Quantitative Stability of Small Doubling Under Pushouts in FR_2^0

Liz Lemma　　Future Detective

January 16, 2026

Abstract

The Freiman categories FR_k and FR_k^0 introduced by Blanco–Haghverdi provide a categorical home for finite additive sets and Freiman homomorphisms; in FR_k^0 one has genuine limits and colimits, including pushouts. The source paper initiates a quantitative study of these (co)limits via the doubling constant but obtains only coarse cardinality bounds for pullbacks/quotients. We develop a quantitative "gluing theory" for $k = 2$: under a natural non-degeneracy hypothesis that the interface C is syndetic in each of A and B (bounded translate-covering), we prove that the pushout P of $A \leftarrow C \to B$ has doubling constant bounded polynomially in the original doubling parameter. Equivalently, bounded doubling is stable under categorical amalgamation along a large structured interface. The proof combines the explicit pushout model in FR_2^0 (as a quotient inside $G \oplus H$) with Ruzsa covering and Plünnecke-type inequalities, and reduces the general case to the subgroup case via passage to universal ambient groups. This provides a reusable quantitative tool for assembling approximate additive structure by categorical constructions, a prerequisite for functorial/categorical formulations of Freiman-type structure theorems.

Table of Contents

4. 4. The subgroup-interface model case: C a subgroup and i, j restrictions of homomorphisms; compute/estimate $\sigma(P)$ in the ambient quotient group $(G \oplus H)/N$.

5. 5. Quotient stability lemmas: bounding $|q(S+T)|$ in terms of $|S+T|$ and the fiber structure induced by N; controlling collapse using syndeticity and small doubling.

6. 6. Reduction to subgroup case via universal ambient groups: replace $(A, G), (B, H), (C, L)$ by Freiman-isomorphic diagrams in universal ambient groups where the cospan extends to group homomorphisms; show colimits respect these replacements.

7. 7. Proof of the main theorem for general C: assemble bounds from Sections 4–6; deduce uniform bounds for $P+P$ and $P-P$.

8. 8. Examples and near-sharpness: (i) gluing progressions along a large progression; (ii) failure without syndeticity (coproduct-like pushout along $\{0\}$); (iii) torsion examples and parameter dependence.

9. 9. Applications and extensions: pullback analogues, iterated gluing, and how this feeds into functorial Freiman models; list of open problems and refinements.

1 Introduction and statement of results

A recurring operation in additive combinatorics is to "glue" two structured finite sets along a common structured subconfiguration. In the most concrete form one has finite sets $A \subseteq G$ and $B \subseteq H$ with small doubling, together with a third set $C \subseteq L$ and Freiman maps $i : C \to A$, $j : C \to B$ describing an "interface" by which elements of A and B are to be identified. One would like a canonical object P that represents the result of identifying $i(c) \in A$ with $j(c) \in B$ for each $c \in C$, and to understand how additive parameters of P (notably doubling) depend on those of A, B, C.

The appropriate universal construction is the pushout of the cospan

$$(A, G) \xleftarrow{i} (C, L) \xrightarrow{j} (B, H)$$

computed in the normalized Freiman category FR_2^0. In this category, objects are finite normalized sets (containing 0) inside some ambient abelian group, and morphisms are Freiman homomorphisms of order 2 preserving 0. The pushout P is characterized by the usual categorical property: it comes equipped with morphisms $A \to P$ and $B \to P$ that agree on C (via i and j), and it is universal among such amalgamations. From the additive-combinatorial point of view, the value of having P is that it is intrinsic: it depends only on the Freiman-2 data and not on arbitrary choices of coordinates or ambient groups. In particular, one can meaningfully ask whether additive stability properties (e.g. polynomial bounds in the doubling constant) are preserved under this gluing operation.

However, without a non-degeneracy hypothesis on the interface C, small doubling is not stable under pushouts. The mechanism is transparent already in the easiest case: if the interface is trivial, then the pushout essentially forms a "wedge" of A and B inside a direct sum ambient group, and mixed sums create a large Cartesian product. Concretely, when $C = \{0\}$ and the maps are forced, the pushout set is

$$P = (A \times \{0\}) \cup (\{0\} \times B) \subseteq G \oplus H.$$

Then

$$P + P \supseteq (A \times \{0\}) + (\{0\} \times B) = A \times B,$$

and so $|P + P| \geq |A||B|$ even when $|A + A| \ll |A|$ and $|B + B| \ll |B|$. If A and B are, say, long arithmetic progressions, then $\sigma(A)$ and $\sigma(B)$ are $O(1)$ while $\sigma(P)$ grows like $\min\{|A|, |B|\}$. Thus, any general statement asserting that $\sigma(P)$ is bounded solely in terms of $\sigma(A), \sigma(B), \sigma(C)$ is false.

This failure indicates that the interface must control the mixed sumset contribution. In the wedge example, the mixed term $(A \times \{0\}) + (\{0\} \times B)$ is unconstrained and is responsible for the blow-up. Conceptually, to prevent such behavior the set C must be "large inside" both A and B, so that every element of A (respectively B) can be reached from $i(C)$ (respectively $j(C)$) by adding one of few possible offsets. We formalize this by assuming that $i(C)$ is syndetic in A and $j(C)$ is syndetic in B, in the concrete covering sense: A is contained in at most K translates of $i(C)$, and likewise B is contained in at most K translates of $j(C)$. This assumption is checkable and robust, and it is precisely what is needed to control mixed sums after amalgamation.

Under this hypothesis, our main result is that pushouts preserve small doubling up to a polynomial loss. More precisely, when A, B, and C each have doubling at most K, and when A and B are each covered by $O(K)$ translates of the corresponding copy of C, the pushout P has doubling at most K^c for an absolute exponent $c > 0$. In addition, the same exponent controls the first nontrivial sum and difference sets of P: the quantities $|P + P|$ and $|P - P|$ are bounded by $K^c|P|$. These estimates express a form of "polynomial stability under categorical amalgamation" that does not appear in the classical sumset toolkit, because the pushout operation is not visible at the level of ordinary subsets in a fixed ambient group.

The proof strategy is guided by an explicit model for the pushout. Writing $\widehat{A} := A \times \{0\}$ and $\widehat{B} := \{0\} \times B$ inside $G \oplus H$, we form the subgroup

$$N := \big\langle (i(c), -j(c)) : c \in C \big\rangle \leq G \oplus H$$

and the quotient map $q : G \oplus H \to (G \oplus H)/N$. Then the pushout set may be represented as

$$P = q(\widehat{A} \cup \widehat{B}).$$

This model allows us to study $P \pm P$ by splitting into three contributions:

$$q(\widehat{A} \pm \widehat{A}), \qquad q(\widehat{B} \pm \widehat{B}), \qquad q(\widehat{A} \pm \widehat{B}).$$

The first two terms inherit small doubling from A and B (since taking quotients cannot create more additive relations among elements already in $\widehat{A}$ or $\widehat{B}$ than those coming from N, and the relevant images remain controlled by $|A \pm A|$ and $|B \pm B|$). The central difficulty is the mixed term $q(\widehat{A}) + q(\widehat{B})$, which, in the degenerate interface case, produces $A \times B$.

Here the syndetic covering assumption becomes decisive. If $A \subseteq X + i(C)$ and $B \subseteq Y + j(C)$ for small sets $X \subseteq G$, $Y \subseteq H$, then after embedding into $G \oplus H$ and passing to the quotient one obtains translate-covering information of the form

$$q(\widehat{A}) \subseteq q(X \times \{0\}) + q(\widehat{i(C)}), \qquad q(\widehat{B}) \subseteq q(\{0\} \times Y) + q(\widehat{j(C)}),$$

where $q(\widehat{i(C)}) = q(\widehat{j(C)})$ by construction of N. Thus both $q(\widehat{A})$ and $q(\widehat{B})$ lie in few translates of a common "backbone" subset in the quotient. A direct sumset calculation then shows that the mixed sum $q(\widehat{A}) + q(\widehat{B})$ lies in few translates of $q(\widehat{i(C)}) + q(\widehat{i(C)})$, giving a bound that is polynomial in K once we know that $q(\widehat{i(C)})$ itself has small doubling in the quotient. In this way, the mixed term is controlled by the same structural object along which we glue.

The preceding outline is cleanest in the subgroup-interface regime, where C is itself a subgroup and the maps i and j are restrictions of homomorphisms; then $q(\widehat{i(C)})$ is a subgroup of the pushout ambient group, and the pushout set becomes a union of few cosets of this subgroup. In that case one can bound $|P+P|$ in terms of the number of cosets involved by elementary considerations, and obtain explicit dependence on the translate-covering numbers of A and B by the interface subgroup.

In the general Freiman setting, C need not be a subgroup, and the maps i, j need not extend to homomorphisms on the ambient groups. We address this by replacing the given diagram in FR_2^0 by an equivalent diagram in universal ambient groups. This replacement turns the Freiman-2 morphisms into genuine homomorphisms at the level of ambient groups while preserving the relevant sumset cardinalities at order 2. After this step, quotient arguments and subgroup methods apply as if we were in the subgroup-interface model case, at the cost of polynomial losses that are absorbed into K^c. We then transfer the estimates back to the original pushout using the invariance properties built into the universal ambient group construction.

The resulting theorem may be viewed as a categorical analogue of familiar stability phenomena: small doubling is stable under taking images by homomorphisms, under passing to large subsets, and under many forms of controlled union. The pushout combines several such operations at once: it is simultaneously a quotient in an ambient direct sum and a union of two embedded pieces. The novelty is that the gluing relation introduces new additive coincidences between A and B, and syndeticity is precisely what prevents these new coincidences from generating an uncontrolled mixed sumset.

Finally, we note that the syndetic covering formulation is not merely convenient but necessary at the level of hypotheses. Alternative "regularity" conditions that quantify expansion of $i(C)$ inside A can be vacuous in the normalized setting, since $0 \in i(C)$ implies $X \subseteq i(C) + X$ for every X. The translate-covering condition avoids this pitfall: it enforces that $i(C)$ and $j(C)$ are large enough inside A and B to serve as genuine interfaces, and it excludes precisely the degenerate pushouts that exhibit large mixed sumsets. With this non-degeneracy in place, pushouts in FR_2^0 become a controlled operation from the point of view of doubling, and one can treat them as a legitimate tool for constructing new approximate additive objects from old ones.

2 Preliminaries

We fix some basic conventions and recall the universal constructions that will be used throughout. All ambient groups are abelian and written additively.

For a finite nonempty set S in an abelian group we write

$$\sigma(S) := \frac{|S+S|}{|S|}$$

for the doubling constant, and we adopt the normalization convention $0 \in S$ whenever S is an object of our category.

2.1 Freiman 2-maps and the category FR_2^0

Let $A \subseteq G$ and $B \subseteq H$ be finite subsets of abelian groups with $0 \in A, B$. A map $\phi : A \to B$ is a *Freiman homomorphism of order* 2 if for all $a_1, a_2, a_3, a_4 \in A$ with

$$a_1 + a_2 = a_3 + a_4$$

in G, we have

$$\phi(a_1) + \phi(a_2) = \phi(a_3) + \phi(a_4)$$

in H. We say that ϕ is *normalized* if $\phi(0) = 0$. A *Freiman* 2*-isomorphism* is a bijection $\phi : A \to B$ such that both ϕ and ϕ^{-1} are Freiman homomorphisms of order 2. In particular, if ϕ is a Freiman 2-isomorphism, then it preserves all additive relations of the form $a_1 + a_2 = a_3 + a_4$ and hence preserves cardinalities of sumsets and difference sets at "order 2" (e.g. $|A \pm A| = |B \pm B|$ after identification).

The normalized Freiman category FR_2^0 has as objects pairs (A, G) where G is an abelian group and $A \subseteq G$ is finite, nonempty, and contains 0. Morphisms $(A, G) \to (B, H)$ are normalized Freiman homomorphisms of order 2, i.e. maps $\phi : A \to B$ with $\phi(0) = 0$ preserving all relations $a_1 + a_2 = a_3 + a_4$. We will freely suppress the ambient group from notation when it is clear, but it is useful to keep in mind that morphisms are *not* required to extend to group homomorphisms between G and H.

For later use we record the elementary but crucial observation that Freiman 2-maps preserve the structure of "two-term" sum/difference sets: if $\phi : A \to B$ is a Freiman 2-homomorphism, then $\phi(a_1) + \phi(a_2)$ depends only on $a_1 + a_2$, and similarly $\phi(a_1) - \phi(a_2)$ depends only on $a_1 - a_2$. Thus ϕ induces well-defined maps on $A + A$ and $A - A$, and consequently

$$|\phi(A) \pm \phi(A)| \leq |A \pm A|.$$

When ϕ is injective, $|\phi(A)| = |A|$, so bounds on $|A \pm A|$ transfer directly to $\phi(A)$.

2.2 Sumset notation

For finite subsets S, T of an abelian group we write

$$S + T := \{s + t : s \in S,\ t \in T\}, \qquad S - T := \{s - t : s \in S,\ t \in T\}.$$

For integers $\ell, m \geq 0$ we use the standard shorthand

$$\ell S - mS := \{s_1 + \cdots + s_\ell - t_1 - \cdots - t_m : \ s_i, t_j \in S\},$$

with $0S = \{0\}$. In particular, $2S = S + S$ and $1S - 1S = S - S$. We emphasize that in this paper only the regimes $\ell + m \leq 2$ are directly forced by Freiman 2-information; higher iterated sumsets will only appear after we pass to auxiliary ambient groups where genuine homomorphisms exist.

We will repeatedly use the monotonicity of sumsets under homomorphisms and quotients: if $\pi : G \to G'$ is a group homomorphism and $S, T \subseteq G$ are finite, then

$$\pi(S) + \pi(T) = \pi(S + T), \qquad \pi(S) - \pi(T) = \pi(S - T),$$

and therefore $|\pi(S \pm T)| \leq |S \pm T|$. In particular, taking quotients cannot increase the cardinality of a sumset.

2.3 Explicit pushouts in FR_2^0

We recall the concrete model for pushouts in FR_2^0 (Proposition 4.10 in the source), specialized to the normalized situation. Suppose we are given a cospan

$$(A, G) \xleftarrow{i} (C, L) \xrightarrow{j} (B, H)$$

in FR_2^0, where i and j are injective and preserve 0. We form the embedded copies

$$\widehat{A} := A \times \{0\} \subseteq G \oplus H, \qquad \widehat{B} := \{0\} \times B \subseteq G \oplus H,$$

and we consider the subgroup

$$N := \big\langle (i(c), -j(c)) : \ c \in C \big\rangle \leq G \oplus H.$$

Let $q : G \oplus H \to (G \oplus H)/N$ be the quotient homomorphism. The pushout ambient group is $(G \oplus H)/N$, and the pushout set can be taken to be

$$P := q(\widehat{A} \cup \widehat{B}) \subseteq (G \oplus H)/N.$$

The canonical morphisms $A \to P$ and $B \to P$ are the restrictions of q to $\widehat{A}$ and $\widehat{B}$, respectively. The relation identifying $i(c) \in A$ with $j(c) \in B$ is encoded by the fact that

$$q(i(c), 0) = q(0, j(c)) \qquad \text{for all } c \in C,$$

since $(i(c), 0) - (0, j(c)) = (i(c), -j(c)) \in N$.

For sumset estimates we will repeatedly use the trivial decomposition

$$P \pm P \subseteq q(\widehat{A} \pm \widehat{A}) \ \cup \ q(\widehat{A} \pm \widehat{B}) \ \cup \ q(\widehat{B} \pm \widehat{B}),$$

and the corresponding cardinality bound

$$|P \pm P| \leq |q(\widehat{A} \pm \widehat{A})| + |q(\widehat{A} \pm \widehat{B})| + |q(\widehat{B} \pm \widehat{B})|.$$

The terms $q(\widehat{A} \pm \widehat{A})$ and $q(\widehat{B} \pm \widehat{B})$ are immediate images of $A \pm A$ and $B \pm B$, and so are bounded by $|A \pm A|$ and $|B \pm B|$. The mixed term $q(\widehat{A} \pm \widehat{B})$ is the main new feature created by the pushout; controlling it is exactly where non-degeneracy of the interface enters in the next section.

We also record the identity of the glued copy of C in the quotient. Writing $\widehat{i(C)} := i(C) \times \{0\}$ and $\widehat{j(C)} := \{0\} \times j(C)$, we have

$$q(\widehat{i(C)}) = q(\widehat{j(C)}),$$

and we will denote this common subset by

$$Q := q(\widehat{i(C)}) = q(\widehat{j(C)}) \subseteq (G \oplus H)/N.$$

In particular, any translate-covering information for A by $i(C)$ or for B by $j(C)$ descends to translate-covering information for $q(\widehat{A})$ and $q(\widehat{B})$ by the same set Q.

2.4 Ruzsa distance and covering tools

We will use standard consequences of small doubling in the form of Ruzsa-type inequalities. For finite nonempty sets S, T in an abelian group, the (logarithmic) *Ruzsa distance* is

$$d(S,T) := \log \frac{|S - T|}{|S|^{1/2}|T|^{1/2}}.$$

It satisfies the triangle inequality

$$d(S,U) \leq d(S,T) + d(T,U),$$

which is equivalent to the Ruzsa triangle inequality

$$|S - U| \leq \frac{|S - T|\ |T - U|}{|T|}.$$

In the special case $S = T = U$ this yields no information, but in combination with small doubling hypotheses it allows one to convert bounds on one sumset into bounds on related sumsets.

We also recall the Ruzsa covering lemma in a form tailored to our later applications. If A, B are finite nonempty subsets of an abelian group with

$$|A + B| \leq K|A|,$$

then there exists a subset $X \subseteq B$ with $|X| \le K$ such that

$$B \subseteq X + (A - A).$$

We will use this both in ambient groups and in quotient groups; the statement is invariant under homomorphisms in the sense that if π is a homomorphism then $|\pi(A) + \pi(B)| \le |A + B|$, so any covering obtained upstairs yields a covering downstairs after applying π.

Finally, we will use basic Plünnecke–Ruzsa estimates to control low-order iterated sumsets. For example, if $|A + A| \le K|A|$, then

$$|A - A| \le K^2|A|$$

and more generally (by standard Plünnecke inequalities)

$$|\ell A - mA| \le K^{\ell+m}|A|$$

for fixed ℓ, m when the ambient setting supports the usual combinatorial graph argument. Since our primary setting is FR_2^0, we will only invoke such bounds either for $\ell + m \le 2$ directly, or after passing to auxiliary ambient groups in which the relevant Freiman maps become genuine homomorphisms. In particular, the reader may safely interpret any appearance of $K^{O(1)}$ as arising from repeated applications of these classical inequalities together with the fact that quotient maps do not increase sumset cardinalities.

These preliminaries reduce most of our later work to two tasks: (i) expressing the pushout set P and its first sum/difference sets inside the quotient $(G \oplus H)/N$, and (ii) converting structural hypotheses (translate coverings and small doubling) into explicit bounds on the mixed sumset contribution. The second task is where the non-degeneracy conditions on the interface will enter.

3 Non-degeneracy hypotheses for the interface

The pushout construction necessarily creates a "mixed" contribution $q(\widehat{A} \pm \widehat{B})$ to $P \pm P$. Without further assumptions on how C sits inside A and B, this mixed term can be as large as $|A||B|$ (cf. Proposition D), and there is no hope for a polynomial bound on $\sigma(P)$ in terms of $\sigma(A), \sigma(B), \sigma(C)$ alone. We therefore impose a non-degeneracy hypothesis on the embeddings i, j which formalizes the idea that C is a large "backbone" inside each of A and B. In this paper we adopt a translate-covering (syndeticity) condition, which is simple to verify and behaves well under quotients, and we briefly compare it with other possible regularity notions.

3.1 Translate coverings and syndetic embeddings

Let S, T be finite subsets of an abelian group G with $0 \in T$. We define the *translate-covering number* of S by T to be

$$\kappa(S;T) \ := \ \min\bigl\{|X| : \ X \subseteq G, \ S \subseteq X + T\bigr\}.$$

Thus $\kappa(S;T) \leq M$ means that S is covered by at most M translates of T. When $T = i(C)$ (or $T = j(C)$) we interpret this as a quantitative non-degeneracy of the embedding.

Definition 3.1 (Syndetic embedding)**.** An injective normalized Freiman 2-map $i : (C, L) \to (A, G)$ is *K-syndetic* if $\kappa(A; i(C)) \leq K$, i.e. if there exists $X \subseteq G$ with $|X| \leq K$ such that $A \subseteq X + i(C)$. Similarly j is K-syndetic if $\kappa(B; j(C)) \leq K$.

We emphasize two basic features of this condition.

(i) It prevents collapse to the coproduct regime. If $C = \{0\}$ then $i(C) = \{0\}$ and $\kappa(A; i(C)) = |A|$, so K-syndeticity forces $|A| \leq K$; in particular it rules out precisely the degenerate gluing in which the mixed sumset dominates.

(ii) It is stable under the operations used in the pushout argument. If $A \subseteq X + i(C)$ then automatically

$$\widehat{A} \subseteq (X \times \{0\}) + \widehat{i(C)} \subseteq G \oplus H,$$

and applying the quotient map $q : G \oplus H \to (G \oplus H)/N$ yields the covering

$$q(\widehat{A}) \subseteq q(X \times \{0\}) + Q,$$

where $Q = q(\widehat{i(C)}) = q(\widehat{j(C)})$. This is exactly the form needed to control $q(\widehat{A} \pm \widehat{B})$ by Lemma 3. In particular, the non-degeneracy hypothesis is designed to descend cleanly to the quotient group where P lives.

We record a few immediate consequences which will be used implicitly. If $A \subseteq X + i(C)$ with $|X| \leq K$, then

$$|A| \leq |X| \cdot |i(C)| \leq K|C|. \tag{1}$$

Moreover, for any choice of signs one has the crude containment

$$A \pm A \ \subseteq \ (X \pm X) + \bigl(i(C) \pm i(C)\bigr), \tag{2}$$

and hence

$$|A \pm A| \ \leq \ |X \pm X| \cdot |i(C) \pm i(C)|.$$

Since $|X \pm X| \leq |X|^2 \leq K^2$ and $|i(C) \pm i(C)| \leq |C \pm C|$, the small-doubling hypothesis on C yields bounds of the form $|A \pm A| \leq K^{O(1)}|C|$, consistent with $\sigma(A) \leq K$. These estimates are not sharp and are not used as input; they simply illustrate that syndeticity ties the sizes of A and C together and prevents extremely sparse embeddings.

3.2 Equivalent viewpoints and quick tests

The definition $\kappa(A; i(C)) \leq K$ is the most direct and is the one we assume. Nevertheless, it is useful to keep in mind several equivalent or near-equivalent formulations.

(1) Coset-index intuition in the subgroup case. If $i(C) \leq G$ happens to be a subgroup and A is a union of M cosets of $i(C)$, then $\kappa(A; i(C)) = M$. Thus in the subgroup-interface model case, K-syndeticity simply means that A meets only $O_K(1)$ cosets of the glued subgroup. This is the regime treated explicitly in the next section, where one can replace translate coverings by genuine quotient considerations.

(2) A partial converse: large intersections with some translate. If $A \subseteq X + i(C)$ with $|X| = M$, then by pigeonhole there exists $x \in X$ with

$$|A \cap (x + i(C))| \ \geq \ \frac{|A|}{M}.$$

Thus syndeticity implies that A has a translate of $i(C)$ capturing a positive proportion ($\geq 1/K$) of its elements. This is a convenient quick sanity check in examples: if every translate of $i(C)$ intersects A in $o(|A|)$ points, then $\kappa(A; i(C))$ must be large.

(3) Relative growth tests. A commonly occurring regularity condition in additive combinatorics is that a pair (A, T) has small sumset, e.g. $|A+T| \leq M|A|$ or $|A+T| \leq M|T|$. Such conditions do not literally imply $A \subseteq X+T$, but they often lead to coverings after passing to difference sets. For instance, if $|A + i(C)| \leq M|i(C)|$, then by the Ruzsa covering lemma there exists $X \subseteq A$ with $|X| \leq M$ such that

$$A \ \subseteq \ X + \big(i(C) - i(C)\big).$$

If $i(C)$ is a subgroup then $i(C) - i(C) = i(C)$ and this recovers syndeticity. In general $i(C) - i(C)$ may be larger than $i(C)$, but under $\sigma(C) \leq K$ we still have $|i(C) - i(C)| \leq K^{O(1)}|C|$, so coverings by $i(C) - i(C)$ can sometimes be upgraded to workable bounds in the quotient. We do not pursue this systematically here, preferring the simpler direct hypothesis $A \subseteq X + i(C)$, but it is worth noting that syndeticity can often be verified by establishing a suitable small relative growth estimate.

(4) A quotient-size heuristic. Even outside the subgroup case, one may think of $\kappa(A; i(C))$ as a proxy for the size of a quotient $A/i(C)$: a small covering number means that the "transversal complexity" of A mod the $i(C)$-direction is bounded. This is precisely the information required to control mixed sums in the pushout, since mixed sums are sensitive to how many different translates of the glued piece are present in A and B.

3.3 Comparison with alternative regularity notions

We briefly discuss other possible non-degeneracy assumptions and explain why we choose syndeticity.

Density assumptions. One might require that $i(C)$ has positive density inside A, e.g. $|i(C)| \geq \delta|A|$ for some $\delta = \delta(K) > 0$. This is weaker than syndeticity: a dense subset need not cover A by few translates. In particular, density does not prevent A from splitting into many far-separated translates of $i(C)$, which is exactly what feeds the mixed term in the pushout. Moreover, density is not stable under Freiman isomorphisms in any useful quantitative sense, whereas translate coverings are.

Energy assumptions. Another option is to assume large additive energy between A and $i(C)$, for example

$$E(A, i(C)) := \left|\{(a_1, a_2, c_1, c_2) \in A^2 \times C^2 : a_1 + i(c_1) = a_2 + i(c_2)\}\right| \geq \frac{1}{M}|A|^2|C|$$

for some $M = K^{O(1)}$. Such a hypothesis does encode that A correlates additively with $i(C)$, and by Balog–Szemerédi–Gowers type arguments it can yield structured subsets and translate-coverings. However, invoking BSG introduces additional layers (passing to large subsets, losing parameters, and tracking the interaction with the pushout quotient). Since our aim is a robust polynomial bound under transparent hypotheses, we avoid an energy-based formulation.

Small relative growth and "regularity" in the original formulation. One may attempt to impose conditions of the form $|i(C)+X| \geq |X|/K$ for all $X \subseteq A$, meant to prevent collapse. In our normalized setting this condition is vacuous: since $0 \in i(C)$, we always have $X \subseteq i(C) + X$ and hence $|i(C) + X| \geq |X|$. More generally, any regularity notion must genuinely constrain how many $i(C)$-translates are needed to cover A, not merely prevent $i(C)$ from shrinking sets.

Why syndeticity is the right hypothesis for pushouts. The pushout glues $i(c)$ and $j(c)$ and then allows sums across the two sides. The resulting mixed term is controlled precisely by knowing that each side lies in few translates of the glued subset Q after passing to the quotient. Syndeticity is exactly the minimal hypothesis that provides this control with no additional machinery: it is checkable on the nose in the original ambient groups and descends immediately to the quotient where the pushout is computed.

Having fixed syndeticity as our non-degeneracy condition, we now turn to the most rigid setting, namely the subgroup-interface model case, where C is a subgroup and i, j arise from homomorphisms. In that regime the translate-covering numbers become literal coset counts, and one can compute the pushout inside the ambient quotient group $(G \oplus H)/N$ with explicit bounds.

4 The subgroup-interface model case

We now isolate the most rigid regime, in which the interface is an actual subgroup and the Freiman data comes from honest homomorphisms. In this case the pushout ambient group is the usual abelian pushout, and the pushout *set* can be controlled by elementary coset combinatorics.

4.1 Standing hypotheses in the model case

Assume that $C \leq L$ is a subgroup, and that i, j are restrictions of injective group homomorphisms on C. In particular $i(C) \leq G$ and $j(C) \leq H$ are subgroups, and the subgroup

$$N \ := \ \big\langle (i(c), -j(c)) : \ c \in C \big\rangle \ \leq \ G \oplus H$$

is exactly the image of C under the homomorphism $c \mapsto (i(c), -j(c))$. We continue to write $q : G \oplus H \to (G \oplus H)/N$ for the quotient map and

$$P \ := \ q(\widehat{A} \cup \widehat{B}), \qquad \widehat{A} = A \times \{0\}, \ \widehat{B} = \{0\} \times B,$$

for a concrete model of the pushout set.

A key simplification in this regime is that the quotient does not "collapse" either side internally.

Lemma 4.1 (No internal collapse on each side)**.** *In the subgroup-interface case one has*

$$N \cap (G \times \{0\}) = \{(0,0)\}, \qquad N \cap (\{0\} \times H) = \{(0,0)\}.$$

Consequently, q is injective on $G \times \{0\}$ and on $\{0\} \times H$, and in particular on $\widehat{A}$ and on $\widehat{B}$.

Proof. Suppose $(x, 0) \in N$. Then $(x, 0) = \sum_{r=1}^{m} (i(c_r), -j(c_r)) = (i(c), -j(c))$ where $c := \sum_r c_r \in C$. The second coordinate gives $j(c) = 0$, hence $c = 0$ by injectivity of $j|_C$, and therefore $x = i(0) = 0$. The argument for $(0, y) \in N$ is identical. □

One immediate consequence is that sumsets on each side pass through the quotient without distortion:

$$|q(\widehat{A} \pm \widehat{A})| \ = \ |\widehat{A} \pm \widehat{A}| \ = \ |A \pm A|, \qquad |q(\widehat{B} \pm \widehat{B})| \ = \ |B \pm B|. \tag{3}$$

Thus the only genuinely new term in $P \pm P$ is the mixed contribution $q(\widehat{A} \pm \widehat{B})$, as anticipated.

4.2 The glued subgroup and a size formula for P

Let

$$Q \;:=\; q(\widehat{i(C)}) \;=\; q(\widehat{j(C)}).$$

In this subgroup-interface case, Q is a *subgroup* of $(G \oplus H)/N$, and $|Q| = |C|$ by Lemma 4.1 (since q is injective on $\widehat{i(C)} \subseteq G \times \{0\}$).

Moreover, the intersection $q(\widehat{A}) \cap q(\widehat{B})$ is exactly Q.

Lemma 4.2 (Intersection is the glued part)**.** *One has* $q(\widehat{A}) \cap q(\widehat{B}) = Q$. *In particular,*

$$|P| \;=\; |q(\widehat{A}) \cup q(\widehat{B})| \;=\; |A| + |B| - |C|.$$

Proof. The inclusion $Q \subseteq q(\widehat{A}) \cap q(\widehat{B})$ is clear from $\widehat{i(C)} \subseteq \widehat{A}$ and $\widehat{j(C)} \subseteq \widehat{B}$.

Conversely, suppose $q(a, 0) = q(0, b)$ with $a \in A$ and $b \in B$. Then $(a, -b) \in N$, so $(a, -b) = (i(c), -j(c))$ for some $c \in C$. Hence $a = i(c)$ and $b = j(c)$, so the common value is $q(i(c), 0) \in Q$.

The cardinality identity follows from Lemma 4.1, which implies $|q(\widehat{A})| = |A|$ and $|q(\widehat{B})| = |B|$, together with $|q(\widehat{A}) \cap q(\widehat{B})| = |Q| = |C|$. □

While we will not rely heavily on the explicit formula $|P| = |A|+|B|-|C|$, it is useful as a sanity check: in the subgroup case, the pushout set has the "expected" size, and the quotient map does not introduce large identifications beyond the intended gluing along C.

4.3 Syndeticity becomes a coset-count, and P is a union of few cosets

Because $i(C) \leq G$ is a subgroup, a translate covering $A \subseteq X + i(C)$ may be interpreted as saying that A lies in the union of $|X|$ cosets of $i(C)$. Let

$$M_A \;:=\; \kappa(A; i(C)), \qquad M_B \;:=\; \kappa(B; j(C)),$$

and fix witnesses $X \subseteq G$, $Y \subseteq H$ with $|X| = M_A$, $|Y| = M_B$ such that

$$A \subseteq X + i(C), \qquad B \subseteq Y + j(C).$$

Pushing this information into the quotient gives a particularly transparent description of P.

Lemma 4.3 (Coset description in the pushout quotient)**.** *With* $Q = q(\widehat{i(C)})$, *one has*

$$q(\widehat{A}) \subseteq q(X \times \{0\}) + Q, \qquad q(\widehat{B}) \subseteq q(\{0\} \times Y) + Q,$$

and hence

$$P \;\subseteq\; T + Q, \qquad T \;:=\; q(X \times \{0\}) \;\cup\; q(\{0\} \times Y),$$

with $|T| \leq M_A + M_B$.

Proof. If $a \in A$ then $a = x + i(c)$ for some $x \in X$ and $c \in C$. Hence

$$q(a,0) \;=\; q(x+i(c),0) \;=\; q(x,0)+q(i(c),0) \;\in\; q(X\times\{0\})+Q.$$

The argument for $q(\widehat{B})$ is the same. The final claim follows by taking unions. $\square$

Thus, in the subgroup-interface case, P is contained in the union of at most $M_A + M_B$ cosets of the subgroup Q. Since Q is a subgroup, it has doubling constant 1, and coset combinatorics immediately control the doubling of any set contained in few cosets of Q.

4.4 Doubling bounds for P in the subgroup-interface case

We now give an explicit estimate for $\sigma(P)$ in terms of the translate-covering numbers, thereby proving the model-case bound advertised earlier.

Proposition 4.4 (Subgroup-interface pushouts are stable)**.** *Assume the subgroup-interface hypotheses. Let $M_A = \kappa(A; i(C))$ and $M_B = \kappa(B; j(C))$. Then*

$$|P+P| \;\le\; (M_A+M_B)^2\,|Q| \;\le\; (M_A+M_B)^2\,|P|,$$

and similarly

$$|P-P| \;\le\; (M_A+M_B)^2\,|Q| \;\le\; (M_A+M_B)^2\,|P|.$$

In particular,

$$\sigma(P) \;=\; \frac{|P+P|}{|P|} \;\le\; (M_A+M_B)^2.$$

Under the standing K-syndeticity hypothesis (so $M_A, M_B \le K$) this yields $\sigma(P) \le (2K)^2$.

Proof. By Lemma 4.3 we have $P \subseteq T + Q$ with $|T| \le M_A + M_B$. Since Q is a subgroup, $(T+Q)+(T+Q) = (T+T)+Q$, whence

$$P+P \;\subseteq\; (T+Q)+(T+Q) \;=\; (T+T)+Q.$$

Therefore

$$|P+P| \;\le\; |(T+T)+Q| \;\le\; |T+T|\cdot|Q| \;\le\; |T|^2\,|Q| \;\le\; (M_A+M_B)^2\,|Q|.$$

Because $Q \subseteq P$ (indeed $Q \subseteq q(\widehat{i(C)}) \subseteq q(\widehat{A}) \subseteq P$), we have $|Q| \le |P|$, giving the stated bound in terms of $|P|$. The estimate for $P-P$ is identical, using $P-P \subseteq (T-T)+Q$ and $|T-T| \le |T|^2$. $\square$

Several remarks are in order.

First, the argument uses only the *covering* information and the fact that Q is a subgroup; it does not require the small-doubling hypotheses on A, B, C. This reflects the fact that in the subgroup case the dominant obstruction to small doubling of P is transversal complexity (how many cosets occur), not internal additive structure within each coset.

Second, Lemma 4.1 and Lemma 4.2 highlight why this model case is substantially cleaner than the general Freiman setting: there is no collapse of $\widehat{A}$ or $\widehat{B}$ in the quotient, so the only nontrivial bookkeeping concerns the mixed term, which is handled by the coset-covering description.

Finally, the coset computation above is the template for the general case. Outside the subgroup regime, $Q = q(\widehat{i(C)})$ need not be a subgroup and the quotient map q may identify many points even within $\widehat{A}$ or $\widehat{B}$. In the next stage of the argument we will therefore develop quotient-stability tools that bound $|q(S+T)|$ and control the size of fibers of q using syndeticity and small doubling.

5 Quotient stability lemmas

We now return to the general Freiman setting, where C need not be a subgroup and i, j need not extend to homomorphisms on the ambient groups. The quotient homomorphism

$$q : G \oplus H \to (G \oplus H)/N$$

is always a genuine group homomorphism, but unlike the subgroup-interface model case it may have large fibers on $\widehat{A}$ or $\widehat{B}$. In particular, even if $|A \pm A|$ is small, the image $q(\widehat{A})$ may be much smaller than $\widehat{A}$, and we must understand how such "collapse" can and cannot occur under our syndeticity and small-doubling hypotheses.

The guiding principle is that the only identifications we can afford are those controlled by the glued part $q(\widehat{i(C)}) = q(\widehat{j(C)})$. We therefore isolate several elementary lemmas that bound sizes of quotient-images by tracking fiber multiplicities, and we formulate a bounded-collapse criterion under which the desired polynomial doubling estimates follow formally.

5.1 Fibers of a quotient map

Let Γ be an abelian group, let $N \leq \Gamma$ be a subgroup, and let $q : \Gamma \to \Gamma/N$ be the quotient map. For a finite set $S \subseteq \Gamma$ we define the (maximal) fiber multiplicity

$$\mu_q(S) := \max_{y \in q(S)} |S \cap q^{-1}(y)| = \max_{x \in \Gamma} |S \cap (x + N)|.$$

Then trivially

$$|q(S)| \;\geq\; \frac{|S|}{\mu_q(S)}. \tag{4}$$

The following simple observation relates $\mu_q(S)$ to the intersection of N with the difference set $S - S$.

Lemma 5.1 (A fiber sits in a translate of $(S-S)\cap N$)**.** *Let $S \subseteq \Gamma$ be finite. Then*

$$\mu_q(S) \;\leq\; |(S-S)\cap N|.$$

Proof. Fix a coset $x+N$, and let $F := S\cap(x+N)$. If $F=\emptyset$ there is nothing to prove. Otherwise choose $s_0 \in F$. For any $s \in F$ we have $s - s_0 \in N$ (since $s, s_0 \in x+N$) and also $s - s_0 \in S - S$. Thus the map $s \mapsto s - s_0$ sends F injectively into $(S-S)\cap N$. Hence $|F| \leq |(S-S)\cap N|$. Taking the maximum over cosets gives the claim. □

In our application $\Gamma = G \oplus H$, and S will be one of $\widehat{A}, \widehat{B}, \widehat{i(C)}, \widehat{j(C)}$ or a sumset built from them. Lemma 5.1 provides a convenient, purely set-theoretic way to measure collapse: large fibers force many differences of elements of S to lie in N.

5.2 Propagation of multiplicity through translate coverings

The syndeticity hypothesis gives translate coverings $A \subseteq X + i(C)$ and $B \subseteq Y + j(C)$ with $|X|, |Y| \leq K$. We will repeatedly use the fact that bounded multiplicity on the backbone $i(C)$ propagates to bounded multiplicity on all of A.

Lemma 5.2 (Multiplicity under a translate cover)**.** *Let $S, R, U \subseteq \Gamma$ be finite, and suppose $S \subseteq U + R$. Then*

$$\mu_q(S) \;\leq\; |U|\,\mu_q(R).$$

Proof. Fix a coset $x+N$. Then

$$S\cap(x+N) \;\subseteq\; (U+R)\cap(x+N) \;=\; \bigcup_{u\in U}\Big((u+R)\cap(x+N)\Big).$$

For each fixed u, we have $(u+R)\cap(x+N) = u + \big(R\cap((x-u)+N)\big)$, whose cardinality is at most $\mu_q(R)$ by definition. Therefore $|S\cap(x+N)| \leq |U|\,\mu_q(R)$. Taking the maximum over x yields the claim. □

Specializing to our situation (with $\Gamma = G \oplus H$), we will apply this to $S = \widehat{A}$, $R = \widehat{i(C)}$, and $U = X \times \{0\}$, and similarly on the B-side. Thus, any quantitative control on $\mu_q(\widehat{i(C)})$ immediately yields quantitative control on $\mu_q(\widehat{A})$, and likewise for B.

5.3 A bounded-collapse hypothesis and its consequences

We record a convenient hypothesis that isolates precisely what we will need from the quotient map q.

Definition 5.3 (Bounded collapse along the interface)**.** Let

$$Q \;:=\; q(\widehat{i(C)}) \;=\; q(\widehat{j(C)}) \;\subseteq (G \oplus H)/N.$$

We say that the interface collapses by at most M if

$$\mu_q(\widehat{i(C)}) \;\le\; M \qquad \text{and} \qquad \mu_q(\widehat{j(C)}) \;\le\; M.$$

In the subgroup-interface model case, Lemma 4.1 gives $M = 1$. In general, M need not be 1, but the point of the next lemma is that once M is known to be $K^{O(1)}$, the desired polynomial bounds for $P \pm P$ follow by a formal computation.

Lemma 5.4 (Size and doubling of the glued part under bounded collapse)**.** *Assume $\mu_q(\widehat{i(C)}) \le M$. Then*

$$|Q| \;\ge\; \frac{|C|}{M}, \qquad |Q{+}Q| \;\le\; |C{+}C| \;\le\; K|C| \qquad \textit{and hence} \qquad \sigma(Q) \;\le\; KM.$$

The same conclusions hold with i replaced by j.

Proof. Since $|\widehat{i(C)}| = |C|$ and $q(\widehat{i(C)}) = Q$, (4) gives $|Q| \ge |C|/M$.

Moreover $Q + Q = q(\widehat{i(C)}) + q(\widehat{i(C)}) = q(\widehat{i(C)} + \widehat{i(C)})$, so $|Q + Q| \le |\widehat{i(C)} + \widehat{i(C)}| = |i(C) + i(C)|$. Because i is a Freiman-2 homomorphism and injective on C, we have $|i(C) + i(C)| \le |C + C| \le K|C|$. Dividing by $|Q| \ge |C|/M$ yields $\sigma(Q) \le KM$. □

Next we relate $|P|$ to $|A|$ and $|B|$ under the same bounded-collapse hypothesis.

Lemma 5.5 (Lower bounds for $|q(\widehat{A})|$ and $|q(\widehat{B})|$)**.** *Assume $A \subseteq X + i(C)$ and $B \subseteq Y + j(C)$ with $|X|, |Y| \le K$, and assume $\mu_q(\widehat{i(C)}) \le M$ and $\mu_q(\widehat{j(C)}) \le M$. Then*

$$|q(\widehat{A})| \;\ge\; \frac{|A|}{KM}, \qquad |q(\widehat{B})| \;\ge\; \frac{|B|}{KM}, \qquad \textit{and in particular} \qquad |P| \;\ge\; \max\Big\{\frac{|A|}{KM}, \frac{|B|}{KM}\Big\}.$$

Proof. We view $\widehat{A} \subseteq (X \times \{0\}) + \widehat{i(C)}$. By Lemma 5.2 we have $\mu_q(\widehat{A}) \le |X|\, \mu_q(\widehat{i(C)}) \le KM$. Applying (4) gives $|q(\widehat{A})| \ge |\widehat{A}|/(KM) = |A|/(KM)$. The argument for $q(\widehat{B})$ is identical. Finally $P = q(\widehat{A} \cup \widehat{B})$ contains both $q(\widehat{A})$ and $q(\widehat{B})$, giving the stated lower bound for $|P|$. □

We now combine these ingredients to obtain quotient-stable bounds for $P \pm P$. The proof is the same decomposition as in the subgroup-interface case, except that we keep track of multiplicities via Lemmas 5.4 and 5.5.

Proposition 5.6 (Doubling of the pushout under bounded collapse)**.** *Assume the standing hypotheses (small doubling for A, B, C and K-syndeticity of i, j), and assume in addition that the interface collapses by at most M in the sense of Definition 5.3. Then*

$$|P + P| \;\le\; K^{O(1)}\, M^{O(1)}\, |P| \qquad \textit{and} \qquad |P - P| \;\le\; K^{O(1)}\, M^{O(1)}\, |P|.$$

In particular $\sigma(P) \le K^{O(1)} M^{O(1)}$.

Proof. Write $Q = q(\widehat{i(C)}) = q(\widehat{j(C)})$. By syndeticity and the quotient-covering observation (Lemma 2 in the global list), we may choose $U_A, U_B \subseteq (G \oplus H)/N$ with $|U_A|, |U_B| \le K$ such that

$$q(\widehat{A}) \subseteq U_A + Q, \qquad q(\widehat{B}) \subseteq U_B + Q.$$

For the mixed term we then have

$$q(\widehat{A}) + q(\widehat{B}) \;\subseteq\; (U_A + U_B) + (Q + Q),$$

so

$$|q(\widehat{A}) + q(\widehat{B})| \;\le\; |U_A + U_B|\, |Q + Q| \;\le\; K^2\, |Q + Q|. \tag{5}$$

By Lemma 5.4, $|Q + Q| \le KM\, |Q| \le KM\, |P|$, since $Q \subseteq P$. Substituting into (5) gives

$$|q(\widehat{A}) + q(\widehat{B})| \;\le\; K^3 M\, |P|.$$

For the pure terms we use only small doubling of A and B, together with the lower bound $|P| \gg |A|/(KM)$ from Lemma 5.5. Indeed,

$$|q(\widehat{A} + \widehat{A})| \;\le\; |\widehat{A} + \widehat{A}| \;=\; |A + A| \;\le\; K|A| \;\le\; K^2 M\, |P|,$$

and similarly $|q(\widehat{B} + \widehat{B})| \le K^2 M\, |P|$. Since

$$P + P \;\subseteq\; q(\widehat{A} + \widehat{A}) \;\cup\; q(\widehat{A} + \widehat{B}) \;\cup\; q(\widehat{B} + \widehat{B}),$$

we obtain $|P + P| \le K^{O(1)} M\, |P|$, which is of the asserted form.

The bound for $P - P$ is identical: we replace $+$ by $-$ throughout, use $|A - A| \le K^{O(1)}|A|$ and $|B - B| \le K^{O(1)}|B|$ (a standard consequence of $\sigma(A), \sigma(B) \le K$), and use Lemma 5.4 to bound $|Q - Q| \ll_K |Q + Q| \le KM|Q| \le KM|P|$. □

Proposition 5.6 reduces the main problem to establishing a polynomial bound $M \le K^{O(1)}$ for the interface collapse. In the subgroup-interface model case we have $M = 1$ by direct inspection of the kernel on each side, whereas in the general Freiman setting we do not have such rigidity in the given ambient groups. The next step is therefore to replace the entire cospan by a Freiman-isomorphic diagram in universal ambient groups, where the maps extend to honest homomorphisms and the quotient behaves as in the model case. We carry this out in the next section.

6 Reduction to the subgroup case via universal ambient groups

In the preceding section we isolated a purely "quotient-stability" obstruction to bounding the doubling of the pushout: one needs polynomial control on the fiber multiplicities of the quotient map on the glued backbone. In the subgroup-interface model case this control is automatic (indeed $\mu_q = 1$ on each side), but in the general Freiman setting the maps i, j are not assumed to extend to homomorphisms on the given ambient groups, and the corresponding quotient may collapse large portions of $\widehat{A}$ or $\widehat{B}$.

We now explain how to remove this difficulty by replacing the entire cospan by a Freiman-isomorphic cospan in *universal ambient groups*, in which the maps do extend to honest homomorphisms. The key point is that pushouts in FR_2^0 are invariant under such replacements (up to Freiman isomorphism), and the sum/difference-set cardinalities $|\ell P - mP|$ with $\ell + m \leq 2$ are preserved. Consequently, any doubling estimate proved in the universal model transfers back to the original pushout.

6.1 Universal ambient groups as a diagram-level replacement

For each normalized additive set (S, Γ) in FR_2^0 there is a universal ambient group $U(S)$ and an embedding

$$\eta_S : S \hookrightarrow U(S)$$

such that η_S is a Freiman-2 isomorphism onto its image and $U(S)$ is generated by $\eta_S(S)$. The universal property we use is the following: if $\Phi : S \to \Lambda$ is any Freiman-2 homomorphism into an abelian group Λ with $\Phi(0) = 0$, then there exists a *unique* group homomorphism $\widetilde{\Phi} : U(S) \to \Lambda$ such that $\widetilde{\Phi} \circ \eta_S = \Phi$. In particular, any morphism $f : (S, \Gamma) \to (T, \Delta)$ in FR_2^0 induces a group homomorphism

$$U(f) : U(S) \to U(T) \qquad \text{with} \qquad U(f) \circ \eta_S = \eta_T \circ f.$$

Applying this construction objectwise to our cospan $(A, G) \xleftarrow{i} (C, L) \xrightarrow{j} (B, H)$ gives a new cospan

$$(A', U(A)) \xleftarrow{i'} (C', U(C)) \xrightarrow{j'} (B', U(B)),$$

where $A' := \eta_A(A)$, $B' := \eta_B(B)$, $C' := \eta_C(C)$, and

$$i' := \eta_A \circ i \circ \eta_C^{-1} : C' \to A', \qquad j' := \eta_B \circ j \circ \eta_C^{-1} : C' \to B'.$$

By construction, η_A, η_B, η_C are Freiman-2 isomorphisms of normalized sets, so the new cospan is isomorphic to the original cospan in FR_2^0.

Two basic features are immediate and will be used implicitly:

- Small doubling is preserved: since Freiman-2 isomorphisms preserve the size of $S \pm S$, we have $\sigma(A') = \sigma(A)$, $\sigma(B') = \sigma(B)$, $\sigma(C') = \sigma(C)$.
- Syndeticity is preserved: if $A \subseteq X + i(C)$ with $|X| \leq K$, then

  $$A' = \eta_A(A) \subseteq \eta_A(X) + i'(C')$$

  and $|\eta_A(X)| = |X| \leq K$. Similarly B' is covered by $\leq K$ translates of $j'(C')$.

Thus all standing hypotheses remain true after replacement. The advantage of this replacement is that the maps i', j' now arise from group homomorphisms on the ambient groups:

$$U(i) : U(C) \to U(A), \qquad U(j) : U(C) \to U(B),$$

whose restrictions to C' coincide with i', j' respectively.

6.2 Pushouts commute with Freiman-isomorphic replacement

Let P be the pushout of the original cospan and let P' be the pushout of the replaced cospan. We claim that P and P' are Freiman-2 isomorphic in a manner compatible with the structure maps from A, B and from A', B'. Concretely, the pushout is characterized by its universal property in FR_2^0, and Freiman isomorphisms are isomorphisms in this category; therefore pushouts are invariant under replacing a diagram by an isomorphic diagram.

More explicitly, denote the embeddings in the original pushout by

$$\alpha : A \to P, \qquad \beta : B \to P,$$

and similarly for the replaced pushout

$$\alpha' : A' \to P', \qquad \beta' : B' \to P'.$$

Since $\eta_A : A \to A'$, $\eta_B : B \to B'$, $\eta_C : C \to C'$ form an isomorphism of cospans in FR_2^0, the universal property of the pushout yields a unique morphism

$$\Theta : P \to P'$$

such that $\Theta \circ \alpha = \alpha' \circ \eta_A$ and $\Theta \circ \beta = \beta' \circ \eta_B$. Applying the same argument to the inverse cospan isomorphism gives a morphism $\Psi : P' \to P$, and uniqueness forces $\Psi \circ \Theta = \mathrm{id}_P$ and $\Theta \circ \Psi = \mathrm{id}_{P'}$. Hence Θ is an isomorphism in FR_2^0, i.e. a Freiman-2 isomorphism of normalized sets.

In particular, for $\ell + m \leq 2$ we have cardinality preservation

$$|\ell P - mP| = |\ell P' - mP'|,$$

since Freiman-2 isomorphisms preserve all additive relations of length at most 2 and therefore preserve the size of each set $\ell S - mS$ with $\ell + m \leq 2$. This is the precise sense in which colimits (here, pushouts) “respect” universal ambient group replacement for the purposes of doubling estimates.

6.3 The quotient model in universal ambient groups and the subgroup-interface behavior

To connect with the quotient-stability framework, we recall the concrete model for the pushout. In the replaced diagram, consider the abelian group

$$\Gamma' \ := \ U(A) \oplus U(B),$$

the subsets $\widehat{A'} := A' \times \{0\}$, $\widehat{B'} := \{0\} \times B'$, and the subgroup

$$N' \ := \ \Big\langle \big(U(i)(u), -U(j)(u)\big) \ : \ u \in U(C) \Big\rangle \ \le \ \Gamma'.$$

Let $q' : \Gamma' \to \Gamma'/N'$ be the quotient homomorphism. As in the general construction of pushouts in FR_2^0, we may take the underlying subset of P' to be

$$P' \ = \ q'(\widehat{A'} \cup \widehat{B'}).$$

The crucial point is that, unlike the original ambient groups, *the gluing data now comes from homomorphisms.* This forces the interface to behave as in the subgroup case with respect to collapse.

To make this precise, define the glued subset

$$Q' \ := \ q'(\widehat{i'(C')}) \ = \ q'(\widehat{j'(C')}) \ \subseteq \Gamma'/N'.$$

We will use the following "no-collapse along the interface" observation.

Lemma 6.1 (No collapse on the interface in the universal model)**.** *In the universal ambient group replacement, the quotient map q' is injective on $\widehat{i'(C')}$ and on $\widehat{j'(C')}$. Equivalently,*

$$\mu_{q'}(\widehat{i'(C')}) \ = \ \mu_{q'}(\widehat{j'(C')}) \ = \ 1.$$

Proof. We treat $\widehat{i'(C')}$; the argument for $\widehat{j'(C')}$ is identical. Suppose

$$q'\big(i'(c_1), 0\big) \ = \ q'\big(i'(c_2), 0\big) \qquad (c_1, c_2 \in C').$$

Then $\big(i'(c_1 - c_2), 0\big) \in N'$. By definition of N', there exist $u_t \in U(C)$ and integers n_t such that

$$\big(i'(c_1 - c_2), 0\big) \ = \ \sum_t n_t\big(U(i)(u_t), -U(j)(u_t)\big) \ = \ \Big(U(i)\Big(\sum_t n_t u_t\Big), \ -U(j)\Big(\sum_t n_t u_t\Big)\Big).$$

Comparing second coordinates gives $U(j)(u) = 0$, where $u := \sum_t n_t u_t \in U(C)$. Hence also $U(i)(u) = i'(c_1 - c_2)$. Since $U(C)$ is generated by C' and $U(j)$ extends the injective Freiman map j' on C', the only element of $U(C)$ mapped to 0 by $U(j)$ is $u = 0$ (this is the standard faithfulness of the universal ambient group construction on injective morphisms in FR_2^0). Therefore $u = 0$, so $i'(c_1 - c_2) = U(i)(u) = 0$. Finally, i' is injective on C' and $0 \in C'$, so $c_1 = c_2$, proving injectivity of q' on $\widehat{i'(C')}$. □

Lemma 6.1 exactly places us back in the bounded-collapse regime with $M = 1$ (Definition 5.3, applied to the replaced diagram). Combining this with the preservation of all hypotheses under replacement and the pushout invariance discussed above, we may henceforth argue *in the universal model* whenever we need a subgroup-interface type statement about the quotient. In particular, any doubling bound for P' obtained by treating the interface as non-collapsing transfers immediately to the original pushout P via the Freiman-2 isomorphism $\Theta : P \to P'$.

The remaining task is therefore purely organizational: we apply the quotient-stability machinery to P', obtain polynomial bounds for $|P' \pm P'|$ in terms of K, and transport them back to P. This assembly is carried out in the next section.

7 Proof of the main theorem in the general interface case

We now assemble the reductions and estimates from the previous sections to complete the proof for an arbitrary normalized interface C, assuming only that $\sigma(A), \sigma(B), \sigma(C) \le K$ and that i, j are K-syndetic embeddings. The only remaining issue after the quotient-stability discussion is that, in the original ambient groups G, H, L, the Freiman maps i, j need not extend to homomorphisms, so the quotient model of the pushout may exhibit uncontrolled collapse. The universal ambient group replacement removes this obstruction and reduces us to an essentially subgroup-interface behavior on the glued backbone.

7.1 Reduction to the universal model

Let

$$(A', U(A)) \xleftarrow{i'} (C', U(C)) \xrightarrow{j'} (B', U(B))$$

be the universal ambient group replacement of our cospan, with pushout P', as constructed previously. We retain the notation

$$\Gamma' := U(A) \oplus U(B), \qquad N' := \big\langle (U(i)(u), -U(j)(u)) : u \in U(C) \big\rangle,$$

and let $q' : \Gamma' \to \Gamma'/N'$ be the quotient map. As before, we realize the underlying set of the pushout as

$$P' \;=\; q'(\widehat{A'} \cup \widehat{B'}), \qquad \widehat{A'} := A' \times \{0\},\; \widehat{B'} := \{0\} \times B'.$$

By pushout invariance under Freiman-isomorphic replacement, there is a Freiman-2 isomorphism $\Theta : P \to P'$ compatible with the structure maps from A, B. Consequently, for $\ell + m \le 2$ we have

$$|\ell P - mP| \;=\; |\ell P' - mP'|. \tag{6}$$

In particular, any bound of the form $|\ell P' - mP'| \le K^{O(1)}|P'|$ transfers immediately to P, since also $|P| = |P'|$ (take $(\ell, m) = (1, 0)$ in (6)). Thus it suffices to prove the desired polynomial estimates for P'.

7.2 A common backbone in the quotient and translate coverings

Define the glued backbone in the quotient by

$$Q' := q'(\widehat{i'(C')}) = q'(\widehat{j'(C')}) \subseteq \Gamma'/N'.$$

By the no-collapse lemma in the universal model, the restriction of q' to $\widehat{i'(C')}$ (and to $\widehat{j'(C')}$) is injective. Hence $q' \circ \widehat{i'} : C' \to Q'$ is a bijection of normalized sets. Moreover, since i' is a Freiman-2 homomorphism and q' is a group homomorphism, the composition $q' \circ \widehat{i'}$ is a Freiman-2 homomorphism; bijectivity then implies that it is a Freiman-2 isomorphism. Therefore the order-2 additive statistics of Q' agree with those of C', and in particular

$$\sigma(Q') = \sigma(C') = \sigma(C) \le K, \qquad |Q' \pm Q'| = |C' \pm C'| \le K\,|C'| = K\,|Q'|. \tag{7}$$

Next we exploit syndeticity. Choose $X \subseteq G$ with $|X| \le K$ such that $A \subseteq X + i(C)$, and transport it to $U(A)$ via the Freiman embedding η_A, obtaining $X' := \eta_A(X) \subseteq U(A)$ with $|X'| = |X| \le K$ and

$$A' \subseteq X' + i'(C').$$

Passing to the quotient and using that q' is a homomorphism, we obtain the translate covering

$$q'(\widehat{A'}) \subseteq q'(X' \times \{0\}) + q'(\widehat{i'(C')}) = U_A + Q',$$

where we set $U_A := q'(X' \times \{0\})$ and note $|U_A| \le |X'| \le K$. Similarly, from $B \subseteq Y + j(C)$ with $|Y| \le K$, we get

$$q'(\widehat{B'}) \subseteq U_B + Q'$$

with $|U_B| \le K$. Since $P' = q'(\widehat{A'} \cup \widehat{B'})$, we conclude that

$$P' \subseteq (U_A \cup U_B) + Q'. \tag{8}$$

Let $T := U_A \cup U_B$; then $|T| \le 2K$ and (8) reads $P' \subseteq T + Q'$.

Two immediate size comparisons will be used repeatedly. First, $Q' \subseteq P'$ since $i'(C') \subseteq A'$ and hence $\widehat{i'(C')} \subseteq \widehat{A'}$, giving

$$|Q'| \le |P'|. \tag{9}$$

Second, (8) implies the crude upper bound

$$|P'| \le |T + Q'| \le |T|\,|Q'| \le 2K\,|Q'|. \tag{10}$$

Combining (9) and (10), we have $|P'| \asymp_K |Q'|$ with explicit losses bounded by $2K$.

7.3 Uniform control of $\ell P' - mP'$ for $\ell + m \leq 2$

From the translate covering $P' \subseteq T + Q'$ and the elementary containment

$$(T + Q') \pm (T + Q') \;\subseteq\; (T \pm T) + (Q' \pm Q'),$$

we obtain

$$P' \pm P' \;\subseteq\; (T \pm T) + (Q' \pm Q'). \tag{11}$$

Taking cardinalities and using the trivial bound $|X + Y| \leq |X|\,|Y|$ for finite sets in an abelian group gives

$$|P' \pm P'| \;\leq\; |T \pm T|\,|Q' \pm Q'| \;\leq\; |T|^2\,|Q' \pm Q'| \;\leq\; (2K)^2\,|Q' \pm Q'|. \tag{12}$$

Now apply (7) to bound $|Q' \pm Q'| \leq K|Q'|$, and then (9) (or (10) if we prefer to eliminate $|Q'|$ in favor of $|P'|$). We get

$$|P' \pm P'| \;\leq\; 4K^2 \cdot K\,|Q'| \;\leq\; 4K^3\,|P'|. \tag{13}$$

This proves the desired polynomial bound for both $P' + P'$ and $P' - P'$. More generally, for any $(\ell, m) \in \{(2,0), (1,1), (0,2)\}$ we have the same argument with $Q' \pm Q'$ and $T \pm T$, yielding

$$|\ell P' - mP'| \;\leq\; 4K^3\,|P'|. \tag{14}$$

In particular,

$$\sigma(P') = \frac{|P' + P'|}{|P'|} \;\leq\; 4K^3,$$

so the pushout has polynomially bounded doubling in the universal model.

7.4 Transfer back to the original pushout

Finally, we return to the original pushout P. By (6) and (14), for each $(\ell, m) \in \{(2,0), (1,1), (0,2)\}$ we have

$$|\ell P - mP| \;=\; |\ell P' - mP'| \;\leq\; 4K^3\,|P'| \;=\; 4K^3\,|P|.$$

In particular, $\sigma(P) \leq 4K^3$, which establishes the main theorem (with an explicit exponent in this presentation) and the uniform bounds for $P + P$ and $P - P$. Since all steps were stable under the finitely generated/torsion reductions already absorbed into the universal ambient group framework, no further case distinctions are required here.

8 Examples and near-sharpness

We record three families of examples illustrating (a) the stability predicted by the theorem when the interface is genuinely large, (b) the necessity of a non-degeneracy hypothesis such as syndeticity, and (c) the role of torsion and what dependence on parameters one should (and should not) expect.

8.1 Gluing progressions along a long progression

We begin with a torsion-free model in which all objects live in $\mathbb{Z}$. Fix integers $N_A, N_B \geq M \geq 2$ and consider

$$A = \{0, 1, \ldots, N_A - 1\} \subseteq \mathbb{Z}, \qquad B = \{0, 1, \ldots, N_B - 1\} \subseteq \mathbb{Z}, \qquad C = \{0, 1, \ldots, M - 1\} \subseteq \mathbb{Z},$$

with i, j the inclusions. Then $\sigma(A), \sigma(B), \sigma(C) \leq 2$. Moreover A is covered by $\lceil N_A/M \rceil$ translates of C, and similarly for B; thus the hypothesis of K-syndeticity holds with $K \asymp \max\{\lceil N_A/M \rceil, \lceil N_B/M \rceil\}$. When M is a fixed proportion of both N_A and N_B, we have $K = O(1)$.

To compute the pushout explicitly, set $G = H = L = \mathbb{Z}$. The subgroup $N \leq \mathbb{Z} \oplus \mathbb{Z}$ is generated by $(c, -c)$ for $c \in C$; since $1 \in C$, we have $N = \langle (1, -1) \rangle$, the diagonal copy of $\mathbb{Z}$. Hence $(\mathbb{Z} \oplus \mathbb{Z})/N \cong \mathbb{Z}$ via the homomorphism

$$\phi : \mathbb{Z} \oplus \mathbb{Z} \to \mathbb{Z}, \qquad \phi(x, y) = x + y,$$

which has kernel N. Under this identification,

$$q(\widehat{A}) = \phi(A \times \{0\}) = A, \qquad q(\widehat{B}) = \phi(\{0\} \times B) = B,$$

and the glued subset $q(\widehat{i(C)})$ is just C. Therefore the pushout set is

$$P = q(\widehat{A} \cup \widehat{B}) = A \cup B = \{0, 1, \ldots, \max\{N_A, N_B\} - 1\},$$

an interval. In particular $\sigma(P) \leq 2$, and likewise $|P - P| \leq 2|P|$. This is the simplest manifestation of the general theorem: once the interface contains a genuine generator, the quotient forces the two axes to cohere into a single one-dimensional object.

A slightly less tautological variant, still in $\mathbb{Z}$, is obtained by taking i to be inclusion and j to be the dilation $j(c) = dc$ for some integer $d \geq 2$, which is a Freiman-2 homomorphism on C and preserves 0. We then require $B \supseteq dC = \{0, d, 2d, \ldots, (M-1)d\}$; for instance we may take $B = \{0, 1, \ldots, N_B - 1\}$ with $N_B \geq d(M-1) + 1$. The subgroup N is generated by $(1, -d)$, so $(\mathbb{Z} \oplus \mathbb{Z})/N \cong \mathbb{Z}$ via

$$\phi_d(x, y) = dx + y,$$

which is invariant under $(x, y) \mapsto (x + t, y - dt)$. Then

$$q(\widehat{A}) = dA = \{0, d, 2d, \ldots, d(N_A - 1)\}, \qquad q(\widehat{B}) = B,$$

so $P = dA \cup B \subseteq \mathbb{Z}$. In particular P is the union of one interval and one d-spaced progression. One checks directly that $P + P \subseteq [0, 2N_B - 2] \cup d[0, 2N_A - 2] \cup ([0, N_B - 1] + d[0, N_A - 1])$, and hence $|P + P| = O(|P|)$ with an absolute constant as soon as N_B is comparable to dN_A (which is precisely the regime where the covering numbers of A by C and of B by dC are bounded). This illustrates that the quotient model can convert a two-axis union into a one-dimensional set even when the two axes have different scales, provided the common backbone is sufficiently large.

8.2 Failure without syndeticity: the coproduct-like pushout along $\{0\}$

We next isolate the basic obstruction. Let $C = \{0\}$ and let i, j be the unique 0-preserving maps. The subgroup N generated by $\{(i(0), -j(0))\}$ is trivial, so the quotient is $\Gamma = G \oplus H$. The pushout set is then simply

$$P = \widehat{A} \cup \widehat{B} = (A \times \{0\}) \cup (\{0\} \times B) \subseteq G \oplus H.$$

Even when A and B have very small doubling, the mixed sums force $\sigma(P)$ to be large. Indeed,

$$P + P \supseteq (A \times \{0\}) + (\{0\} \times B) = A \times B,$$

so $|P + P| \geq |A||B|$. On the other hand $|P| = |A| + |B| - 1$. Taking, for example, $G = H = \mathbb{Z}$ and

$$A = \{0, 1, \dots, n-1\}, \qquad B = \{0, 1, \dots, n-1\},$$

we have $\sigma(A), \sigma(B) \leq 2$ but

$$|P| = 2n - 1, \qquad |P + P| \geq n^2, \qquad \text{hence} \qquad \sigma(P) \geq \frac{n^2}{2n-1} \asymp n.$$

Thus no bound of the form $\sigma(P) \leq K^{O(1)}$ can hold in complete generality without an assumption preventing the interface from being too small. In the present formulation, syndeticity fails maximally: A cannot be covered by $O(1)$ translates of $\{0\}$ unless $|A| = O(1)$, and similarly for B.

This example also indicates what goes wrong conceptually: the pushout set P behaves like a coproduct (a disjoint union of axes) rather than an amalgamated sum, and the mixed sumset $A \times B$ has size comparable to the product of the sizes. Any mechanism that forces $q(\widehat{A})$ and $q(\widehat{B})$ to sit inside few translates of a common backbone eliminates precisely this product growth.

8.3 Torsion examples and parameter dependence

We finally discuss finite ambient groups, where torsion may create additional identifications in quotients and where it is useful to separate phenomena that are genuinely torsion-theoretic from those that are already present in $\mathbb{Z}^d$.

Cyclic torsion with progression data. Let p be prime and work in $\mathbb{Z}/p\mathbb{Z}$. Take $A = B = \{0, 1, \dots, n-1\}$ and $C = \{0, 1, \dots, m-1\}$ with $2 \leq m \leq n \leq p$. With i, j inclusions, the subgroup $N \leq (\mathbb{Z}/p\mathbb{Z})^2$ is generated by $(1, -1)$ and the quotient is again $\mathbb{Z}/p\mathbb{Z}$ via $\phi(x, y) = x + y$. Hence $P = A \cup B = A$, so $\sigma(P) = \sigma(A) \leq 2$. Here torsion is irrelevant: even if n is a positive fraction of p, the doubling of an interval in $\mathbb{Z}/p\mathbb{Z}$ remains ≤ 2 since $|A + A| = \min\{2n - 1, p\}$.

Subgroup backbone and coset-union growth. A more telling torsion example is the subgroup-interface regime. Let $G = H = (\mathbb{Z}/p\mathbb{Z})^r$, let $C \le G$ be a fixed subgroup of size $|C| = p^{r-1}$, and let i, j identify C with the same subgroup in G and H. Choose sets $T_A, T_B \subseteq G/C$ of sizes $|T_A| = M_A$, $|T_B| = M_B$, and set

$$A = \bigcup_{t \in T_A} (t + C), \qquad B = \bigcup_{t \in T_B} (t + C).$$

Then $\sigma(C) = 1$, and $\sigma(A)$ and $\sigma(B)$ are controlled in terms of M_A, M_B (indeed $A + A = (T_A + T_A) + C$, so $\sigma(A) = |T_A + T_A|/|T_A| \le M_A$, and similarly for B). Moreover A is M_A-syndetic over C and B is M_B-syndetic over C. In the pushout quotient, the glued subgroup remains a subgroup, and P is contained in a union of $\le M_A + M_B$ cosets of it. Consequently $\sigma(P)$ is $\ll M_A + M_B$, and this linear dependence is in general sharp: if T_A is a dissociated set in G/C, then $|T_A + T_A| \asymp M_A^2$ and $\sigma(A) \asymp M_A$, while $\sigma(P)$ is $\asymp M_A$ as well. This provides a genuine lower bound of order K (up to constants) for any theorem of the form $\sigma(P) \le K^c$.

What torsion can force. In finite groups, additional collapse can occur if the subgroup generated by $\{(i(c), -j(c))\}$ is larger than one would predict from the formal rank of C as a Freiman object, simply because ambient torsion introduces relations not visible at order 2. The universal ambient group replacement is designed precisely to avoid drawing incorrect conclusions from such accidental relations: it upgrades the Freiman data to honest homomorphisms in a universal setting and only then passes to a quotient. From the perspective of bounds, the effect of torsion is therefore confined to bounded factors that can be absorbed into $K^{O(1)}$ once one reduces to finitely generated models and separates torsion-free and torsion parts.

8.4 Near-sharpness and possible improvements of exponents

The proof presented earlier yields an explicit bound $\sigma(P) \le 4K^3$ in the universal model by a sequence of worst-case estimates. The examples above show that one cannot hope, in general, for a bound independent of K: already in the subgroup-interface regime, $\sigma(P)$ can grow linearly in the covering numbers (and hence linearly in K under the syndeticity hypothesis). On the other hand, these same subgroup examples suggest that the exponent 3 is not best possible in structured situations: when P is genuinely a union of $O(K)$ cosets of a backbone with $\sigma \le K$, one expects $\sigma(P)$ to be at most $K^{1+O(1)}$, and often $\ll K$.

What obstructs improving the exponent in the general statement is not the behavior of Q' (whose doubling is already $\le K$), but the potential inefficiency in passing from a covering $P' \subseteq T + Q'$ to lower bounds on $|P'|$ in terms of $|T||Q'|$: heavy overlap among the translates $t + Q'$ may make $|P'|$

closer to $|Q'|$ than to $|T||Q'|$. Any refinement that quantitatively controls this overlap (for instance, via additional regularity assumptions on the coverings, or via energy bounds for the set of translate representatives) would feed directly into a better exponent. We leave such improvements to the applications section and to the open problems listed thereafter.

9 Applications and extensions

We record several directions in which the pushout stability statement (Theorem A and Corollary B) can be used as a basic "gluing lemma" in additive combinatorics, and we formulate a number of natural refinements that appear accessible once one improves the bookkeeping of translate overlap in the quotient model.

9.1 Pullback analogues and fiber products

The pushout construction amalgamates two sets along a common interface. A formally dual operation is the pullback (fiber product) of a span

$$(A, G) \xrightarrow{f} (C, L) \xleftarrow{g} (B, H)$$

in FR_2^0, which (when it exists in a suitable ambient category) should model the set of pairs (a, b) with $f(a) = g(b)$. In the subgroup-interface case, when $C \leq L$ is a subgroup and f, g extend to homomorphisms $G \to L$, $H \to L$, the ambient pullback group is the subgroup

$$G \times_L H := \{(x, y) \in G \oplus H : \; f(x) = g(y)\},$$

and the underlying pullback set is contained in $(A \times B) \cap (G \times_L H)$.

Unlike pushouts, pullbacks can easily be too large even when A and B have small doubling: if f and g are constant maps, the pullback is essentially $A \times B$. Thus any pullback stability statement must impose a non-degeneracy hypothesis of a genuinely different flavor. One natural condition is a *bounded fiber* hypothesis: for some $M \geq 1$, every $c \in C$ has at most M preimages in A under f, and at most M preimages in B under g. Another is a *co-syndetic image* hypothesis: $C \subseteq U + f(A)$ and $C \subseteq V + g(B)$ with $|U|, |V| \leq M$, so that $f(A)$ and $g(B)$ cover C up to M translates.

In the subgroup-interface regime one can combine these hypotheses with the standard estimate $|S \times T| = |S||T|$ and the observation that $G \times_L H$ is the kernel of the homomorphism $(x, y) \mapsto f(x) - g(y)$ to obtain crude control of the doubling of the pullback set in terms of the doubling of A and B and the fiber parameters. We record the guiding heuristic as an informal principle: bounded fibers prevent product growth in the same way that syndetic interfaces prevent mixed-sum growth in pushouts.

Problem. Formulate and prove a pullback analogue of Theorem A: find checkable hypotheses on f, g (for instance, bounded fibers together with small doubling of $f(A)$ and $g(B)$) under which the pullback set P^* satisfies $\sigma(P^*) \le K^{O(1)} M^{O(1)}$.

In the general (non-subgroup) setting, the universal ambient group replacement (Lemma 5) appears again to be the correct first step: one upgrades f, g to homomorphisms between universal models and only then forms the fiber product subgroup. The main technical obstruction is that fiber bounds are not invariant under arbitrary Freiman isomorphism; one should work with a diagram-level notion of bounded fibers that is intrinsic to FR_2^0.

9.2 Iterated gluing and complexity control

A standard pattern in additive combinatorics is to build large structured sets by repeatedly adjoining pieces along substantial overlaps. The categorical pushout formalizes this procedure. Suppose we have a finite tree of normalized additive sets $\{(A_v, G_v)\}_{v \in V}$ and normalized interfaces (C_e, L_e) on edges $e \in E$, with injections $C_e \to A_v$ into the incident vertices, each satisfying the hypotheses of Theorem A with parameter K. Iteratively pushing out along the edges produces a single object P_T equipped with compatible Freiman embeddings of each A_v.

Two issues arise: (i) whether the resulting object depends on the order of gluing, and (ii) how the doubling constant propagates. Associativity of pushouts in a fixed category suggests that different orders yield canonically isomorphic colimits; in our concrete model $P = q(\widehat{A} \cup \widehat{B})$ this is reflected by the fact that successive quotients by subgroups generated by interface relations commute up to canonical isomorphism. Thus the relevant quantity is the cumulative loss in the doubling constant, not the ambiguity of the construction.

A naive iteration of Theorem A yields a bound of the form

$$\sigma(P_T) \le K^{c|E|}$$

for a constant $c > 0$, since each gluing step introduces a polynomial loss. In many applications the number of gluings is itself bounded by a constant (e.g. bounded-complexity decompositions), and this estimate suffices. When $|E|$ is large, it becomes important to identify hypotheses under which the losses do not multiply. In the subgroup-interface case, Proposition C already suggests a more stable behavior: if each vertex set A_v is covered by $O(K)$ cosets of the glued backbone, then the final colimit remains covered by $O(K|V|)$ cosets of the common image of that backbone, and one expects

$$\sigma(P_T) \ll K^{O(1)} |V|^{O(1)}$$

rather than $K^{O(|E|)}$. Proving such a statement in the general Freiman setting amounts to controlling overlap among many families of translates in a common quotient.

> *Problem.* Develop an "iterated syndetic gluing lemma" giving $\sigma(P_T) \le K^{O(1)}\mathrm{poly}(|V|)$ under additional regularity assumptions on the translate coverings (for instance, bounded additive energy among the translate representatives at each stage).

9.3 Functorial Freiman models from colimits

One motivation for introducing categorical pushouts is that many arguments in additive combinatorics are diagrammatic: one compares several sets linked by Freiman homomorphisms, passes to quotient models, and deduces structural information that is not attached to any single set in isolation. The universal ambient group replacement (Lemma 5) already provides a functorial way to turn a Freiman object (S, Γ) into a group $U(S)$ together with a Freiman embedding $S \hookrightarrow U(S)$. Pushouts allow one to extend this functoriality to *diagrams*: given a finite diagram D in FR_2^0 satisfying syndetic interface hypotheses on its morphisms, one may form its colimit object $\mathrm{colim}(D)$ and obtain a single ambient group in which all objects of D embed compatibly.

In practice, one uses this as follows. Suppose $A \subseteq G$ and $B \subseteq H$ are small-doubling sets known to share a large Freiman-structured subset C (for example, a common progression-like model or a common large energy component). The pushout P provides a canonical way to identify A and B along C while keeping explicit control of doubling in the resulting ambient. Once P has small doubling, one may apply any available structure theorem to P (for example, Freiman–Ruzsa type modeling results in torsion-free groups) and then pull back the obtained structure simultaneously to A and B. The key point is that this avoids having to choose models for A and B independently and then check compatibility: compatibility is built into the colimit.

A related use is the construction of *functorial approximate group envelopes*. If one attaches to each small-doubling set S a controlled set $E(S)$ (say, a coset progression containing S with bounded rank and bounded size inflation), then one would like E to behave well with respect to gluing: whenever A and B are glued along a large interface, $E(P)$ should be comparable to the pushout of $E(A)$ and $E(B)$ in an appropriate category of structured sets. Establishing such functoriality appears to require precisely the kind of polynomial doubling control provided by Theorem A, together with quantitative stability of the chosen envelope construction under Freiman isomorphism.

9.4 Open problems and refinements

We close by listing several concrete problems suggested by the proofs and examples.

1. **Exponent improvement.** Determine the optimal growth rate of $\sigma(P)$ in terms of K under syndeticity. The subgroup-interface examples force $\sigma(P) \gg K$ in general, while the present proof yields $\sigma(P) \leq K^{O(1)}$ with a non-optimized exponent. It is natural to ask whether $\sigma(P) \leq K^{1+o(1)}$ holds under additional mild hypotheses on the coverings.

2. **Higher sumsets.** Extend Corollary B from $\ell + m \leq 2$ to uniform bounds $|\ell P - mP| \leq K^{O_{\ell,m}(1)}|P|$ for fixed ℓ, m. In the quotient model $P = q(\widehat{A} \cup \widehat{B})$, this requires controlling mixed expressions with many alternations between the two axes, and appears to demand more than small doubling of A, B, C.

3. **Characterizing non-degeneracy.** Syndeticity is sufficient and checkable, but not necessary. Find an intrinsic condition on the cospan $A \leftarrow C \rightarrow B$ that is equivalent (up to polynomial losses) to small doubling of the pushout. Any such characterization should rule out the coproduct-like behavior of Proposition D while allowing interfaces that are not literally translate-large.

4. **Non-abelian variants.** Develop an analogue for non-abelian approximate groups, replacing Freiman-2 homomorphisms by suitable partial homomorphisms and replacing the quotient $(G \oplus H)/N$ by an amalgamated free product modulo relations. Even formulating the correct category in which a controlled pushout exists is nontrivial.

5. **Algorithmic questions.** In finite ambient groups, compute or approximate $|P|$ and $|P + P|$ from combinatorial data of A, B, C and the coverings X, Y. This is relevant to the empirical search for extremizers and for testing conjectured sharp exponents.

These problems all share a common theme: Theorem A reduces the stability of gluing to the control of overlap among a bounded number of translates in a quotient. Any method that quantifies such overlap more efficiently than the present worst-case bounds should immediately yield sharper exponents and stronger functorial statements.