Weak L^3-control versus full control in additive combinatorics

Liz Lemma   Future Detective

January 16, 2026

Abstract

Bloom introduced (and renamed) an L^3-uniform convolution bound ("control") that propagates to quantitative improvements in problems such as sum-product, additive growth of convex sets, and Balog–Szemerédi–Gowers (BSG). Bloom also notes that many threshold-breaking arguments only require a weaker hypothesis—roughly, an L^3 bound on $\mathbf{1}_A \circ \mathbf{1}_A$ together with L^2 bounds on $\mathbf{1}_A \circ \mathbf{1}_B$—but no meaningful examples are known that separate this weak notion from full control. We formalise weak control via two parameters: a uniform L^2 cross-correlation bound and a self L^3 correlation bound. Our main theorem upgrades these weak assumptions to full L^3-control (up to polylogarithmic loss), yielding an equivalence of invariants and allowing Bloom's control-based machinery to be invoked under simpler, more checkable hypotheses. We also discuss candidate separating constructions in high-rank finite vector spaces and outline computational tests that could certify separation if equivalence fails.

Table of Contents

1. 1. Introduction: Bloom control, weak control, and why equivalence matters; statement of main results and corollaries for exporting control-based bounds.

2. 2. Preliminaries: convolution and difference-convolution conventions; representation functions; basic inequalities (Hölder, Cauchy–Schwarz, dyadic decomposition) and polylog notation.

3. 3. Weak control parameters $K_2(A), K_3(A)$: definition, monotonicity, scaling, and immediate relations to Bloom control; comparison with Shakan/Shkredov parameters.

4. 4. From set tests to function tests: layer-cake upgrades of (i) to weighted statements; Lorentz-space viewpoint; removing large tail losses (no κ^{100}-type terms).

5. 5. The upgrade theorem (weak $\Rightarrow$ full): reduction to bounding $\|\mathbf{1}_A * \mathbf{1}_B\|_3$ for arbitrary B via a multi-scale decomposition; proof of the main implication with explicit exponents.
6. 6. Equivalence up to polylogs: full $\Rightarrow$ weak (easy) + weak $\Rightarrow$ full (main theorem) gives a robust identification of the right invariant; discussion of tightness and where polylogs enter.
7. 7. Applications: simplifying hypotheses in Bloom-style propagation theorems; black-box replacements in convex-set arguments, sum-product decompositions, and BSG-from-control arguments.
8. 8. Examples and non-examples: explicit computations for random sets, arithmetic progressions, convex sets/images under convex maps; sanity checks showing K_2^4 and K_3 scale correctly.
9. 9. Separation program (if needed): explicit candidate constructions in $\mathbb{F}_p^n$; what would need computer verification; diagnostics for failure of equivalence.
10. 10. Further questions: stability/inverse problems for weak control; extension to L^4 control in finite fields; connections to incidence bounds and operator norms.

1 Introduction

A recurring theme in additive combinatorics is that quantitative bounds for a set $A \subset G$ are most robust when they are expressed through an *invariant* that controls how A correlates additively with *every* other set B. One convenient way to package such information is via higher moments of the representation function

$$r_{A+B}(x) = \big(\mathbf{1}_A * \mathbf{1}_B\big)(x),$$

whose size measures how often x is realised as a sum $a + b$. Among these, the third moment

$$\sum_{x \in G} r_{A+B}(x)^3$$

plays a special role: it is strong enough to force substantial additive structure when it is large, yet sufficiently stable under the arguments that propagate additive information (energy increment schemes, Balog–Szemerédi–Gowers type reductions, and various growth/expansion mechanisms). Bloom formalised this by introducing a *full L^3-control* parameter $\kappa(A)$: it is the least constant for which one has, uniformly for all finite $B \subset G$,

$$\sum_x \big(\mathbf{1}_A * \mathbf{1}_B(x)\big)^3 \;\le\; \kappa(A)\, |A|^2 |B|^2.$$

In practice (and throughout this paper) we allow polylogarithmic losses in $|A|$, since essentially all downstream applications tolerate such losses; thus the correct scale is $\kappa(A)$ up to factors of $\log(2|A|)^{O(1)}$.

The main point of this work is that the full L^3 control condition, while natural from Bloom's perspective, is often stronger than what one can verify directly in concrete examples. In many situations, the available input takes the form of two weaker pieces of information. The first is a uniform L^2 bound for cross-correlations

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2,$$

equivalently a uniform upper bound for the additive energy $E(A, B)$. The second is a self-correlation bound at exponent 3, namely control of $\|\mathbf{1}_A \circ \mathbf{1}_A\|_3$. These are precisely the weak parameters $K_2(A)$ and $K_3(A)$ introduced in the global context: they measure, respectively, how large $E(A, B)$ can be as B varies and how concentrated the difference representation function r_{A-A} can be at high multiplicities.

The motivating question is then the following.

To what extent is Bloom's full control $\kappa(A)$ determined (up to polylogarithmic factors) by the weaker invariants $K_2(A)$ and $K_3(A)$?

This question matters because the known consequences of full control are numerous and powerful, but the hypothesis $\kappa(A) \le \kappa$ is not always the most convenient interface with incidence geometry, Fourier-analytic bounds, or combinatorial structure theorems. By contrast, estimates for energies and for a single self-correlation moment often appear as direct outputs of such methods. For instance, an argument may naturally yield an inequality of the form

$$E(A,B) \lesssim |A|^{3/2}|B|^{3/2} \qquad \text{for all } B,$$

together with a separate bound showing that r_{A-A} has few very popular differences, which is exactly what an L^3 bound on $\mathbf{1}_A \circ \mathbf{1}_A$ quantifies. If these two inputs already imply full L^3 control, then one may immediately import the entire package of "control-based" results without re-running Bloom's framework from scratch.

Our principal result establishes precisely such an implication: the weak hypotheses (i) and (ii) from the enclosing scope upgrade to full L^3 control, with the expected dependence on parameters. The dependence is dictated by scaling. Indeed, the simplest way to see that K_2 must enter as K_2^4 is that full control bounds an L^3 quantity, and by Hölder one can pass from L^3 to L^2 with an exponent loss of 1/4; conversely, one expects that an L^2 hypothesis must be iterated (in a dyadic or energy-increment sense) to reach a third moment, producing a fourth power. Similarly, the parameter K_3 is already an L^3 self-correlation bound, so it should enter linearly. Our theorem confirms that, up to polylogarithmic factors, $\kappa(A)$ is controlled by $\max\{K_2(A)^4, K_3(A)\}$ and, conversely, full control automatically implies the weak bounds. In particular, the three invariants are polylog-equivalent in the sense that

$$\kappa(A) \asymp_{\text{polylog}} \max\{K_2(A)^4,\ K_3(A)\}.$$

We emphasise two consequences of this equivalence.

Exporting results. Any statement in the literature whose hypothesis is formulated in terms of $\kappa(A)$ can be reformulated with hypotheses on $K_2(A)$ and $K_3(A)$ only, at the cost of a polylogarithmic degradation. This is not merely cosmetic: in concrete settings one may be able to bound $K_2(A)$ and $K_3(A)$ by direct estimates (energies and popular-differences tails) while $\kappa(A)$ remains opaque. Thus the equivalence provides a systematic translation layer between older "energy-type" inputs and Bloom's control framework.

Conceptual simplification. Bloom's control parameter is, by definition, uniform over all sets B, and a priori it is not clear whether it is genuinely stronger than uniform energy bounds plus a single self-correlation constraint. The equivalence shows that, at least up to polylogarithmic losses, no additional hidden obstruction exists: the third moment of $\mathbf{1}_A * \mathbf{1}_B$ can only become large through mechanisms already detected by (i) and (ii). In particular, the "high multiplicity" obstruction is captured by the L^3 behaviour of

r_{A-A} (parameter K_3), and the "medium multiplicity" obstruction is captured by uniform energy control (parameter K_2).

At the level of proof strategy, our argument proceeds by decomposing the third moment of r_{A+B} into contributions from multiplicity ranges. The most basic decomposition is dyadic: for t ranging over powers of 2, let

$$S_t := \{x \in G : \; r_{A+B}(x) \in [t, 2t)\}.$$

Then

$$\sum_x r_{A+B}(x)^3 \;\sim\; \sum_t \sum_{x \in S_t} r_{A+B}(x)^3,$$

and it suffices to bound the contribution of each S_t in a manner summable over t. The burden is to relate information about r_{A+B} to information about r_{A-A} and to mixed energies involving A and portions of B.

A technical obstacle appears immediately: the hypothesis (i) is stated only for indicators $\mathbf{1}_B$, while a dyadic decomposition naturally introduces *weights* (for instance, one is led to consider functions measuring multiplicities of fibres). We therefore begin by upgrading (i) to a stable inequality for general nonnegative functions f:

$$\|\mathbf{1}_A \circ f\|_2 \;\ll\; \log(2|\mathrm{supp}(f)|)^{O(1)}\, K_2\, |A|^{3/4}\, \|f\|_{4/3}.$$

This is achieved by a layer-cake (dyadic) decomposition of f into level sets and an efficient summation across scales. While elementary, this step is the analytic backbone of the argument: it allows us to apply (i) in situations where B is replaced by a structured multiset extracted from B and the level set S_t.

With this functional upgrade in hand, we bound each dyadic contribution by separating into three regimes.

- In the *low multiplicity* regime (small t), trivial inequalities (Young and Hölder) already give acceptable estimates, and the contribution is summable with room to spare.
- In the *medium multiplicity* regime, we relate the size of S_t and the distribution of r_{A+B} on S_t to mixed additive energies, and then invoke the uniform energy hypothesis encoded by K_2. Here the functional form of (i) is used to control the relevant correlation terms without losing more than logarithmic factors.
- In the *high multiplicity* regime, one cannot hope to control S_t solely from energy bounds: very large values of r_{A+B} can arise from concentration on a small set of popular differences. This is precisely where

the self L^3 hypothesis (ii) enters. From $\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 \le K_3|A|^4$ we deduce a tail bound for popular differences,

$$|\{d: \ r_{A-A}(d) \ge t\}| \ \ll \ K_3 \, \frac{|A|^4}{t^3},$$

which forces high-multiplicity phenomena to be rare enough for the dyadic sum to converge.

Combining the three regimes yields a bound of the desired shape, with the dyadic summation contributing a factor $\log(2|A|)^{O(1)}$ and the controlling parameter appearing as $\max\{K_2^4, K_3\}$. The reverse implication, namely that full control implies the weak bounds, is straightforward and follows from Hölder-type inequalities and specialisation to $B = -A$. Taken together, these establish the polylog-equivalence of invariants.

Finally, we record a conceptual alternative. If the upgrade from weak to full control were false, then one should be able to construct explicit counterexamples in groups with rich subgroup geometry, such as $\mathbb{F}_p^n$. In that setting one can attempt to decouple moment conditions by taking unions of cosets (to engineer energy behaviour) together with pseudorandom perturbations (to manipulate higher moments). While we do not pursue such a construction to completion here, we outline a candidate program and a finite verification approach: search over structured witness sets B for which $\|\mathbf{1}_A * \mathbf{1}_B\|_3$ is maximised, subject to constraints enforcing (i) and (ii). The present theorem may be viewed as ruling out this separation program in full generality, up to the polylogarithmic losses inherent in the method.

In the next section we fix notation and collect basic analytic tools: convolution and difference-convolution conventions, representation functions, and the standard inequalities and dyadic decompositions used repeatedly in the proof.

2 Preliminaries

2.1 Ambient group, measures, and norms

Throughout we work in an abelian group G written additively. All sets and functions under consideration are finitely supported, and all sums are taken with respect to the counting measure on G. For $1 \le p < \infty$ and finitely supported $f\colon G \to \mathbb{C}$ we write

$$\|f\|_p := \Big(\sum_{x\in G} |f(x)|^p\Big)^{1/p}, \qquad \|f\|_\infty := \sup_{x\in G} |f(x)|.$$

If $A \subset G$ is finite we denote by $\mathbf{1}_A$ its indicator function and note $\|\mathbf{1}_A\|_p = |A|^{1/p}$ for $1 \le p < \infty$.

We will repeatedly use the convention that implicit constants are absolute. When polylogarithmic losses in $|A|$ are permitted we write $X \lesssim Y$ to mean

$$X \le C \log(2|A|)^C Y$$

for some absolute constant $C > 0$ (which may vary from line to line). In contexts where $|A|$ is not the only relevant size parameter (e.g. when a function f has support of size $|\operatorname{supp}(f)|$), we may instead record the dependence as $\log(2|\operatorname{supp}(f)|)^{O(1)}$.

2.2 Convolution and difference-convolution

For finitely supported functions $f, g\colon G \to \mathbb{C}$ we define the (additive) convolution

$$(f * g)(x) := \sum_{y \in G} f(x-y)g(y),$$

and the *difference-convolution* (or cross-correlation)

$$(f \circ g)(x) := \sum_{y \in G} f(x+y)g(y).$$

These two operations are related by the involution $\tilde{g}(y) := g(-y)$, since

$$f \circ g \;=\; f * \widetilde{g}.$$

In particular, for indicators one has

$$(\mathbf{1}_A * \mathbf{1}_B)(x) = |\{(a,b) \in A \times B : a+b = x\}|, \qquad (\mathbf{1}_A \circ \mathbf{1}_B)(x) = |\{(a,b) \in A \times B : a-b = x\}|.$$

We will freely pass between sum and difference language using this identity, and we occasionally write $-B := \{-b : b \in B\}$ so that $\mathbf{1}_{-B}(x) = \mathbf{1}_B(-x)$ and $\mathbf{1}_A \circ \mathbf{1}_B = \mathbf{1}_A * \mathbf{1}_{-B}$.

The basic algebraic properties are standard: $f * g = g * f$, $(f * g) * h = f * (g * h)$, while for $\circ$ we have $f \circ g = \widetilde{g} \circ \widetilde{f}$ and $(f \circ g)(x) = (g \circ f)(-x)$. We use these only at the level of harmless rearrangements of sums.

2.3 Representation functions and additive energy

Given finite sets $A, B \subset G$ we define the representation functions

$$r_{A+B}(x) := (\mathbf{1}_A * \mathbf{1}_B)(x), \qquad r_{A-B}(x) := (\mathbf{1}_A \circ \mathbf{1}_B)(x).$$

Thus $r_{A+B}(x)$ counts the number of representations $x = a + b$, and $r_{A-B}(x)$ counts the number of representations $x = a - b$.

The (additive) energy between A and B is

$$E(A,B) := \|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2 = \sum_{x \in G} r_{A-B}(x)^2.$$

Expanding the square shows the familiar quadruple-counting interpretation:

$$E(A,B) = |\{(a,a',b,b') \in A^2 \times B^2 :\ a - b = a' - b'\}|.$$

A second expansion yields the correlation identity

$$E(A,B) = \sum_{d \in G} r_{A-A}(d)\, r_{B-B}(d), \tag{1}$$

obtained by grouping solutions by the common difference $d = a - a' = b - b'$. We will use (1) to transfer information between $A - A$ and mixed energies involving B.

Since $\mathbf{1}_A \circ \mathbf{1}_B = \mathbf{1}_A * \mathbf{1}_{-B}$, the same energy also controls the second moment of r_{A+B}:

$$\|\mathbf{1}_A * \mathbf{1}_B\|_2^2 = E(A, -B).$$

This trivial observation is the bridge between hypotheses formulated in terms of difference-convolutions and the third moment of r_{A+B} that we ultimately seek to control.

2.4 Standard inequalities

We record the analytic inequalities we use repeatedly.

Hölder and Cauchy–Schwarz. For $1 \le p, q, r \le \infty$ with $1/r = 1/p + 1/q$ we have

$$\|fg\|_r \le \|f\|_p \|g\|_q.$$

In particular, $\sum_x |f(x)g(x)| \le \|f\|_2 \|g\|_2$ (Cauchy–Schwarz), and the monotonicity $\|f\|_p \le \|f\|_q$ holds whenever $p \ge q$ and f is supported on a set of finite size, with the usual dependence on $|\operatorname{supp}(f)|$.

Young's inequality. If $1 \le p, q, r \le \infty$ satisfy $1 + 1/r = 1/p + 1/q$, then

$$\|f * g\|_r \le \|f\|_p \|g\|_q.$$

We will most frequently use the cases

$$\|f * g\|_2 \le \|f\|_1 \|g\|_2, \qquad \|f * g\|_\infty \le \|f\|_1 \|g\|_\infty, \qquad \|f * g\|_3 \le \|f\|_{3/2} \|g\|_1,$$

as well as the corresponding statements for $\circ$ via $f \circ g = f * \widetilde{g}$.

Moment interpolation. When we have control of $\|h\|_2$ and $\|h\|_\infty$ we may bound intermediate moments by interpolation; for instance

$$\|h\|_3^3 \le \|h\|_2^2\, \|h\|_\infty, \qquad \|h\|_2^2 \le \|h\|_1\, \|h\|_\infty.$$

We will use such inequalities in the "low multiplicity" range of the dyadic decomposition, where r_{A+B} is uniformly small.

2.5 Dyadic decompositions and layer-cake bookkeeping

A recurring device is to decompose a nonnegative function into dyadic level sets. If $f \geq 0$ is finitely supported, we define sets

$$E_j := \{x \in G : \ 2^j \leq f(x) < 2^{j+1}\},$$

so that

$$f(x) \leq \sum_j 2^{j+1}\,\mathbf{1}_{E_j}(x), \qquad f(x) \geq \sum_j 2^j\,\mathbf{1}_{E_j}(x), \tag{2}$$

where the sums range over those integers j for which $E_j \neq \varnothing$. Since f is finitely supported, only $O(\log(2\|f\|_\infty))$ such j occur.

Two simple estimates will be used for summing across scales. First, by disjointness of the E_j we have

$$\|f\|_p^p \sim_p \sum_j 2^{jp}\,|E_j|,$$

with implicit constants depending only on p. Second, Cauchy–Schwarz across the dyadic index gives a generic polylogarithmic loss: for nonnegative coefficients a_j, b_j,

$$\sum_j a_j b_j \leq \Big(\sum_j a_j^2\Big)^{1/2}\Big(\sum_j b_j^2\Big)^{1/2} \leq (\#\{j\})^{1/2}\,\max_j a_j\,\Big(\sum_j b_j^2\Big)^{1/2},$$

and $\#\{j\} \ll \log(2|\operatorname{supp}(f)|)$ once one normalises f dyadically on its support. This is the mechanism by which our arguments introduce only polylogarithmic losses when we replace indicator hypotheses by weighted inequalities.

We apply the same decomposition to representation functions. Given finite A, B and a dyadic parameter t (a power of 2), we define

$$S_t := \{x \in G : \ t \leq r_{A+B}(x) < 2t\}.$$

Then

$$\sum_{x\in G} r_{A+B}(x)^3 \sim \sum_t \sum_{x \in S_t} r_{A+B}(x)^3,$$

and on each level set S_t we have the crude but useful comparability

$$t^3\,|S_t| \ \leq \ \sum_{x\in S_t} r_{A+B}(x)^3 \ < \ (2t)^3\,|S_t|.$$

The proof of the main theorem is organised by estimating the contribution of each S_t with bounds that are summable over t.

2.6 Popularity bounds

Finally, we will repeatedly pass between an ℓ^p bound and a tail estimate. The basic inequality is Markov's: for $h \geq 0$ and $p \geq 1$,

$$|\{x : h(x) \geq \tau\}| \leq \frac{\|h\|_p^p}{\tau^p} \qquad (\tau > 0).$$

Applied with $h = r_{A-A}$ and $p = 3$, this turns an ℓ^3 bound on $\mathbf{1}_A \circ \mathbf{1}_A$ into a quantitative statement that there are few very popular differences. We will typically apply this only for dyadic τ, so that the additional logarithmic bookkeeping is absorbed by our $\lesssim$ notation.

All subsequent arguments rely only on the above conventions and inequalities, together with straightforward rearrangements of finitely supported sums. In particular, no assumptions on G beyond commutativity are used at this stage.

3 Weak control parameters $K_2(A)$ and $K_3(A)$

In Bloom's notion of full L^3-control one asks for a uniform bound on the third moment of $\mathbf{1}_A * \mathbf{1}_B$ for *every* test set B. Many arguments in the literature (notably those originating in work of Shakan and Shkredov) naturally produce weaker information, typically of two kinds: a uniform L^2 bound for mixed correlations and a self-correlation bound for A in an L^3 (or higher-energy) norm. We isolate these as quantitative invariants.

3.1 Definitions and normalisations

For a finite set $A \subset G$ we define $K_2(A)$ to be the least constant $K_2 \geq 0$ such that for every finite $B \subset G$,

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2 \leq K_2\, |A|^{3/4} |B|^{3/4}. \tag{3}$$

Equivalently, in energy language,

$$K_2(A)^2 = \sup_{B \neq \varnothing} \frac{E(A,B)}{|A|^{3/2}|B|^{3/2}}, \qquad E(A,B) = \|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2. \tag{4}$$

We similarly define $K_3(A)$ to be the least constant $K_3 \geq 0$ such that

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 \leq K_3\, |A|^4. \tag{5}$$

Since $\mathbf{1}_A \circ \mathbf{1}_A = r_{A-A}$, the left-hand side is the *third additive energy*

$$E_3(A) := \sum_{d \in G} r_{A-A}(d)^3 = \|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3,$$

and therefore

$$K_3(A) = \frac{E_3(A)}{|A|^4}. \tag{6}$$

The exponents in (3)–(5) are chosen so that the parameters are dimensionless and typically lie in $(0,1]$ (indeed, always ≤ 1 by a trivial estimate). For $K_2(A)$, we use

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2 \le \|\mathbf{1}_A \circ \mathbf{1}_B\|_1 \, \|\mathbf{1}_A \circ \mathbf{1}_B\|_\infty = |A||B| \cdot \|\mathbf{1}_A \circ \mathbf{1}_B\|_\infty,$$

and $\|\mathbf{1}_A \circ \mathbf{1}_B\|_\infty \le \min\{|A|,|B|\} \le (|A||B|)^{1/2}$, which yields

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2 \le |A|^{3/4}|B|^{3/4},$$

hence $K_2(A) \le 1$. For $K_3(A)$, we similarly have

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 \le \|\mathbf{1}_A \circ \mathbf{1}_A\|_\infty \, \|\mathbf{1}_A \circ \mathbf{1}_A\|_2^2 \le |A| \cdot E(A,A) \le |A| \cdot |A|^3 = |A|^4,$$

so $K_3(A) \le 1$.

We also record the elementary lower bounds

$$K_2(A) \ge |A|^{-1/4} \quad (\text{take } B = \{0\}), \qquad K_3(A) \ge |A|^{-1} \quad (\text{since } r_{A-A}(0) = |A|). \tag{7}$$

Thus $K_2(A)$ and $K_3(A)$ are small precisely when A exhibits a degree of additive pseudorandomness in the corresponding moment.

3.2 Invariances and monotonicity

The parameters are invariant under the natural symmetries of the ambient group. If $x \in G$ then translation does not affect difference representations, hence

$$K_2(A+x) = K_2(A), \qquad K_3(A+x) = K_3(A),$$

and similarly $K_2(-A) = K_2(A)$, $K_3(-A) = K_3(A)$. More generally, if $\phi\colon G \to G$ is a group automorphism then $r_{\phi(A)-\phi(B)}(d) = r_{A-B}(\phi^{-1}d)$, so both K_2 and K_3 are invariant under ϕ.

Monotonicity under restriction is immediate and will be used repeatedly when passing to subsets produced by dyadic decompositions. If $A' \subseteq A$ then $\mathbf{1}_{A'} \le \mathbf{1}_A$ pointwise, hence for every B,

$$\|\mathbf{1}_{A'} \circ \mathbf{1}_B\|_2 \le \|\mathbf{1}_A \circ \mathbf{1}_B\|_2, \qquad \|\mathbf{1}_{A'} \circ \mathbf{1}_{A'}\|_3^3 \le \|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3,$$

and therefore

$$K_2(A') \le K_2(A), \qquad K_3(A') \le K_3(A). \tag{8}$$

We will also use crude subadditivity statements for decompositions $A = \bigsqcup_i A_i$. While mixed correlations between the pieces may contribute positively to $\mathbf{1}_A \circ \mathbf{1}_A$, the triangle inequality at the level of norms yields bounds of the shape

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2 \le \sum_i \|\mathbf{1}_{A_i} \circ \mathbf{1}_B\|_2, \qquad \|\mathbf{1}_A \circ \mathbf{1}_A\|_3 \le \sum_i \|\mathbf{1}_{A_i} \circ \mathbf{1}_A\|_3, \tag{9}$$

which are sufficient for the bookkeeping encountered later: when a construction or argument splits A into $O(\log|A|)$ structured layers, the resulting loss can be absorbed into $\log(2|A|)^{O(1)}$.

3.3 Immediate relations to Bloom control

We now relate the weak parameters to Bloom's full L^3-control parameter $\kappa(A)$. Suppose A satisfies Bloom control, meaning that for all finite B,

$$\|\mathbf{1}_A * \mathbf{1}_B\|_3^3 = \sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3 \le \kappa(A)\,|A|^2|B|^2. \tag{10}$$

Then (3) follows from interpolation between L^1 and L^3. Indeed, for any finitely supported $h \ge 0$ we have

$$\|h\|_2 \le \|h\|_1^{1/4}\|h\|_3^{3/4}, \qquad \text{equivalently} \qquad \|h\|_2^4 \le \|h\|_1\,\|h\|_3^3,$$

since $1/2 = (1/4)\cdot 1 + (3/4)\cdot(1/3)$. Applying this with $h = \mathbf{1}_A \circ \mathbf{1}_B = \mathbf{1}_A * \mathbf{1}_{-B}$ gives

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2^4 \le \|\mathbf{1}_A \circ \mathbf{1}_B\|_1\,\|\mathbf{1}_A \circ \mathbf{1}_B\|_3^3 = |A||B|\cdot\|\mathbf{1}_A * \mathbf{1}_{-B}\|_3^3 \le |A||B|\cdot\kappa(A)|A|^2|B|^2,$$

hence

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2 \le \kappa(A)^{1/4}|A|^{3/4}|B|^{3/4}, \qquad \text{so} \qquad K_2(A) \le \kappa(A)^{1/4}. \tag{11}$$

Likewise, taking $B = -A$ in (10) yields

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \|\mathbf{1}_A * \mathbf{1}_{-A}\|_3^3 \le \kappa(A)\,|A|^4, \qquad \text{so} \qquad K_3(A) \le \kappa(A). \tag{12}$$

Thus full control dominates the weak parameters at the expected scales K_2^4 and K_3. The main content of our upgrade theorem is that, up to polylogarithmic factors, the converse holds: the information encoded by (3) and (5) already forces (10) with $\kappa(A)$ comparable to $\max\{K_2(A)^4, K_3(A)\}$.

3.4 Comparison with parameters in Shakan–Shkredov

The inequalities (3)–(5) are closely related to the energy-based parameters used by Shakan and Shkredov. In much of that literature one encounters quantities of the form

$$d^+(A) := \sup_{B \ne \varnothing} \frac{E(A,B)}{|A|\,|B|^{3/2}}, \qquad \text{or more symmetrically} \qquad \sup_{B \ne \varnothing} \frac{E(A,B)}{|A|^{3/2}|B|^{3/2}},$$

as well as higher-energy normalisations based on $E_3(A) = \sum_d r_{A-A}(d)^3$. Our choice is precisely the symmetric normalisation for the mixed energy together with the scale-invariant normalisation $E_3(A)/|A|^4$ for the third energy:

$$K_2(A)^2 = \sup_{B \ne \varnothing} \frac{E(A,B)}{|A|^{3/2}|B|^{3/2}}, \qquad K_3(A) = \frac{E_3(A)}{|A|^4}.$$

If one prefers the asymmetric $d^+(A)$, then (4) implies $d^+(A) \le K_2(A)^2\,|A|^{1/2}$, while conversely $K_2(A)^2 \le d^+(A)\,|A|^{-1/2}$ by restricting to $|B| \le |A|$ (the range typically relevant in applications). Similarly, $K_3(A)$ is exactly the third-energy density, and tail bounds for popular differences that are often assumed in "Szemerédi–Trotter type" hypotheses follow from $K_3(A)$ by Markov's inequality:

$$|\{d:\ r_{A-A}(d) \ge t\}| \le \frac{E_3(A)}{t^3} = K_3(A)\frac{|A|^4}{t^3}.$$

From our perspective, these parameters are the minimal hypotheses one can hope to propagate to full L^3-control: (3) supplies uniform second-moment bounds for mixed sums/differences, while (5) prevents the obstruction coming from excessively many very popular differences of A.

In the next section we begin the analytic mechanism that converts the set-testing hypothesis (3) into a weighted inequality for $\mathbf{1}_A \circ f$, which is the input needed for the dyadic analysis of $\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3$ without incurring catastrophic losses from hard truncations.

4 From set tests to function tests: a layer-cake upgrade of (3)

The hypothesis (3) is stated for indicators of sets. In the proof of the main upgrade theorem, however, we cannot remain in the category of sets: after decomposing the representation function $r_{A+B} = \mathbf{1}_A * \mathbf{1}_B$ into dyadic pieces, we are naturally led to expressions in which $\mathbf{1}_A$ is correlated not with a set but with a *weight* encoding the local density of B on certain fibres. A robust analytic mechanism is therefore needed to convert the set-testing inequality (3) into a weighted inequality for $\mathbf{1}_A \circ f$ with f an arbitrary finitely supported function.

A naive attempt would be to approximate f by a large union of level sets and apply (3) repeatedly, but if one performs this approximation by hard truncations (for instance, cutting off all values above a chosen threshold and bounding the tail trivially), then the resulting dependence on parameters is typically unstable: when the tail is later estimated in a crude way, one incurs large powers of the control parameters (the familiar "κ^{100}-type losses"). The point of the layer-cake viewpoint is that we do *not* separate a "main part" and a "tail part"; instead we keep all scales simultaneously, paying only for the number of relevant dyadic scales, which is logarithmic in a natural size parameter.

4.1 Dyadic layer-cake decomposition

Let $f\colon G \to [0,\infty)$ be finitely supported. We write $S = \mathrm{supp}(f)$ and decompose f dyadically as

$$f = \sum_{j\in\mathbb{Z}} 2^j \, \mathbf{1}_{E_j}, \qquad E_j := \{x \in G : \ 2^j \le f(x) < 2^{j+1}\}. \tag{13}$$

Only finitely many E_j are nonempty. This decomposition is a discrete form of the layer-cake identity $f(x) = \int_0^\infty \mathbf{1}_{\{f>t\}}(x)\,dt$, with the advantage that it keeps track of scales explicitly.

Applying (3) to each E_j gives

$$\|\mathbf{1}_A \circ \mathbf{1}_{E_j}\|_2 \le K_2 \, |A|^{3/4} \, |E_j|^{3/4}. \tag{14}$$

Since $\mathbf{1}_A \circ f = \sum_j 2^j(\mathbf{1}_A \circ \mathbf{1}_{E_j})$ and $\|\cdot\|_2$ is a norm, we obtain the crude bound

$$\|\mathbf{1}_A \circ f\|_2 \le \sum_j 2^j \, \|\mathbf{1}_A \circ \mathbf{1}_{E_j}\|_2 \le K_2 \, |A|^{3/4} \sum_j 2^j |E_j|^{3/4}. \tag{15}$$

Thus the problem reduces to estimating the dyadic sum $\sum_j 2^j |E_j|^{3/4}$ in terms of a natural norm of f. The correct exponent is forced upon us: since $2^j|E_j|^{3/4} = (2^{4j/3}|E_j|)^{3/4}$, the quantity $\sum_j 2^{4j/3}|E_j|$ is comparable to $\sum_x f(x)^{4/3} = \|f\|_{4/3}^{4/3}$. We therefore expect an estimate in terms of $\|f\|_{4/3}$, up to a factor accounting for the number of active dyadic scales.

4.2 A Lorentz-space inequality and a polylog loss

Define $a_j := 2^{4j/3}|E_j|$. Then (15) becomes

$$\|\mathbf{1}_A \circ f\|_2 \le K_2 \, |A|^{3/4} \sum_j a_j^{3/4}.$$

By Hölder on the sequence (a_j) with exponents $4/3$ and 4, we have

$$\sum_j a_j^{3/4} \le \Big(\sum_j a_j\Big)^{3/4} \Big(\sum_j 1\Big)^{1/4} = \Big(\sum_j 2^{4j/3}|E_j|\Big)^{3/4} (\#\{j : \ E_j \ne \varnothing\})^{1/4}. \tag{16}$$

Moreover,

$$\sum_j 2^{4j/3}|E_j| \le \sum_x (2f(x))^{4/3} \ll \|f\|_{4/3}^{4/3}, \tag{17}$$

since $f(x) \in [2^j, 2^{j+1})$ on E_j. Thus

$$\|\mathbf{1}_A \circ f\|_2 \ll K_2 \, |A|^{3/4} \, \|f\|_{4/3} \, (\#\{j : \ E_j \ne \varnothing\})^{1/4}. \tag{18}$$

The factor $\#\{j:\ E_j\neq\varnothing\}$ is controlled by a logarithm. Indeed, for each nonempty E_j we have $2^j\le\|f\|_\infty$, and also $2^j\ge\min\{f(x): x\in S\}$. While $\min f$ may be very small, we may harmlessly discard layers below the average scale: if $2^j\le\|f\|_1/|S|$, then

$$\sum_{k\le j}2^k\mathbf{1}_{E_k}\le\frac{\|f\|_1}{|S|}\mathbf{1}_S,$$

and the contribution of these low layers can be bounded by applying (3) with $B=S$ and absorbing the resulting term into the $\|f\|_{4/3}$ contribution. Concretely, after such a normalization one may assume the active scales lie in an interval of length $O(\log(2|S|))$, and hence

$$\#\{j:\ E_j\neq\varnothing\}\ \ll\ \log(2|S|). \tag{19}$$

Substituting (19) into (18) yields the desired weighted inequality with a polylogarithmic loss.

It is convenient to phrase this in a Lorentz-space language. The sum $\sum_j 2^j|E_j|^{3/4}$ is (up to absolute constants) the discrete analogue of the Lorentz norm $\|f\|_{L^{4/3,1}}$, while $\|f\|_{4/3}$ is the $L^{4/3}$ norm. On a finite set, $L^{4/3,1}$ embeds into $L^{4/3}$ with a logarithmic loss, reflecting precisely the scale-counting argument above. We will not need the formalism, but it is conceptually useful: (3) is an endpoint estimate that naturally upgrades to a Lorentz bound, and the passage from Lorentz back to Lebesgue costs only $\log^{O(1)}$.

4.3 Indicator-to-function upgrade

We summarize the outcome as the following proposition, which is the analytic input for the multi-scale argument in the next section.

Proposition 4.1. *Assume* (3) *holds with constant K_2. Then for any finitely supported $f\ge 0$,*

$$\|\mathbf{1}_A\circ f\|_2\ \ll\ \log(2|\mathrm{supp}(f)|)^{O(1)}\,K_2\,|A|^{3/4}\,\|f\|_{4/3}. \tag{20}$$

If f is signed, then the same bound holds with f replaced by $|f|$, and hence also for general f after splitting into positive and negative parts.

Proof. For $f\ge 0$ we use the dyadic decomposition (13) and apply (3) on each level set E_j to obtain (15). We then estimate the dyadic sum by (16) and (17). Finally, we bound the number of nonempty dyadic levels by (19), losing a factor $\log(2|\mathrm{supp}(f)|)^{O(1)}$. For signed f, write $f=f_+-f_-$ and apply the nonnegative case to f_+ and f_-, using $\|\mathbf{1}_A\circ f\|_2\le\|\mathbf{1}_A\circ f_+\|_2+\|\mathbf{1}_A\circ f_-\|_2$ and $\|f_+\|_{4/3}+\|f_-\|_{4/3}\le 2\|f\|_{4/3}$. ☐

Two comments are in order. First, the exponent 4/3 is not incidental: it is exactly the dual exponent that will appear when we pair $\mathbf{1}_A \circ f$ against another function in L^4, and it is the exponent for which the dyadic summation matches the 3/4 power present in (3). Second, the logarithmic loss in (20) is the correct price for an endpoint upgrade of this type; importantly, it is *uniform* and does not depend on the magnitudes taken by f, only on the number of available scales (which will be $O(\log|A|)$ in our applications).

4.4 Why this removes catastrophic tail losses

In the multi-scale analysis of $\|\mathbf{1}_A * \mathbf{1}_B\|_3^3$, one repeatedly encounters weighted functions of the form

$$f = \sum_{y \in B} w(y)\, \delta_y,$$

or more generally weights obtained by selecting a level set of r_{A+B} and projecting it back to B. The key difficulty is that such weights may have long tails: a small part of B can carry disproportionately large weight, and a hard cutoff at an arbitrary threshold introduces an artificial dichotomy between "structured" and "error" parts. Proposition 4.1 avoids this by treating each dyadic level on its own scale and summing the contributions with the correct exponent. The only cumulative loss comes from counting dyadic scales, and hence is polylogarithmic. This stability is what allows the next section to run the decomposition of r_{A+B} across many multiplicity regimes while keeping the final dependence on $\max\{K_2^4, K_3\}$ at the natural scale.

5 The weak-to-full upgrade: a multi-scale bound for $\|\mathbf{1}_A * \mathbf{1}_B\|_3$

We now prove the main implication that the weak hypotheses (3) and (5) upgrade to Bloom's full L^3-control. Fix an arbitrary finite $B \subset G$ and set

$$r := r_{A+B} = \mathbf{1}_A * \mathbf{1}_B.$$

Our task is to show

$$\sum_{x \in G} r(x)^3 \ \ll\ \log(2|A|)^{O(1)}\, \max\{K_2^4, K_3\}\, |A|^2 |B|^2.$$

The argument proceeds by decomposing r into dyadic multiplicity scales and bounding the contribution of each scale by a combination of the L^2 cross-correlation information (hence K_2) and the popular-differences information coming from the self L^3 bound (hence K_3). The only cumulative loss comes from summing over scales.

5.1 Dyadic decomposition of the representation function

For dyadic $t = 2^k$ we define the level sets

$$S_t := \{x \in G : \ t \le r(x) < 2t\}.$$

Then S_t is empty unless $1 \le t \le \min\{|A|, |B|\}$, and we have the disjoint decomposition

$$r = \sum_t r\,\mathbf{1}_{S_t}.$$

Since $r(x) \asymp t$ on S_t, we obtain

$$\sum_x r(x)^3 = \sum_t \sum_{x \in S_t} r(x)^3 \ll \sum_t t^3 |S_t|. \tag{21}$$

Thus it suffices to bound $t^3|S_t|$ uniformly in t, up to polylogarithmic losses, in such a way that the sum over dyadic t converges to $\max\{K_2^4, K_3\}|A|^2|B|^2$.

5.2 From a level set to a weighted fibre function

To each level set S_t we associate a weight on G that records how often an element of G appears as a B-component of a representation of a sum in S_t. Define

$$f_t(y) := |\{(a, b) \in A \times B : \ a + b \in S_t, \ b = y\}| = \sum_{x \in S_t} \mathbf{1}_A(x - y)\mathbf{1}_B(y).$$

Equivalently, $f_t = \mathbf{1}_B \cdot (\mathbf{1}_A \circ \mathbf{1}_{S_t})(-\,\cdot\,)$, so $\operatorname{supp}(f_t) \subseteq B$ and $\|f_t\|_1$ is exactly the number of representations (a, b) whose sum lies in S_t:

$$\|f_t\|_1 = \sum_y f_t(y) = \sum_{x \in S_t} r(x). \tag{22}$$

In particular, since $r(x) \ge t$ on S_t,

$$\|f_t\|_1 \ge t|S_t|. \tag{23}$$

At the same time, the level contribution to the third moment can be expressed in terms of f_t. Indeed,

$$\sum_{x \in S_t} r(x)^3 \asymp t^2 \sum_{x \in S_t} r(x) = t^2 \|f_t\|_1,$$

where we used $r(x) \asymp t$ on S_t and (22). Thus

$$\sum_{x \in S_t} r(x)^3 \ \ll \ t^2\,\|f_t\|_1. \tag{24}$$

The point is that $\|f_t\|_1$ can be bounded by analytic means from an L^2 control on $\mathbf{1}_A \circ f_t$, and Proposition 4.1 provides exactly such a bound in terms of $\|f_t\|_{4/3}$ with only a polylogarithmic loss.

5.3 Medium multiplicity: two uses of the K_2 hypothesis

We first treat the range of t for which the K_2-information is decisive. By Cauchy–Schwarz on B we have

$$\|f_t\|_1 \le |B|^{1/2}\, \|f_t\|_2. \tag{25}$$

To bound $\|f_t\|_2$, we relate f_t back to S_t and use the weighted upgrade Proposition 4.1. Observing $f_t \le (\mathbf{1}_A \circ \mathbf{1}_{S_t})(-\,\cdot\,)$, we may estimate (up to absolute constants)

$$\|f_t\|_2 \le \|\mathbf{1}_A \circ \mathbf{1}_{S_t}\|_2.$$

Applying (3) with $B = S_t$ yields

$$\|f_t\|_2 \ \ll\ K_2\, |A|^{3/4}\, |S_t|^{3/4}. \tag{26}$$

Substituting (26) into (25) and then into (24) gives

$$\sum_{x\in S_t} r(x)^3 \ \ll\ t^2\, |B|^{1/2}\, K_2\, |A|^{3/4}\, |S_t|^{3/4}. \tag{27}$$

At this stage the level size $|S_t|$ must be eliminated. Here we use the L^2 control on r coming from the energy form of (3) (Lemma **??** in the global notation): by (3) applied to $-B$,

$$\sum_x r(x)^2 = \|\mathbf{1}_A * \mathbf{1}_B\|_2^2 = E(A,-B) \le K_2^2\, |A|^{3/2}|B|^{3/2}.$$

Since $r(x) \ge t$ on S_t, we have

$$t^2|S_t| \le \sum_{x\in S_t} r(x)^2 \le \sum_x r(x)^2 \le K_2^2\, |A|^{3/2}|B|^{3/2}, \tag{28}$$

and hence $|S_t| \ll K_2^2|A|^{3/2}|B|^{3/2}t^{-2}$. Inserting this into (27) gives

$$\sum_{x\in S_t} r(x)^3 \ \ll\ t^2\, |B|^{1/2}\, K_2\, |A|^{3/4} \left(K_2^2|A|^{3/2}|B|^{3/2}t^{-2}\right)^{3/4}.$$

Collecting exponents, we obtain

$$\sum_{x\in S_t} r(x)^3 \ \ll\ K_2^4\, |A|^2|B|^2 \cdot \left(\frac{t}{|A|^{1/2}|B|^{1/2}}\right)^{1/2}. \tag{29}$$

In particular, whenever $t \le |A|^{1/2}|B|^{1/2}$ (which covers all but the extreme high-multiplicity levels), we have the uniform bound

$$\sum_{x\in S_t} r(x)^3 \ \ll\ K_2^4\, |A|^2|B|^2. \tag{30}$$

Summing (30) over the $O(\log(2|A|))$ dyadic values of t in this medium range contributes at most a polylogarithmic factor, as desired.

We emphasize that the exponent K_2^4 is forced by this scheme: the level-set estimate (27) uses K_2 once, and (28) uses K_2^2 (energy), which is then raised to the 3/4 power, producing an overall $K_2 \cdot (K_2^2)^{3/4} = K_2^4$.

5.4 High multiplicity: K_3 and popular differences in $A - A$

It remains to bound the contribution of those (few) dyadic levels with t comparable to $\min\{|A|, |B|\}$, where (29) is no longer uniform. In this regime, large multiplicity forces repeated differences inside A, and hypothesis (5) precisely limits how many such popular differences can exist.

We use the following heuristic (made precise by a Katz–Koester style containment argument): if $r(x) \geq t$, then there are $\gg t^2$ ordered pairs of representations $(a_1, b_1), (a_2, b_2) \in A \times B$ with $a_1 + b_1 = a_2 + b_2 = x$, hence $\gg t^2$ difference relations

$$a_1 - a_2 = b_2 - b_1 \in (A - A) \cap (B - B).$$

Thus, a large level set S_t produces many incidences between $B - B$ and the set of popular differences of A. Lemma 2 in the global notation bounds the popular differences of A: from (5), for all $u \geq 1$,

$$|\{d : \ r_{A-A}(d) \geq u\}| \ \ll \ K_3 \, \frac{|A|^4}{u^3}.$$

Taking $u \asymp t$ and summing over dyadic t yields a convergent series in t once weighted by t^3, exactly matching the third-moment scaling in (21). Concretely, the contribution of the high-multiplicity levels can be bounded by

$$\sum_{t \text{ high}} t^3 |S_t| \ \ll \ \log(2|A|)^{O(1)} \, K_3 \, |A|^2 |B|^2, \tag{31}$$

where the factor $|B|^2$ enters only through the trivial bound $r_{B-B}(d) \leq |B|$ when we convert incidences with $B - B$ into a count of representations.

5.5 Conclusion of the upgrade

Combining the medium estimate (30) summed over dyadic t with the high-multiplicity estimate (31), we obtain

$$\sum_x r(x)^3 \ \ll \ \log(2|A|)^{O(1)} \big(K_2^4 + K_3\big) \, |A|^2 |B|^2 \ \ll \ \log(2|A|)^{O(1)} \, \max\{K_2^4, K_3\} \, |A|^2 |B|^2,$$

which is the asserted full L^3-control bound. We stress that the polylogarithmic factor comes from two sources only: the dyadic decomposition (21) (a logarithmic number of nonempty scales) and the endpoint nature of the L^2 testing inequality (3) when upgraded to weighted functions via Proposition 4.1. No loss of the form K_2^{-c} or K_3^{-c} is incurred.

5.6 Equivalence of weak and full control up to polylogarithmic factors

Theorem A identifies $\max\{K_2^4, K_3\}$ (up to polylogarithmic losses) as the correct quantitative invariant governing Bloom-style L^3 propagation. Conceptually, the point is that Bloom's full control parameter $\kappa(A)$ is defined by a *uniform* third-moment inequality over all finite B, whereas many arguments in the literature (including older threshold-breaking results) assume only *weaker* information: a uniform L^2 cross-correlation estimate and a single self-correlation L^3 estimate. Theorem A shows that, after tolerating logarithmic losses in $\log(2|A|)$, these weaker hypotheses already force the full uniform third-moment bound. Together with the easy reverse implication (full $\Rightarrow$ weak), this gives a robust equivalence and justifies treating $\max\{K_2^4, K_3\}$ as interchangeable with $\kappa(A)$ in applications.

We first record the "easy" direction, which is essentially bookkeeping. Assume that A has full control $\kappa(A)$, meaning that for every finite B we have

$$\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3 \le \kappa(A)\,|A|^2|B|^2.$$

By choosing $B = -A$, we directly obtain the self-correlation estimate

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \sum_x (\mathbf{1}_A * \mathbf{1}_{-A}(x))^3 \le \kappa(A)\,|A|^4,$$

so $K_3(A) \le \kappa(A)$. Similarly, full L^3-control implies the uniform L^2 cross-correlation bound (i) with $K_2 \le \kappa(A)^{1/4}$. One convenient way to see the scaling is to interpolate between the second and third moments of $\mathbf{1}_A * \mathbf{1}_B$: writing $r = \mathbf{1}_A * \mathbf{1}_B$, we have

$$\|r\|_2^2 \le \|r\|_3^{3/2}\,\|r\|_1^{1/2}, \qquad \|r\|_1 = |A||B|, \qquad \|r\|_3^3 \le \kappa(A)|A|^2|B|^2,$$

hence

$$\|r\|_2^2 \le \kappa(A)^{1/2}\,|A|^{3/2}|B|^{3/2}.$$

Using the energy reformulation $E(A, -B) = \|r\|_2^2$ and the symmetry between sum and difference forms, this yields precisely the K_2-type hypothesis with $K_2^2 \le \kappa(A)^{1/2}$, i.e. $K_2 \le \kappa(A)^{1/4}$. In particular,

$$\max\{K_2(A)^4, K_3(A)\} \;\le\; \kappa(A). \tag{32}$$

This calibrates the normalisations and shows that the combination $\max\{K_2^4, K_3\}$ is not an artefact of our proof: it is forced already by the elementary consequences of full control.

The substantive direction is Theorem A: assuming only (i) and (ii), we recover full control with a loss of at most $\log(2|A|)^{O(1)}$. When combined with (32), we obtain the polylogarithmic equivalence

$$\kappa(A) \;\asymp_{\text{polylog}}\; \max\{K_2(A)^4, K_3(A)\}.$$

This equivalence is structurally stable and therefore well-suited to black-box use: one may freely replace any hypothesis expressed in terms of $\kappa(A)$ with hypotheses expressed in terms of $K_2(A)$ and $K_3(A)$, losing only polylogarithmic factors in the final constants.

It is worth isolating why the exponent 4 on K_2 is the correct one. In the medium-multiplicity regime of the dyadic decomposition, our argument uses the uniform L^2 hypothesis twice: once in a direct estimate of $\|\mathbf{1}_A \circ \mathbf{1}_{S_t}\|_2$ (contributing a factor K_2), and once in the energy bound $\|\mathbf{1}_A * \mathbf{1}_B\|_2^2 \leq K_2^2|A|^{3/2}|B|^{3/2}$ (contributing K_2^2, subsequently raised to the $3/4$ power because $|S_t|$ enters as $|S_t|^{3/4}$). Thus the multiplicative structure of the argument forces

$$K_2 \cdot (K_2^2)^{3/4} = K_2^4.$$

Conversely, one cannot generally hope to replace K_2^4 by a smaller power without invoking additional structure, because the third moment is genuinely more sensitive to level-set concentration than the second moment, and the energy bound controls *only* the L^2 mass of r_{A+B}, not its distribution.

The role of K_3 is complementary. The uniform L^2 information alone does not exclude the possibility that the third moment is dominated by a small number of extremely popular sums; such a situation corresponds, via Katz–Koester type containments, to the existence of many popular differences in $A - A$. Hypothesis (ii) precisely rules out this obstruction by controlling the size of the level sets $\{d : r_{A-A}(d) \geq t\}$ with a t^{-3} tail, which is exactly the decay needed to make the third-moment summation converge at the high end. This explains why the final invariant is a maximum: in any given configuration, either the medium levels are dominant (leading to K_2^4) or the extreme levels are dominant (leading to K_3), and the stronger of the two constraints dictates the outcome.

We also indicate explicitly where the polylogarithmic losses enter. There are two sources and both are of the same, essentially unavoidable, nature in an endpoint multi-scale argument.

1. The dyadic decomposition of r_{A+B} introduces a sum over $O(\log(2|A|))$ relevant multiplicity scales t. Even when each scale is controlled uniformly, summing the bounds incurs a logarithmic factor. This is the same phenomenon that appears in many additive-combinatorial "energy increment" or "popularity" arguments.

2. The passage from indicator testing (hypothesis (i) for all sets B) to weighted testing (needed to treat fibre functions such as f_t) is performed by a layer-cake decomposition and Cauchy–Schwarz across dyadic layers. This is an endpoint substitute for a strong-type inequality of the form $\|\mathbf{1}_A \circ f\|_2 \lesssim K_2|A|^{3/4}\|f\|_{4/3}$ with no loss. In general such a lossless upgrade need not hold in this level of generality, and our Proposition C quantifies the best available bound in terms of $\log(2|\mathrm{supp}(f)|)$.

Importantly, there is no loss in negative powers of K_2 or K_3: once the hypotheses hold with $K_2, K_3 \leq 1$, the argument never requires a density increment or a decomposition that worsens these parameters. This stability is what makes the equivalence useful in practice: the polylogarithmic losses are absorbed in the same way as in Bloom's original framework note, and do not disrupt threshold phenomena governed by power savings in $|A|$.

As a final remark on tightness, we note that the equivalence is compatible with the standard test cases. For highly structured sets (such as arithmetic progressions), one expects $\kappa(A)$, $K_2(A)$, and $K_3(A)$ to be bounded below by absolute constants, consistent with the absence of genuine L^p improvement. For sets enjoying incidence-geometric control (such as convex sets in $\mathbb{R}$), the known bounds give $K_3(A) \lesssim |A|^{-1}$ and $K_2(A) \lesssim |A|^{-1/4}$, so $K_2(A)^4$ matches $K_3(A)$ at the correct scale and hence recovers the expected full control. In pseudorandom regimes, $K_2(A)$ may be close to 1 while $K_3(A)$ is dictated by density considerations, again aligning with the philosophy that $\kappa(A)$ should be read off from the dominant obstruction to uniformity.

We therefore treat $\max\{K_2^4, K_3\}$ as the effective control parameter for A, with the understanding that all statements are stable up to factors of $\log(2|A|)^{O(1)}$. In the next section we exploit this by rewriting Bloom-style propagation theorems so that their hypotheses may be checked using only the weaker L^2 and self-L^3 inputs, without changing the conclusions beyond polylogarithmic losses.

5.7 Applications: replacing full control by weak control as a black box

Many propagation arguments in additive combinatorics are stated under Bloom's full L^3-control hypothesis, i.e. an a priori bound on

$$\sup_{B \subset G \text{ finite}} \frac{1}{|A|^2|B|^2} \sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3.$$

Corollary B allows us to regard this hypothesis as interchangeable (up to polylogarithmic loss) with the pair of weaker inputs (i)–(ii). Concretely, whenever a theorem assumes $\kappa(A) \leq K$, we may instead assume

$$K_2(A) \;\leq\; K^{1/4} \qquad \text{and} \qquad K_3(A) \;\leq\; K,$$

and obtain the same conclusion with K replaced by $\log(2|A|)^{O(1)} \max\{K_2(A)^4, K_3(A)\}$. Conversely, if one proves bounds for $K_2(A)$ and $K_3(A)$ by geometric or Fourier-analytic means, then Theorem A upgrades these to the full uniform third-moment estimate required by Bloom's framework.

We emphasize that this replacement is genuinely "black box": the proofs of most propagation lemmas use the full control hypothesis only through inequalities of the type

$$\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3 \ \ll\ \kappa(A)\,|A|^2|B|^2$$

for specific auxiliary sets B (or weighted variants thereof), together with the stability of $\kappa(\cdot)$ under simple operations. Theorem A supplies precisely such estimates from (i)–(ii), and Proposition C supplies the needed indicator-to-weighted upgrade at the cost of a polylogarithmic factor in $\log(2|A|)$. Thus one may re-run the original arguments with $\kappa(A)$ replaced everywhere by $\log(2|A|)^{O(1)}\max\{K_2(A)^4, K_3(A)\}$, without changing the combinatorial skeleton.

Convex-set and incidence-geometric inputs. A common pattern in applications is that one can obtain a self-correlation bound of the form (ii) by incidence geometry (e.g. Szemerédi–Trotter), while the full uniform L^3-control over all B is not checked directly. For instance, when $A \subset \mathbb{R}$ is convex (or more generally an image of an interval under a strictly convex function), the known incidence machinery yields strong control on popular differences, which can be stated in our language as

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 \ \ll\ |A|^3,$$

i.e. $K_3(A) \ll |A|^{-1}$. In many arguments one also has an energy estimate consistent with (i), for example

$$E(A,B) \ =\ \|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2 \ \ll\ |A|\,|B|^{3/2} \qquad \text{(typical convex-set energy behaviour)},$$

which implies $K_2(A) \ll |A|^{-1/4}$ (up to absolute constants, and with the natural $|B|^{3/2}$ scaling). At this point Theorem A gives

$$\kappa(A) \ \ll\ \log(2|A|)^{O(1)}\max\{|A|^{-1}, |A|^{-1}\} \ \ll\ \log(2|A|)^{O(1)}|A|^{-1},$$

recovering the full control statement normally quoted as a starting point for Bloom-style propagation. Thus, in convex-set arguments, it is enough to verify the two moment bounds (i)–(ii), which are often closer to the native output of incidence theory than the uniform third-moment inequality.

Sum-product decompositions and "structured+random" reductions. In additive-multiplicative settings (e.g. subsets of a field, or $\mathbb{F}_p^n$ with two operations), one frequently decomposes a set A into pieces $A = \bigsqcup_i A_i$ such that each A_i behaves well with respect to one of the two operations. In such proofs, the full L^3-control of each piece is conceptually convenient but often not directly accessible: one may have (i) available for A_i as a consequence of a uniform energy bound against arbitrary test sets (coming from Fourier uniformity or pseudorandomness), and (ii) available because the self-correlation is controlled for structural reasons (e.g. A_i is contained in a low-dimensional progression, or has an incidence interpretation).

Here Corollary B interacts well with the monotonicity properties of K_2 and K_3: if $A' \subseteq A$ then $K_2(A') \leq K_2(A)$ and $K_3(A') \leq K_3(A)$, while for disjoint unions one has a subadditivity at the level relevant for moment bounds (Lemma 5). Thus, if a decomposition argument produces pieces A_i each satisfying (i)–(ii) with parameters $K_{2,i}, K_{3,i}$, then each A_i automatically enjoys full control with parameter

$$\kappa(A_i) \ll \log(2|A_i|)^{O(1)} \max\{K_{2,i}^4, K_{3,i}\},$$

and one can import any Bloom-style propagation result on each piece. This removes the need to prove a uniform third-moment inequality separately for each A_i, which is typically the most awkward part of such decompositions.

Balog–Szemerédi–Gowers type outputs from control. Another standard use of full control is as a hypothesis in BSG-from-control statements: roughly, if a set (or pair of sets) has large additive energy, then one can find large subsets with small doubling, with quantitative losses governed by the control parameter. In Bloom's framework the role of $\kappa(A)$ is to provide, uniformly in auxiliary sets, an L^3 bound that converts energy information into structural information by a popularity argument on level sets of r_{A+B} and repeated applications of Cauchy–Schwarz.

Theorem A implies that the same BSG conclusion holds under the weak hypotheses (i)–(ii). Indeed, the typical BSG-from-control proof needs to bound expressions of the form

$$\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3 \quad \text{and} \quad \|\mathbf{1}_A \circ f\|_2$$

for certain B extracted from popular level sets, and for weights f which are dyadic truncations of representation functions. The first quantity is handled directly by Theorem A, while the second is handled by Proposition C, giving (up to polylogarithmic loss) the same inequalities Bloom obtains from full control. Consequently, any conclusion of the form

$$\text{(large energy)} \quad E(A,B) \geq \eta |A|^{3/2}|B|^{3/2} \quad \Longrightarrow \quad \text{(structured subsets)}$$

with constants depending on $\kappa(A)$ may be re-stated with $\kappa(A)$ replaced by $\log(2|A|)^{O(1)} \max\{K_2(A)^4, K_3(A)\}$. In practical terms, one may run a BSG argument assuming only the uniform L^2 cross-correlation control (i) together with the single self L^3 bound (ii), and obtain the same structured-set output at essentially the same quantitative strength.

Rewriting hypotheses in propagation theorems. We record the general template. Suppose a statement in Bloom's framework is of the schematic form

$$\kappa(A) \leq K \quad \Longrightarrow \quad \mathcal{P}(A) \text{ holds with constants depending on } K,$$

where $\mathcal{P}(A)$ is some propagation property (energy growth bounds, sumset lower bounds, popular sumset structure, or a BSG-type conclusion). By Corollary B, it is equivalent (up to polylogarithmic losses) to assume

$$K_2(A) \;\le\; K^{1/4}, \qquad K_3(A) \;\le\; K,$$

and to conclude $\mathcal{P}(A)$ with K replaced by $\log(2|A|)^{O(1)} \max\{K_2(A)^4, K_3(A)\}$. This restatement is often more natural to verify: (i) is an energy bound against arbitrary B, and (ii) is a single higher-moment self-correlation bound, which can be attacked by direct combinatorial counting, incidence geometry, or Fourier-analytic methods depending on the setting.

In summary, Theorem A and Proposition C allow us to treat $\max\{K_2(A)^4, K_3(A)\}$ as the effective control parameter in essentially any argument that previously used $\kappa(A)$. The resulting reformulations reduce the burden of checking full uniform L^3-control, while preserving (up to polylogarithmic factors) the quantitative strength of Bloom-style propagation theorems.

5.8 Examples and non-examples: sanity checks for the scales of K_2^4 and K_3

We record a few model computations indicating that the normalisations in the definitions of $K_2(A)$ and $K_3(A)$ are consistent with the scale of Bloom's control parameter $\kappa(A)$, and that the combination $\max\{K_2(A)^4, K_3(A)\}$ behaves as the correct effective invariant. Throughout, all implicit constants are absolute, and we freely ignore polylogarithmic losses in $\log(2|A|)$.

1. A baseline: subgroups and cosets (maximally additive structure). Let G be a finite abelian group and let $A = H \le G$ be a subgroup of size $|H| = m$. Then

$$r_{A-A}(d) = \begin{cases} m, & d \in H, \\ 0, & d \notin H, \end{cases} \qquad \text{so} \qquad \|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \sum_{d\in G} r_{A-A}(d)^3 = |H|\, m^3 = m^4,$$

hence $K_3(A) = 1$. For (i), taking $B = H$ gives $\mathbf{1}_A \circ \mathbf{1}_B = \mathbf{1}_H \circ \mathbf{1}_H$ with $\|\mathbf{1}_H \circ \mathbf{1}_H\|_2^2 = |H|\, m^2 = m^3$, so $\|\mathbf{1}_H \circ \mathbf{1}_H\|_2 = m^{3/2}$ and therefore $K_2(A) \ge 1$. Conversely, for any B one has the crude bound $r_{H-B}(x) \le m$ and $\sum_x r_{H-B}(x) = m|B|$, hence $\|\mathbf{1}_H \circ \mathbf{1}_B\|_2^2 \le m \cdot m|B| = m^2|B|$, which is consistent with $K_2(A) \asymp 1$ under the $|A|^{3/2}|B|^{3/2}$ scaling (indeed, the worst case is when $|B|$ is comparable to $|H|$). Finally, Bloom control is also constant: taking $B = H$,

$$(\mathbf{1}_H * \mathbf{1}_H)(x) = m\, \mathbf{1}_H(x) \quad \Rightarrow \quad \sum_x (\mathbf{1}_H * \mathbf{1}_H(x))^3 = |H|\, m^3 = m^4 = |A|^2|B|^2,$$

so $\kappa(H) \ge 1$, and trivially $\kappa(H) \ll 1$ as well. Thus $\kappa(A)$, $K_2(A)^4$, and $K_3(A)$ are all $\asymp 1$ in this structured regime.

2. Arithmetic progressions in $\mathbb{Z}$ (near-maximal energy without subgroup structure). Let $A = \{0, 1, \ldots, n-1\} \subset \mathbb{Z}$, so $|A| = n$ and

$$r_{A-A}(d) = n - |d| \qquad (|d| \leq n-1).$$

We compute

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \sum_{d=-(n-1)}^{n-1} (n-|d|)^3 = n^3 + 2\sum_{k=1}^{n-1} k^3 = n^3 + 2\Big(\frac{(n-1)^2 n^2}{4}\Big) \asymp n^4,$$

so $K_3(A) \asymp 1$. Likewise,

$$E(A, A) = \|\mathbf{1}_A \circ \mathbf{1}_A\|_2^2 = \sum_{d=-(n-1)}^{n-1} (n-|d|)^2 = n^2 + 2\sum_{k=1}^{n-1} k^2 \asymp n^3,$$

hence $\|\mathbf{1}_A \circ \mathbf{1}_A\|_2 \asymp n^{3/2}$, which saturates the $|A|^{3/4}|A|^{3/4} = n^{3/2}$ scaling and therefore forces $K_2(A) \asymp 1$. These computations match the expectation that progressions have $\kappa(A) \asymp 1$: for instance $B = A$ already gives $\sum_x (\mathbf{1}_A * \mathbf{1}_A(x))^3 \asymp n^5$ while $|A|^2|B|^2 = n^4$, yielding a constant-sized control ratio once one notes $\mathbf{1}_A * \mathbf{1}_A$ is triangular of height n supported on an interval of length $2n$. Again K_2^4 and K_3 are both constant and of the correct scale.

3. Random sets in a finite group (pseudorandom behaviour and the role of K_2^4). Let G be a finite abelian group of size N, and let $A \subset G$ be a Bernoulli random set of density $\alpha \in (0, 1)$, so $|A| \approx \alpha N$. Heuristically, for a fixed finite $B \subset G$ with $|B| = \beta N$, the values of $r_{A-B}(x)$ are approximately concentrated around $|A||B|/N \approx \alpha\beta N$, and one expects

$$\|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2 = \sum_x r_{A-B}(x)^2 \approx N\Big(\frac{|A||B|}{N}\Big)^2 \approx \alpha^2\beta^2 N^3.$$

Comparing with the normalisation $|A|^{3/2}|B|^{3/2} \approx (\alpha\beta)^{3/2}N^3$, this suggests a typical value

$$K_2(A)^2 \approx \frac{\alpha^2\beta^2}{(\alpha\beta)^{3/2}} = (\alpha\beta)^{1/2}, \qquad \text{hence} \qquad K_2(A) \approx (\alpha\beta)^{1/4}.$$

Since $K_2(A)$ is defined via a supremum over B, the worst case among densities $\beta \in (0, 1]$ gives $K_2(A) \approx \alpha^{1/4}$, so $K_2(A)^4 \approx \alpha$.

For K_3, one expects $r_{A-A}(d) \approx |A|^2/N \approx \alpha^2 N$ for most d, and hence

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \sum_d r_{A-A}(d)^3 \approx N(\alpha^2 N)^3 = \alpha^6 N^4,$$

so $K_3(A) \approx \alpha^2$. In particular, in the sparse regime $\alpha \ll 1$ we have $K_2(A)^4 \gg K_3(A)$.

This is consistent with the lower bounds forced by Bloom control. Indeed, taking $B = G$, we have $(\mathbf{1}_A * \mathbf{1}_G)(x) = |A|$ for all x, so

$$\sum_x (\mathbf{1}_A * \mathbf{1}_G(x))^3 = N|A|^3, \qquad \text{and} \qquad \frac{1}{|A|^2|G|^2} \sum_x (\mathbf{1}_A * \mathbf{1}_G(x))^3 = \frac{|A|}{N} \approx \alpha.$$

Thus $\kappa(A) \gtrsim \alpha$, which matches $K_2(A)^4 \approx \alpha$ and shows that in this model the parameter K_2^4 is the one that captures the obstruction coming from very large test sets B, even though the self-correlation parameter K_3 is smaller.

4. Sidon-type sets and near-minimal K_3. At the opposite extreme from progressions, suppose A is (approximately) Sidon in the sense that $r_{A-A}(d) \leq 1$ for all $d \neq 0$. Then

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = r_{A-A}(0)^3 + \sum_{d \neq 0} r_{A-A}(d)^3 \leq |A|^3 + |A|^2,$$

so $K_3(A) \ll |A|^{-1}$ (up to negligible lower-order terms). This matches the heuristic that a set with essentially no repeated differences should be close to the "minimal" self-correlation permitted by the diagonal contribution $d = 0$, and it coincides with the convex-set scale discussed next. In such cases, any nontrivial full control must come from K_2^4 if it comes at all, and Theorem A predicts precisely that: the effective control is $\max\{K_2(A)^4, |A|^{-1}\}$.

5. Convex sets and images of intervals under strictly convex maps (incidence-driven bounds). Let $A \subset \mathbb{R}$ be a finite convex set in the standard additive-combinatorial sense (e.g. $A = \{f(1), \ldots, f(n)\}$ with f strictly convex), so $|A| = n$. Incidence geometry (via Szemerédi–Trotter) implies that the difference representation function has few popular values; in a convenient packaged form one has

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \sum_d r_{A-A}(d)^3 \ll n^3,$$

hence $K_3(A) \ll n^{-1}$. Separately, convexity implies favourable energy behaviour against arbitrary test sets $B \subset \mathbb{R}$: one commonly obtains estimates of the shape

$$E(A, B) = \|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2 \ll n\,|B|^{3/2},$$

which is consistent with the folklore principle that convex sets behave additively like sets with "few" repeated differences. Comparing with the required normalisation $K_2(A)^2 n^{3/2} |B|^{3/2}$, this gives $K_2(A)^2 \ll n^{-1/2}$, i.e. $K_2(A) \ll n^{-1/4}$, and therefore $K_2(A)^4 \ll n^{-1}$. In particular, for convex sets we have matching scales

$$K_3(A) \ll n^{-1} \qquad \text{and} \qquad K_2(A)^4 \ll n^{-1},$$

so $\max\{K_2(A)^4, K_3(A)\} \ll |A|^{-1}$, which is the correct order of magnitude for the strongest known L^3-control inequalities in this setting. This is the basic sanity check that our normalisations align with the incidence-geometric regime: the third-moment self-correlation (ii) and the uniform energy control (i) are naturally produced at the same scale.

6. A note on "non-examples" and what they indicate. The preceding cases exhibit two robust phenomena. First, when A carries a large internal additive symmetry (subgroups, progressions), both $K_2(A)$ and $K_3(A)$ are $\asymp 1$, and one should not expect any smallness in $\kappa(A)$. Second, in pseudo-random or incidence-controlled regimes (random sets of small density; convex images), $K_2(A)^4$ and $K_3(A)$ are typically $\ll 1$, and, crucially, they often coincide in scale up to constants (convexity) or one dominates in a predictable way (random sets, where K_2^4 captures the obstruction from large B). In particular, these computations give no evidence for a gap between weak control and full control: in each model, $\kappa(A)$ is naturally of the same order as $\max\{K_2(A)^4, K_3(A)\}$ (up to polylogarithmic slack).

This leaves only a narrow window in which a separation could plausibly occur: one would need sets A for which uniform L^2 cross-correlation remains small against *all* finite B (so $K_2(A)$ is small), and the self L^3 correlation is also small (so $K_3(A)$ is small), yet there exists a carefully tuned B making $\|\mathbf{1}_A * \mathbf{1}_B\|_3$ anomalously large. The next subsection describes a concrete program for searching for such families in $\mathbb{F}_p^n$, where the geometry of cosets and unions of subspaces provides a plausible mechanism for decoupling these moment conditions.

5.9 A conditional separation program in $\mathbb{F}_p^n$

In view of Corollary B, a genuine gap between weak control and full control could only occur if there exist families $A \subset G$ for which both weak parameters $K_2(A)$ and $K_3(A)$ are small, yet $\kappa(A)$ is much larger than $\max\{K_2(A)^4, K_3(A)\}$ (beyond the polylogarithmic slack). We do not know such a family. However, if the upgrade implication were to fail, then it is reasonable to expect that a counterexample can be found inside vector spaces $G = \mathbb{F}_p^n$, where coset geometry provides an explicit way to tune additive statistics and where the relevant quantities can be computed exactly for moderate p^n.

Why $\mathbb{F}_p^n$ is the natural search space. In $G = \mathbb{F}_p^n$, every subset can be represented as a bitstring of length p^n, and convolutions and correlations can be evaluated either by direct counting or via Fourier transform on G. In particular:

$$K_3(A) = \frac{\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3}{|A|^4} = \frac{1}{|A|^4} \sum_{d \in G} r_{A-A}(d)^3$$

is determined solely by the distribution of differences of A, and is straightforward to compute once r_{A-A} is known. Likewise, for any candidate witness set B, both

$$E(A,B) = \|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2 \quad \text{and} \quad \sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3$$

can be computed explicitly, so that one can test whether the ratio defining $\kappa(A)$ is unexpectedly large for some B. The difficulty is that (i) and the definition of $\kappa(A)$ involve quantification over all finite B, so any feasible search must use principled witness classes and diagnostics indicating what form an extremising B should take.

A structural heuristic for what a separation would require. Let us write $F = \mathbf{1}_A * \mathbf{1}_B$. Then $\sum_x F(x)^3$ is the third moment of the degree sequence of the bipartite sum graph between A and B, while $E(A,B) = \|\mathbf{1}_A \circ \mathbf{1}_B\|_2^2$ counts additive 4-cycles (equivalently, common differences). A separation would therefore resemble a combinatorial design: many vertices x with moderately large $F(x)$, but arranged so that difference coincidences (which feed $E(A,B)$ and ultimately K_2) remain scarce, and at the same time the internal difference distribution of A remains flat enough that K_3 is small. In other words, one needs a mechanism producing many length-3 stars in the sum graph without producing too many length-4 cycles and without creating too many popular differences in $A-A$. Vector spaces admit exactly the kind of multi-scale additive decompositions that might allow this.

Candidate templates: unions of cosets with quotient pseudorandomness. Fix a subspace $H \le G$ of size $|H| = h$, and choose a set $T \subset G/H$ of size m. Let

$$A = \bigcup_{t \in T} (t + H), \qquad \text{so} \qquad |A| = mh.$$

If the cosets are disjoint (equivalently, T is a genuine subset of G/H), then for $d \in H$ we have $r_{A-A}(d) = m\, r_{H-H}(d) = mh$, whereas for $d \notin H$ the value of $r_{A-A}(d)$ is controlled by the number of representations of the coset $d + H$ as a difference $t - t'$ in the quotient. In particular, the contribution of differences in H to $K_3(A)$ is

$$\sum_{d \in H} r_{A-A}(d)^3 = |H|\,(mh)^3 = m^3h^4, \qquad \text{so} \qquad \frac{1}{|A|^4} \sum_{d \in H} r_{A-A}(d)^3 = \frac{m^3h^4}{m^4h^4} = \frac{1}{m}.$$

Thus, provided the quotient difference statistics of T are not too concentrated, such multi-coset sets naturally have $K_3(A)$ as small as m^{-1}. This makes them plausible candidates for having simultaneously small $K_3(A)$ and nontrivial internal structure.

However, such A also come with obvious large-energy witnesses B if T itself is additively structured in G/H. The search therefore suggests taking T to be pseudorandom in the quotient (e.g. a random subset of G/H of density θ), so that A is highly structured along H but randomly distributed across cosets. One can then vary $\dim H$ and θ to tune $|A|$, $K_3(A)$, and the apparent extremisers for (i).

Two-scale mixtures: cosets plus sparse noise. A more flexible family is obtained by mixing a multi-coset component and a sparse pseudorandom component:

$$A = \Big(\bigcup_{t\in T} (t + H) \Big) \ \cup \ R,$$

where $R \subset G$ is a Bernoulli set of small density ρ, independent of T. The intended role of R is to "regularise" the distribution of differences and sums so that no single mechanism forces $K_2(A)$ or $K_3(A)$ to be large, while retaining enough structured mass to permit a carefully chosen B to inflate $\|\mathbf{1}_A * \mathbf{1}_B\|_3$. In computations, one should monitor the empirical tail bound

$$|\{d : \ r_{A-A}(d) \geq t\}|$$

and compare it to the t^{-3}-decay suggested by Lemma 2. Any serious candidate for separation must look "as if" it satisfies Lemma 2 with a small constant, while still admitting a B producing unusually heavy tails for r_{A+B}.

Hybrid linear-algebraic constructions beyond cosets. Coset unions are not the only linear objects available in $\mathbb{F}_p^n$. Another natural class is graphs of linear maps (or unions thereof). For a decomposition $G = U \oplus V$ and linear maps $L_i : U \to V$, one may take

$$A = \bigcup_{i=1}^{m} \{(u, L_i u) : u \in U\}.$$

Such sets have controlled intersection patterns between translates, and their sumsets with appropriately chosen B can exhibit structured multiplicity profiles reminiscent of incidence configurations. The hope (in a separation scenario) would be that these sets exhibit small self-correlation in the sense of (ii) because distinct graphs intersect little, while still allowing a witness B (perhaps a union of dual graphs) that causes many sums to have multiplicity on the order of m across a large portion of G. This is precisely the kind of "many stars, few 4-cycles" phenomenon that could potentially evade an L^2-based control while inflating an L^3 moment.

What needs to be verified computationally. Given a candidate $A \subset \mathbb{F}_p^n$, we propose the following finite verification procedure.

(1) *Compute $K_3(A)$ exactly.* Compute $r_{A-A} = \mathbf{1}_A \circ \mathbf{1}_A$ and then $\sum_d r_{A-A}(d)^3$. This yields $K_3(A)$ with no optimisation.

(2) *Upper-bound $K_2(A)$ by searching a witness class for (i).* Since

$$K_2(A)^2 = \sup_{B\neq\emptyset} \frac{E(A,B)}{|A|^{3/2}|B|^{3/2}}, \qquad E(A,B) = \sum_x r_{A-B}(x)^2,$$

a direct supremum over all B is infeasible. One therefore restricts to a structured witness class $\mathcal{W}$ that plausibly contains near-extremisers, such as: subspaces and cosets of each dimension; unions of a bounded number of cosets of a fixed subspace; random subsets of prescribed densities; and level sets of $|\widehat{\mathbf{1}_A}|$ (Fourier-spectrum witnesses). The output is a certified lower bound on $K_2(A)$ and, if no witness is found above a target threshold, evidence (not proof) that $K_2(A)$ is small.

(3) *Search for B inflating the L^3 convolution.* Define

$$\kappa_B(A) := \frac{1}{|A|^2|B|^2} \sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3, \qquad \text{so} \qquad \kappa(A) = \sup_B \kappa_B(A).$$

Again, one searches over a witness class $\mathcal{W}'$ containing: $B = A$, $B = -A$, $B = G$; subspaces and cosets; unions of cosets aligned with the quotient structure used to define A; and greedy "threshold" sets of the form $B = \{x : (\mathbf{1}_A * \mathbf{1}_C)(x) \geq \tau\}$ for auxiliary C (an attempt to approximate extremisers suggested by layer-cake decompositions). Any observed $\kappa_B(A)$ substantially larger than $\max\{K_2(A)^4, K_3(A)\}$ (accounting for polylogarithms) flags A as a potential separator.

(4) *Diagnostics locating the scale of failure.* For a flagged pair (A, B), one should examine the dyadic profile of r_{A+B}: for t dyadic, set $S_t = \{x : r_{A+B}(x) \in [t, 2t)\}$ and record

$$M_3(t) := \sum_{x \in S_t} r_{A+B}(x)^3.$$

A separation would manifest as a range of t where $M_3(t)$ is anomalously large while (a) the energy $E(A, B)$ remains small relative to $|A|^{3/2}|B|^{3/2}$, and (b) the popular-differences tail of r_{A-A} remains small relative to $K_3(A)|A|^4/t^3$. This pinpoints whether the obstruction lives in a medium-multiplicity regime (suggesting a gap in the L^2-to-L^3 transfer) or in a high-multiplicity regime (suggesting that K_3 fails to control the relevant difference concentrations).

Feasibility and expected outcomes. For moderate sizes (say p^n up to 10^6 in total group size), the above steps can be carried out exactly, and one can iterate over families of parameters $(p, n, \dim H, |T|, \rho)$. We emphasise that failure to find a separating B in a witness class is not a proof of equivalence; nonetheless, if the upgrade implication were genuinely false, one would expect counterexamples to exhibit a relatively rigid and reproducible geometry, and hence to be discoverable by such a structured search. Conversely, if extensive searches across these linear-algebraic templates never produce $\kappa_B(A)$ exceeding the predicted scale, this provides strong empirical evidence that weak control and full control are indeed equivalent (up to polylogarithms) in the most plausible finite-field setting for separation.

5.10 Further questions

We record several problems suggested by the upgrade theorem and by the role of the weak parameters $K_2(A)$ and $K_3(A)$ as surrogates for full Bloom control. We focus on three directions: inverse/stability phenomena for weak control, higher-moment analogues (especially L^4 in finite fields), and an operator-norm/incidence perspective that may clarify the mechanism behind the upgrade and the origin of the polylogarithmic losses.

(1) Stability and inverse problems for weak control. The hypotheses (i)–(ii) admit natural "best possible" scales. Indeed, taking $B = \{0\}$ in (i) gives

$$\|\mathbf{1}_A \circ \mathbf{1}_{\{0\}}\|_2 = \|\mathbf{1}_A\|_2 = |A|^{1/2} \leq K_2|A|^{3/4}, \qquad \text{hence} \qquad K_2(A) \geq |A|^{-1/4}.$$

Likewise $r_{A-A}(0) = |A|$ implies

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_3^3 = \sum_d r_{A-A}(d)^3 \geq |A|^3, \qquad \text{hence} \qquad K_3(A) \geq |A|^{-1}.$$

It is therefore natural to renormalise

$$\widetilde{K}_2(A) := K_2(A)\,|A|^{1/4} \geq 1, \qquad \widetilde{K}_3(A) := K_3(A)\,|A| \geq 1,$$

and to ask for structural consequences when $\widetilde{K}_2(A)$ and $\widetilde{K}_3(A)$ are close to 1.

At the extreme $\widetilde{K}_2(A) \approx 1$, condition (i) with $B = A$ forces

$$E(A, A) = \|\mathbf{1}_A \circ \mathbf{1}_A\|_2^2 \leq K_2(A)^2|A|^3 \approx |A|^2,$$

so A has nearly minimal additive energy. In a model case, a Sidon set satisfies $r_{A-A}(d) \leq 1$ for all $d \neq 0$, giving $E(A, A) \asymp |A|^2$ and $\sum_d r_{A-A}(d)^3 \asymp |A|^3$, i.e. $\widetilde{K}_2(A), \widetilde{K}_3(A) \asymp 1$. This suggests the following.

Question 10.1 (removal/stability at the Sidon scale). Assume $\widetilde{K}_2(A) \leq 1 + \varepsilon$ and $\widetilde{K}_3(A) \leq 1 + \varepsilon$. Must A contain a subset $A' \subseteq A$ with $|A'| \geq (1 - O(\varepsilon))|A|$ such that A' is Sidon (or at least $r_{A'-A'}(d) \leq 1$ for all but $O(\varepsilon|A|^2)$ differences d)?

A positive answer would amount to a quantitative removal lemma for additive quadruples/difference collisions at very low density. One might attempt to encode collisions as 4-cycles in a natural bipartite graph and use a C_4-removal mechanism; however, the usual graph removal bounds are too weak to be meaningful at the $|A|^{-1}$ energy scale. A more specialised approach exploiting the algebraic form of additive quadruples may be necessary.

More generally, one may ask for an inverse theory in regimes where $\widetilde{K}_2(A)$ and $\widetilde{K}_3(A)$ are moderately bounded (say $O(1)$), but not necessarily close to 1. Since Theorem A converts weak control to full control up to polylogarithms, any inverse statement for small $\kappa(A)$ transfers (again up to polylogarithms) to weak control. The point is that existing consequences of small $\kappa(A)$ (e.g. growth bounds, energy propagation, BSG-type conclusions) are typically *forward* statements; they do not directly classify sets with small $\kappa(A)$. It would be useful to know whether there is a meaningful dichotomy: either A is essentially Sidon/pseudorandom, or else it correlates with a structured object (cosets, generalized arithmetic progressions, graphs of homomorphisms in $\mathbb{F}_p^n$, etc.).

Question 10.2 (inverse theory for κ and for weak control). Is there a robust classification (even conjectural) of finite $A \subset G$ for which $\kappa(A) \ll 1$ (or equivalently $\max\{K_2(A)^4, K_3(A)\} \ll 1$ up to polylogarithms)? In particular, can one characterise the near-extremisers for the inequality

$$\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3 \ll \kappa(A)\,|A|^2|B|^2 \quad \text{uniformly in } B?$$

A related stability issue concerns the polylogarithmic losses in Theorem A. In our argument, they arise from dyadic decompositions (both in Proposition C and in the layer-cake analysis of r_{A+B}). It is not clear whether these losses are an artefact of the proof or genuinely necessary in full generality.

Question 10.3 (log-free upgrade). Can Theorem A be strengthened to

$$\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^3 \ll \max\{K_2^4, K_3\}\,|A|^2|B|^2$$

with an absolute implied constant (no polylogarithmic factor), or can one produce examples showing that a logarithmic loss is unavoidable?

(2) Higher moments: L^4 control in finite fields. Bloom control is an L^3 statement, and the upgrade theorem shows that, up to polylogarithms, it is equivalent to the pair (K_2, K_3). A natural next step is to ask for analogues at higher moments, particularly in $G = \mathbb{F}_p^n$ where one can test and compute such quantities.

One possible definition is the fourth-moment control parameter

$$\kappa_4(A) := \inf\Big\{\kappa : \sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^4 \le \kappa\,|A|^3|B|^3 \text{ for all finite } B\Big\},$$

which normalises so that $\kappa_4(A)$ is scale-invariant under product-set heuristics (and matches the exponents obtained by the trivial bound $\|\mathbf{1}_A * \mathbf{1}_B\|_4 \le \|\mathbf{1}_A\|_{4/3}\|\mathbf{1}_B\|_{4/3}$). The quantity $\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^4$ counts 8-tuples $(a_i, b_i)_{i=1}^4$ with $a_1 + b_1 = \cdots = a_4 + b_4$, so its control is a higher-uniformity statement on the bipartite sum graph between A and B.

In parallel with $K_3(A)$, one might introduce a self-correlation hypothesis at level 4, for example

$$\|\mathbf{1}_A \circ \mathbf{1}_A\|_4^4 \le K_4\,|A|^5 \qquad \Big(\text{equivalently } \sum_d r_{A-A}(d)^4 \le K_4\,|A|^5\Big),$$

which implies a tail bound $|\{d : r_{A-A}(d) \ge t\}| \ll K_4|A|^5/t^4$ by the same counting argument as Lemma 2. One can then ask whether the combination of a uniform L^2 cross-correlation hypothesis (i) with such an L^4 self-correlation hypothesis yields a full L^4 control statement (with the correct dependence on K_2 and K_4, and perhaps with a manageable polylogarithmic loss).

Question 10.4 (weak $\Rightarrow$ full at L^4). In $G = \mathbb{F}_p^n$ (or in general abelian G), assume (i) with constant K_2 and assume additionally $\|\mathbf{1}_A \circ \mathbf{1}_A\|_4^4 \le K_4|A|^5$. Does it follow that

$$\sum_x (\mathbf{1}_A * \mathbf{1}_B(x))^4 \;\ll\; (\log(2|A|))^{O(1)}\,\max\{K_2^c, K_4\}\,|A|^3|B|^3$$

for some absolute exponent c (and optimally $c = 6$ or another explicit value dictated by scaling considerations)?

Even a partial result (restricted classes of B, or bounds that interpolate between L^3 and L^4) would be relevant for finite-field sum-product and for quantitative incidence estimates, where fourth moments often correspond to counting rectangles/parallelograms and to controlling ℓ^4 norms of Fourier transforms.

(3) Incidence bounds and operator norms. The weak hypothesis (i) is naturally an operator norm bound. Consider the linear operator T_A on finitely supported functions given by

$$T_A f := \mathbf{1}_A \circ f.$$

Proposition C may be viewed as a (polylogarithmically lossy) restricted-to-strong upgrade, asserting that the indicator testing bound (i) implies

$$\|T_A f\|_2 \ll (\log(2|\mathrm{supp}(f)|))^{O(1)}\,K_2\,|A|^{3/4}\,\|f\|_{4/3} \qquad (f \ge 0).$$

Similarly, full L^3 control can be written as a family of bounds for T_A acting on indicators after reflection, since $\mathbf{1}_A * \mathbf{1}_B = \mathbf{1}_A \circ \mathbf{1}_{-B}$. This suggests that the pair (K_2, K_3) is controlling two different aspects of the operator T_A: a global $\ell^{4/3} \to \ell^2$ mapping property and a self-interaction constraint at a higher moment.

A conceptual goal would be to formulate Theorem A as an interpolation principle for T_A with a "nonlinear endpoint" provided by (ii). Such a formulation might (a) clarify which parts of the argument are purely analytic and which are combinatorial, and (b) provide a path to removing logarithmic losses via real interpolation, Lorentz-space refinements, or sparse domination analogues in the discrete setting.

Question 10.5 (operator-norm reformulation). Is there a clean operator-theoretic statement equivalent (up to polylogarithms) to Theorem A, for instance an estimate of the form

$$\sup_{B\neq\emptyset} \frac{\|\mathbf{1}_A \circ \mathbf{1}_B\|_3^3}{|A|^2|B|^2} \lesssim \Phi\Big(\|T_A\|_{\ell^{4/3}\to\ell^2},\, \|\mathbf{1}_A \circ \mathbf{1}_A\|_3\Big),$$

with Φ explicitly comparable to $\max\{K_2^4, K_3\}$ after normalisation?

Finally, the incidence viewpoint remains relevant beyond $\mathbb{R}$. In groups where geometric incidence theorems exist (Euclidean settings via Szemerédi–Trotter, finite fields via point-line/point-plane incidence bounds in various ranges), L^3 convolution bounds are often equivalent to incidence estimates after an appropriate encoding. It would be useful to know to what extent weak control hypotheses (especially (i), which is an energy bound uniform in B) can be deduced from incidence input, and conversely whether full control can yield incidence statements for structured families of sets.

Question 10.6 (incidence mechanisms). In $\mathbb{F}_p^n$, can one derive non-trivial bounds on $K_2(A)$ or $K_3(A)$ for natural algebraic sets A (quadratic surfaces, graphs of polynomials, Cartesian products) using incidence theory, and can Theorem A then be used to propagate these bounds to L^3 convolution control with meaningful consequences (e.g. growth, expansion, or sum-product type estimates)?

We expect that progress on any of the questions above would sharpen the conceptual status of weak control parameters: whether they merely provide a technically convenient gateway to Bloom control, or whether they admit their own inverse theory and their own geometric/computational interpretations.