum_{i=1}^{r} J_i\right) \leq \dim V \leq d + t. \tag{156}$$

Set

$$c_i := d + t - \dim J_i.$$

Then equation 156 gives

$$\begin{aligned} \sum_{i=1}^{r} c_i &= r(d+t) - \sum_i \dim J_i \\ &\geq (r-1)(d+t-u) \geq (r-1)t. \end{aligned}$$

Hence

$$\max_i c_i \geq \left\lceil \frac{(r-1)t}{r} \right\rceil = t - \left\lfloor \frac{t}{r} \right\rfloor.$$

Because H_i is a subspace,

$$\mathcal{N}_V(H_i) = |\pi_V(H_i)| = 2^{\dim H_i - \dim(V \cap H_i)} = 2^{c_i}.$$

This proves the lower bound in equation 155.

For the matching construction, choose integers $s_1, \dots, s_r$ summing to t, each equal to either $\lfloor t/r \rfloor$ or $\lceil t/r \rceil$. Choose subspaces $X_i' \le X_i$ with $\dim X_i' = s_i$, and put

$$V := U \oplus \bigoplus_{i=1}^{r} X_i'.$$

Then $\dim V = d + t$, while

$$V \cap H_i = U \oplus X_i', \qquad \mathcal{N}_V(H_i) = 2^{t - s_i}.$$

Therefore

$$\max_i \mathcal{N}_V(H_i) = 2^{t - \min_i s_i} = 2^{\,t - \lfloor t/r \rfloor},$$

which proves equality. □

Theorem 51 isolates a different source of ambiguity from the two-scale example used for the exact common radius. Here latent cardinality is fixed; what remains hidden is how the support extends along mutually independent quotient directions. When $r > t$, no admissible $(d+t)$-dimensional common subspace can absorb a positive-dimensional piece of every direction, and the full quotient cost 2^t is forced on at least one member of the family.

L A One-Core Size-Only Mechanism Barrier

The moderate-corruption learner certifies clean mass using the observed size of one extracted core. We show that the $K^{-1/2}$ scale is intrinsic to this size-only proof architecture. The bound concerns that mechanism; other recovery algorithms may use additional information.

The mechanism has two numerical ingredients. First, at normalized energy $1/K$, one cannot universally force a bounded-small-doubling core larger than order $K^{-1/2}$ of the observation. Second, if clean mass is inferred only from a core's retained fraction, the elementary purity inequality used in Appendix D is sharp.

Proposition 52 (Sharpness of size-only purity)**.** *Fix $0 < \gamma \le 1/2$ and $\eta \in [0, 1)$. For arbitrarily large finite sets A, there exist B and $Y \subseteq B$ such that*

$$|A \triangle B| \le \eta |A|, \qquad |Y| \ge \gamma |B|,$$

and

$$\frac{|A \cap Y|}{|A|} = \big[\gamma - (1-\gamma)\eta\big]_+ + o_{|A| \to \infty}(1). \tag{157}$$

Consequently, size information of this form alone can force a positive universal clean fraction only when

$$\eta < \frac{\gamma}{1-\gamma}. \tag{158}$$

Proof. Let $|A| = N$. Choose a set D disjoint from A with $|D| = \lfloor \eta N \rfloor$, and put $B := A \cup D$. Thus all corruption consists of insertions and $|A \triangle B| \le \eta |A|$. Choose $Y \subseteq B$ of cardinality $\lceil \gamma |B| \rceil$, taking points from D before points from A. Then

$$|A \cap Y| = \big[\lceil \gamma |B| \rceil - |D|\big]_+.$$

Since $|B| = (1 + \eta + o(1))N$, division by N gives

$$\frac{|A \cap Y|}{N} = \big[\gamma(1+\eta) - \eta\big]_+ + o(1) = \big[\gamma - (1-\gamma)\eta\big]_+ + o(1).$$

The bracket is positive exactly under equation 158. □

The following construction is a disjoint affine-block variant of the coordinate-axis example of Reiher & Schoen (2024). Combining it with the sharp purity bound gives the size-only lifting barrier below.

Lemma 53 (Affine-block core-size obstruction)**.** *Let $r \geq 2$, let $L \geq 1$, and let $\delta > 0$. There is a finite nonempty set $S \subseteq \mathbb{F}_2^n$ such that*

$$E(S) \geq r^{-2}|S|^3, \tag{159}$$

and every nonempty $Y \subseteq S$ satisfying

$$|Y + Y| \leq L|Y|$$

obeys

$$|Y| \leq \frac{1+\delta}{r}|S|. \tag{160}$$

At the same time, S contains a subset of size $|S|/r$ with doubling one.

Proof. Choose a power of two m so large that $m \geq 2L/\delta$. Let

$$G = E \oplus W_1 \oplus \cdots \oplus W_r,$$

where $\dim E = r$, $|W_i| = m$, and $e_1, \ldots, e_r$ is a basis of E. Define pairwise disjoint affine blocks

$$S_i := e_i + W_i, \qquad S := \bigcup_{i=1}^{r} S_i.$$

Then $|S| = rm$, $S_i + S_i = W_i$, and for $i \neq j$,

$$S_i + S_j = e_i + e_j + W_i + W_j$$

has cardinality m^2. The cross-block sumsets belonging to distinct unordered pairs $\{i, j\}$ are disjoint.

We first compute enough energy. At zero, the representation count is $r_S(0) = rm$. For every nonzero $d \in W_i$, there are exactly m ordered representations inside S_i, and for every $d \in S_i + S_j$, $i < j$, there are exactly two ordered cross-block representations. Hence

$$\begin{aligned} E(S) &= (rm)^2 + r(m-1)m^2 + 4\binom{r}{2}m^2 \\ &= rm^3 + 3r(r-1)m^2 \geq rm^3 = r^{-2}|S|^3, \end{aligned}$$

which proves equation 159.

Now let $Y \subseteq S$, write

$$a_i := |Y \cap S_i|, \qquad M := |Y| = \sum_i a_i.$$

For $i < j$, direct-sum uniqueness makes addition injective on $(Y \cap S_i) \times (Y \cap S_j)$, and the resulting cross-block sumsets are disjoint across unordered pairs. Therefore

$$\begin{aligned} |Y + Y| &\geq \sum_{i<j} a_i a_j = \frac{1}{2}\left(M^2 - \sum_i a_i^2\right) \\ &\geq \frac{1}{2}M(M-m), \end{aligned}$$

because $a_i \leq m$ implies $\sum_i a_i^2 \leq mM$. If $|Y + Y| \leq LM$ and $M > 0$, then $M \leq m + 2L$. Thus

$$\frac{|Y|}{|S|} \leq \frac{1}{r}\left(1 + \frac{2L}{m}\right) \leq \frac{1+\delta}{r},$$

which is equation 160. Finally, every full block S_i has size $|S|/r$ and doubling one. □

Theorem 54 (One-core size-only mechanism barrier)**.** *Let $L : [1, \infty) \to [1, \infty)$ be any prescribed small-doubling budget. For every integer $r \geq 2$, every $\delta > 0$, and $K := r^2$, there is a nonempty set $S \subseteq \mathbb{F}_2^n$ with*

$$E(S) \geq K^{-1}|S|^3$$

such that every nonempty $Y \subseteq S$ satisfying

$$|Y + Y| \leq L(K)|Y|$$

has

$$|Y| \leq (1+\delta)K^{-1/2}|S|. \tag{161}$$

Consequently, consider any proof architecture that

1. *extracts one bounded-small-doubling core from an energy-* $1/K$ *observation using only a universal retained-fraction guarantee* $|Y| \geq \gamma(K)|B|$*; and*

2. *certifies clean latent mass in that same preselected core using only* $Y \subseteq B$*, its relative size, and* $|A \triangle B| \leq \eta|A|$.

Then along $K = r^2$ *one necessarily has*

$$\gamma(K) \leq K^{-1/2}, \tag{162}$$

and this size-only route can certify a positive universal clean fraction only below

$$\eta < \frac{1}{\sqrt{K} - 1}. \tag{163}$$

Proof. Apply Lemma 53 with $L = L(K)$. Because $\delta > 0$ is arbitrary, no universally valid retained fraction can exceed $1/r = K^{-1/2}$; this proves equation 162.

For $r \geq 2$, $\gamma(K) \leq 1/r \leq 1/2$. By Proposition 52, a positive clean-mass guarantee based only on retained size requires

$$\eta < \frac{\gamma(K)}{1 - \gamma(K)}.$$

Since $x/(1-x)$ is increasing on $[0, 1)$,

$$\frac{\gamma(K)}{1 - \gamma(K)} \leq \frac{1/r}{1 - 1/r} = \frac{1}{r - 1} = \frac{1}{\sqrt{K} - 1},$$

which proves equation 163. □

The arbitrary-corruption learner instead peels disjoint cores and uses their total captured clean mass. Averaging then finds a serving core for each support, without requiring one preselected core to serve every support on the basis of its size alone.

M Detailed Relation to Prior Work

M.1 Latent identifiability and common structural targets

Latent-variable learning asks which features of an unobserved object can be determined from observations. Nonlinear ICA and related models use auxiliary structure to obtain identifiability (Hyvarinen et al., 2019; Khemakhem et al., 2020); impossibility results for unsupervised disentanglement emphasize the role of inductive bias (Locatello et al., 2019). Under partial observation, Yao et al. (Yao et al., 2024) and Xu et al. (Xu et al., 2024) study shared or sparse latent information. Syrota et al. (Syrota et al., 2025) study relations that remain identifiable when latent coordinates do not.

Here the latent object is a finite support, and the target is a subspace whose covering guarantee holds across all compatible supports. This is a validity requirement over an identification set; several subspaces may meet it. The exact-recovery contrast in Appendix J also applies when that set is a singleton: a common coarse cover can be computed with polynomially many oracle calls while exact reconstruction requires exponentially many.

M.2 Robust, list-decodable, and finite-field subspace learning

List-decodable learning permits a small inlier fraction and returns a short list of hypotheses (Charikar et al., 2017). This approach has led to robust mean and subspace recovery guarantees (Diakonikolas et al., 2018; Raghavendra & Yau, 2020; Bakshi & Kothari, 2021). Our shared insertion–deletion budget makes a strict majority of B genuine under every compatible explanation. This property supports a stronger final requirement: one subspace with a uniform covering guarantee. The query learner uses a list to generate candidates, then selects a common output by fresh samples.

Chen, De, and Vijayaraghavan (Chen et al., 2021) study mixtures of two finite-field subspaces and identify LPN in the nested case. Our reduction addresses quotient covers, for which one parity graph

has a secret-independent two-coset solution. Blocking makes a coarse output reveal a short secret list. The frozen-table coupling then realizes one fixed observed support and separately controls its symmetric-difference budget. The reduction allows polynomial-size output lists and supplied latent cardinality.

M.3 Query access and persistent noise

Query learning treats access beyond passive examples as a computational resource (Angluin, 1988). Local membership queries restrict that resource (Awasthi et al., 2013). Noisy parity learning (Blum et al., 2003), persistent-noise query models (Jackson et al., 1999), and parity learning with queries (Feldman, 2007) provide relevant precedents. Feldman and Shah (Feldman & Shah, 2009) obtain conditional comparisons between faulty-teacher access models using noisy-parity hardness.

We compare samples and membership queries for the same quotient-covering loss on a fixed corrupted support. The statistical theorem establishes that samples contain enough information; the matched comparison in Section 2 isolates the computational contribution of queries under the stated LPN assumption. Persistent-noise analyses already use fixed random functions and collision couplings. Here the table is indexed by complete blocked inputs, and the support promise is checked separately without conditioning the simulator on its validity.

De, Nadimpalli, and Servedio (De et al., 2022) study indicator-oracle access to additive sets and approximate a relaxed sumset-size functional. Our algorithms construct structural hypotheses and reusable implicit cores; the cores must support adaptive membership evaluation and exact sampling from the same realized set.

M.4 Algorithmic additive structure

The polynomial Freiman–Ruzsa theorem supplies the clean covering guarantee (Gowers et al., 2025). Recent algorithms make clean PFR constructive (Arunachalam et al., 2026; Castro-Silva et al., 2026). Appendix C adapts the restricted-homomorphism primitive of Castro–Silva et al. to an input of unknown size, deterministic resource caps, and conditional invocation after an adaptive past.

Reiher & Schoen (2024) establish the essentially optimal $\Theta(\sqrt{\alpha})$ retained-mass scale through high-/low-multiplicity arguments and relation cleaning. Tulsiani & Wolf (2014) give randomized access procedures for implicit BSG-related sets, with set-sandwich guarantees controlling testing noise. For explicit sets, algorithmic BSG appears in structured 3SUM (Chan & Lewenstein, 2015). Song (2026) reaches the optimal output-size scale for integer sets using weighted shift sampling, a shared codegree classifier, and sampled relation cleaning.

Our construction shares these structural ingredients. It works through sample-and-query access to a potentially exponentially large support and returns a finite record defining one fixed subset. Its guarantee combines the sharp retained-mass scale, a bound on invalid nonfailure outputs without an energy promise, high-energy completeness, and exact product sampling conditional on finite-block success. These are the properties used by residual peeling and the subsequent PFR calls. Appendix B proves them.

M.5 Exact affine recovery

The exact decoder uses Goldreich–Levin heavy-Fourier search (Goldreich & Levin, 1989). Related coding and learning results include local list decoding of binary Reed–Muller codes (Gopalan et al., 2008), sample learning of Fourier-sparse Boolean functions (Haviv & Regev, 2015), and property testing of affine-subspace indicators (Bshouty, 2024). After localization, the affine-support indicator has degree equal to its local codimension, which is bounded by the coarse covering quality.

The support-relative corruption model requires several steps before the Fourier primitive can be applied: finding a containing flat, certifying its latent density, and setting an observable coefficient threshold without knowing the latent size. Candidate validation handles that unknown size. The critical $1/3$ result measures the access cost of exact reconstruction in this model. Its lower-bound pair has different latent cardinalities; the known-cardinality refinement extends exact recovery below $1/2$. These thresholds use affine rigidity at $K = 1$. For general small doubling, the main target remains a common structural cover.

N Algorithmic and Access-Model Questions

The statistical and query learners reach the same type of common structural output through different candidate families. This leaves two quantitative questions: how much covering quality is lost by constructing candidates efficiently, and how much computation is needed to represent those candidates. A third question concerns which forms of access can support the construction.

N.1 Matching the statistical covering quality

The statistical learner achieves cost $2Kp(K)/\varepsilon$ by searching over clean PFR descriptions. The general query learner constructs candidates through structured-core extraction and then applies the common-output filter. Its larger bound reflects both the quality of the core covers and the transfer from a serving core to the latent support.

These losses suggest two possible routes to improvement. One is to strengthen the candidate-generation guarantee, so that each compatible support has a serving entry with smaller quotient cost. Another is to use more of the relationship between the candidates when selecting a common output. The present filter uses their individual coverage scores and dimensions; it does not combine their algebraic directions.

The resulting question is whether query access can attain the statistical covering bound with comparable sample complexity. This question can be studied separately from the running-time dependence: better covering quality would already strengthen the result even if the algorithm remained XP.

N.2 Representing residuals without recursive expansion

The current XP exponent comes from implementing the successive residuals $R_j = B \setminus \bigcup_{i<j} Y_i$. The number of peeling rounds depends only on the structural parameters, but evaluating a core predicate invokes the residual oracle from its construction stage. Expanding these calls through earlier stages produces the parameter-dependent polynomial degree.

An FPT implementation would need a representation whose evaluation cost has a parameter-independent degree in n. Possible approaches include sharing intermediate predicate evaluations, compiling several peeling stages jointly, or constructing a candidate family without sequential residual evaluation. For an approach based on reuse, the central quantity is the number of distinct derived query points across the adaptive computation.

Any such representation must preserve the two properties used downstream: membership refers to one fixed realized core, and sampling supplies that same core to the PFR routine. Improving this interface could therefore reduce the computational cost while leaving the outer covering and selection arguments intact.

N.3 Learning with weaker access

The query construction tests additive relations at points formed from previous samples, such as $u + v + w$. This suggests studying access restrictions in terms of which constructed points can be tested. A restricted-query learner would need to generate enough useful additive relations within its allowed query set to support candidate extraction.

Nonuniform sampling changes a different part of the argument. The shared corruption budget controls cardinalities, and exact uniform sampling turns those bounds into probability-mass guarantees. For a general sampling distribution on B, this correspondence needs an additional assumption. A natural starting point is a distribution whose density relative to U_B is bounded above and below, with the goal of tracking how the selection thresholds and core guarantees depend on those bounds.

Approximate membership also interacts with the sampling interface: accepting points through an inaccurate predicate can change the distribution supplied to the next stage. A useful extension would specify the oracle error model and control both the structural validity of the returned core and the distributional error in its downstream samples. These questions connect weaker access directly to the two mechanisms used by the learner: candidate generation and common-output selection.