POWER WINDOWS IN THE NIELSEN–SOELBERG GROUP G_1: THE ELEVENTH STAIRCASE CELL, AND THE BALLS THAT WERE NOT COMPUTED

KOREA SUPERINTELLIGENCE LABS

Abstract. Nielsen and Soelberg proved that a finite subset A of a torsion-free group whose square $A \cdot A$ contains no uniquely represented element has $|A| \geq 8$, and exhibited two groups attaining the bound; the first, here G_1, is a two-generator group that is virtually class-2 nilpotent, with a torsion-free nilpotent normal subgroup of index 32. A recent quantitative study of these extremal groups computes, inside the word balls B_r of the defining generating set of G_1, the *unique-product staircase* $u_{B_r}(n)$: the least number of uniquely represented elements of $A \cdot A$ over all n-element $A \subseteq B_r$. It finds $u_{B_5}(n) = 2$ for $2 \leq n \leq 10$ and leaves the cells $n = 11$ and $n = 12$ open, its solver returning only the upper bound 15 against a proved lower bound of 2. We settle the first of them: $u_{B_5}(11) = 2$.

The instrument is a lemma that appears in neither that paper nor its companion: $n \geq 2$ consecutive powers of an element of infinite order have *exactly two* unique products, the multiplicative reading of the elementary fact that an n-term arithmetic progression has exactly two sums of multiplicity one. It is sharp here in a precise sense: the longest run of consecutive powers of a single element of G_1 inside B_r, normalised to contain the identity, is exactly 11 for $r = 5$ and 13 for $r = 6$. So the method reaches $n = 11$ in B_5 and no further — exactly where the published flat row ends — and beyond it we narrow $u_{B_5}(12)$ from $[2, 15]$ to $[2, 3]$. In the two larger balls, whose staircases are not computed in the source at all, it gives $u_{B_6}(n) \leq 2$ for $2 \leq n \leq 13$ and $u_{B_7}(n) \leq 2$ for $2 \leq n \leq 15$, the latter beside $u_{B_7}(8) = 0$. We also record a parity obstruction: for an inverse-closed set in a torsion-free group the number of unique products is even, so a set realising the value 1 — the object of the source's remaining open cell $u_{B_6}(8) \in \{1, 2\}$ — cannot be inverse-closed. Every count the theorems below rest on is machine-checked in Lean 4; the auxiliary computations that are not are labelled where they occur.

1. Introduction

Unique products, and the smallest failure. A group G has the *unique product property* if for all finite nonempty $A, B \subseteq G$ some element of $A \cdot B$ is represented exactly once as ab with $a \in A$, $b \in B$; a single finite set A is *non-UP* if $A \cdot A$ has no uniquely represented element. These sets are the classical combinatorial obstruction behind Kaplansky's zero-divisor and unit problems for group rings of torsion-free groups [4, 8]; the first torsion-free group without the property was produced by Rips and Segev [10], the standard small example is Promislow's [9], and Gardam's counterexample to the unit conjecture [3] runs through Promislow's group.

Nielsen and Soelberg [7, Thm. 1.2] settled the extremal question at the level of all torsion-free groups: if $A \cdot A$ has no unique product then $|A| \geq 8$, and the bound is sharp. Their proof of sharpness produced two torsion-free groups, called G_1 and G_2 after their Section 3, each containing an explicit 8-element non-UP set. According to [11] these two are, to date, the only known witnesses to the sharpness of the global bound.

The staircase, and the cells left open. The existence statement says nothing about how such a configuration sits inside its ambient group. Tabei [11, 12] makes that question quantitative. For G_1, generated by x and y as in [7, §3], write B_r for the ball of radius r in the word metric of $\{x^{\pm 1}, y^{\pm 1}\}$; [11] records the ball sizes

$$|B_1|, \ldots, |B_7| \ = \ 5,\ 17,\ 53,\ 153,\ 401,\ 933,\ 1935,$$

and, writing $u(n)$ for "the least number of unique products of an n-subset of a stated ball", studies the resulting *staircase*. Two facts from that paper frame this one. Its localization theorem [11, Thm. 3.1] says that for every $r \leq 6$ the ball B_r contains *no* 8-element non-UP set, while B_7 contains the Nielsen–Soelberg witness A_1: the global minimum of G_1 is spread out, and B_7 is the least ball holding it. Its staircase proposition [11, Prop. 5.1] says that for $2 \leq n \leq 10$ the value $u_{B_5}(n)$ is exactly 2, and then adds, verbatim:

> (For $n = 11, 12$ the solver returned minimizers with 15 unique products against a proved lower bound of 2 within budget; we leave those values open.)

A second sentence of the same section leaves a third cell open:

> (Whether the value 1 occurs for G_1 at $n = 8$ in B_6 — where 0 is certified impossible — we could not decide within a six-hour budget; $u_{B_6}(8) \in \{1, 2\}$ remains open.)

Outside B_5 no staircase of G_1 is computed at all in [11]: its work in B_6 is the localization infeasibility and this undecided cell.

What is settled here. We close the first open cell and locate the second and third precisely. All statements are about G_1 with the generating set above.

- $u_{B_5}(11) = 2$ (Theorem 4.1 with the lower bound of [11, Prop. 5.1]): the flat row of [11] extends from $2 \leq n \leq 10$ to $2 \leq n \leq 11$. The witness is the set of the eleven consecutive powers $x^{-5}, \ldots, x^5$, which has exactly two unique products, thirteen fewer than the minimiser the solver of [11] returned at this size.
- $u_{B_5}(12) \in [2, 3]$ (Theorem 4.2), narrowing the interval $[2, 15]$ of [11]. Whether 2 is attained stays open, and we show that the method used for $n \leq 11$ provably cannot reach it.
- $u_{B_6}(n) \leq 2$ for $2 \leq n \leq 13$ and $u_{B_7}(n) \leq 2$ for $2 \leq n \leq 15$ (Theorems 5.1 and 5.2); at $n = 8$ the second bound is not tight, since $u_{B_7}(8) = 0$.
- The longest run of consecutive powers of one element of G_1 inside a ball, normalised to contain the identity, is exactly 11 in B_5 and 13 in B_6 (Theorem 3.4), attained by the four generators. This is what makes the eleventh cell and the twelfth cell different in kind.

The engine is elementary. If g has infinite order and $A = \{g^k : k \in I\}$ for an integer interval I with $|I| = n \geq 2$, then the multiplicities along $A \cdot A$ are $1, 2, \ldots, n, \ldots, 2, 1$, so $u(A) = 2$ exactly (Lemma 3.2). This is the multiplicative reading of the fact that an n-term arithmetic progression $P \subseteq \mathbb{Z}$ has $|P + P| = 2|P| - 1$ with exactly two sums of multiplicity one — the equality case of Freiman's inverse problem for small doubling [1, 6]. We claim no novelty for the lemma; we claim the values it settles, and the fact that it settles them. Neither [11] nor [12] uses it: a search of the LaTeX sources of both papers finds no occurrence at all of "progression", "consecutive" or "powers of", and the single occurrence of "geometric" is in the unrelated phrase "geometric, not notational" about the localization phenomenon of [11, Thm. 3.1]. Yet it reproduces by inspection the whole flat row those papers obtained by constraint solving, and it is predictive, because the length of the longest power window inside a ball is a computable invariant of the group and the generating set.

Two caveats belong in the introduction rather than in a footnote. First, our contributions to the staircase are *upper* bounds; the equality $u_{B_5}(11) = 2$ combines our upper bound with the lower bound of [11, Prop. 5.1], which that paper reports as a constraint solver's verdict and not as an externally checkable certificate — [11] says of its own results that "the staircase values and the two-sided localization remain CP-SAT verdicts". The interval $[2, 3]$ for $n = 12$ inherits its left endpoint from the same source. Second, our computations are carried out in an exact coordinate model of G_1 built from the tables published with [11]; that this model is a faithful copy of G_1 is *that paper's* certificate chain and is not reproved here (Section 2). What we do verify, from those tables and independently of the implementation that produced them, is that the model satisfies both defining relators, that its finite quotient is a group of order 32, that its ball sizes reproduce the published sequence, and that the published witnesses have the published invariants.

What is not claimed. We do not decide $u_{B_5}(12)$, and we do not decide $u_{B_6}(8)$. We do not compute any staircase cell exactly except through a lower bound taken from [11] (for $u_{B_5}(11)$) or through the trivial lower bound 0 (for $u_{B_7}(8)$). We prove no lower bound of our own anywhere. We do not count the non-UP 8-sets of B_7 — the "?" entry in the comparison table of [11] — and we say nothing about G_2, about the universal group of [7, §4], or about the two-sided problem. Section 7 states what each of these would cost.

A note on the word "symmetric". In [11, 12] a *symmetric* non-UP set is one for the one-sided problem $A \cdot A$, as against a two-sided pair $A \cdot B$ (the word does not occur in [7]). Section 6 concerns sets with $A = A^{-1}$, an unrelated condition; to avoid a collision we call those *inverse-closed* throughout and reserve "symmetric" for the sources' meaning.

2. The group, its balls, and the model

The group. Following [7, §3], as transcribed in [11, §2.1], G_1 is the group generated by x, y subject to the two relations

$$yx^{-1}y^{-1}xy^2x^{-1}y^{-2}x^{-1}yxy^{-1}x = 1, \qquad yx^{-1}yxy^{-1}x^{-1}yx^{-1}yx^{-1}y^{-1}x = 1.$$

The subgroup $H = \langle h_1, h_2, h_3, h_4 \rangle$ with

$$h_1 = (xy)^2, \quad h_2 = yx^2yx^{-1}y^{-2}x^{-1}, \quad h_3 = xy^{-1}xyxy^{-2}x^{-1}, \quad h_4 = y^2x^{-1}yxy^{-1}x^{-2}$$

is normal of index 32, with presentation $H = \langle h_1, \dots, h_4 \mid h_1, h_2 \text{ central},\ h_4h_3 = h_2^8h_3h_4 \rangle$, so $H \cong \mathbb{Z} \times N$ with N torsion-free class-2 nilpotent; G_1 is torsion-free [7]. The Nielsen–Soelberg witness is

$$A_1 = \{\, x,\ y,\ y^{-1}x^2,\ y^{-1}xy,\ xy^{-1}x,\ x^{-1}yxy^{-1}x,\ x^{-1}yx,\ x^{-1}yx^{-1}y^{-1}xy^2 \,\}.$$

Unique products and the staircase. For a finite $A \subseteq G$ let

$$\mathcal{U}(A) \;=\; \{\, w \in A \cdot A \;:\; \#\{(a,b) \in A \times A : ab = w\} = 1 \,\}, \qquad u(A) = |\mathcal{U}(A)|,$$

so A is non-UP exactly when $u(A) = 0$. For $B \subseteq G$ and $n \geq 2$ the staircase value is

$$u_B(n) \;=\; \min\{\, u(A) \;:\; A \subseteq B,\ |A| = n \,\}.$$

Throughout, $B_r \subseteq G_1$ is the ball of radius r about the identity in the word metric of $\{x^{\pm 1}, y^{\pm 1}\}$. Every statement below is relative to that generating set and to balls centred at the identity; localization radii and staircases depend on both, as [11] stresses.

The model. Elements of G_1 are represented as pairs (c, h) with c a coset index in $\{0, \dots, 31\}$ for G_1/H and $h \in \mathbb{Z}^4$ a Mal'cev exponent vector for H, multiplied by

$$(c, h) \cdot (c', h') \;=\; \big(m(c, c'),\ h \cdot \alpha_c(h') \cdot s(c, c')\big),$$

where m is the coset multiplication table, α_c the conjugation action of the transversal representative t_c, and $s(c, c') = t_c t_{c'} t_{m(c,c')}^{-1}$ the cocycle, with the class-2 collection law $(a_1, a_2, a_3, a_4)(b_1, b_2, b_3, b_4) = (a_1{+}b_1,\ a_2{+}b_2{+}8a_4b_3,\ a_3{+}b_3,\ a_4{+}b_4)$ on $\mathbb{Z}^4$. The three tables are those extracted by modified Todd–Coxeter rewriting and published as ancillary data with [11]; the arithmetic here is an independent implementation of the law printed in [11, §2.2]. Faithfulness — that the model is a copy of G_1 — is the certificate chain of [11, §2.3], resting on the index-32 theorem and the torsion-freeness argument of [7]; it is not reproved here, and every theorem below should be read as a theorem about G_1 modulo that chain.

One implementation point is worth recording, because it is not visible in the printed law. In α_c the four published rows are the images of $h_1, \dots, h_4$; the action on a general $h = (a_1, a_2, a_3, a_4)$ is the Mal'cev word $\alpha_c(h_1)^{a_1}\alpha_c(h_2)^{a_2}\alpha_c(h_3)^{a_3}\alpha_c(h_4)^{a_4}$ and *not* the linear map with those rows as columns. Reading the rows linearly breaks associativity: on 20 000 random triples drawn from B_4 the linear reading fails 938 times and the Mal'cev reading fails none.

The reproduction gate. Before any new value we reproduce what [11] publishes for G_1.

Proposition 2.1 (ball sizes)**.** *Breadth-first search from the identity in the model gives*

$$|B_0|,\ |B_1|, \dots, |B_7| \ = \ 1,\ 5,\ 17,\ 53,\ 153,\ 401,\ 933,\ 1935,$$

whose entries $|B_1|, \dots, |B_7|$ are the sequence recorded in [11, §3].

Proposition 2.2 (the published data, re-derived)**.** *In the model: both defining relators of G_1 evaluate to the identity, and the coset table is the multiplication table of a group of order* 32*; the Nielsen–Soelberg witness A_1 has $u(A_1) = 0$ and consists of eight distinct elements exhibited by words of length at most* 7*, while its words do not fit in radius* 6*; and the* 8*-element subset of B_5 shipped as ancillary data with* [11] *as a minimiser for the value $u_{B_5}(8) = 2$ of its Proposition 5.1 has $u = 2$.*

Proposition 2.2 is a check of the model against its source, not a new result: every item is a value [11] publishes. Its point is that the values survive an independent implementation. Two further reproductions were made in the reference implementation outside the machine-checked development, and we label them as such: the 64 ordered products of A_1 fall into equal-product classes of sizes $\{3^2, 2^{29}\}$ — two classes of size three and twenty-nine of size two, the fingerprint described in [7, §4] and reported in [11, §2.3] — and the geodesic lengths of the eight elements of A_1 form the multiset $\{1, 1, 3, 3, 3, 3, 5, 7\}$, matching the tuple $(1, 1, 3, 3, 3, 5, 3, 7)$ printed in [11, §3]. All eleven staircase minimisers shipped with [11] were also re-evaluated there, giving $u = 2$ at $n = 2, \dots, 10$ and $u = 15$ at $n = 11, 12$, the values that paper reports.

3. Power windows

Definition 3.1. Let $g \in G$ have infinite order and let $I \subseteq \mathbb{Z}$ be an interval of integers. The set $\{g^k : k \in I\}$ is the *power window* of g on I; its length is $|I|$.

Lemma 3.2. *Let g have infinite order in a group G, let $I = \{a, a+1, \dots, a+n-1\}$ with $n \geq 2$, and put $A = \{g^k : k \in I\}$. Then $|A| = n$, $A \cdot A = \{g^m : m \in I + I\}$ has $2n-1$ elements, the multiplicity of g^m in $A \cdot A$ is $n - |m - (2a+n-1)|$, and*

$$\mathcal{U}(A) = \{g^{2a},\ g^{2a+2n-2}\}, \qquad u(A) = 2.$$

Proof. Since g has infinite order, $k \mapsto g^k$ is injective, so $|A| = n$ and products of elements of A are indexed by sums: $g^i g^j = g^{i+j}$ with $i, j \in I$, and $g^m = g^{m'}$ only if $m = m'$. The number of representations of g^m is $\#\{(i,j) \in I \times I : i + j = m\}$, which for $m \in I + I = \{2a, \dots, 2a+2n-2\}$ equals $n - |m - (2a+n-1)|$ and is 0 otherwise. This is 1 exactly at the two endpoints $m = 2a$ and $m = 2a + 2n - 2$, which are distinct because $n \geq 2$. □

The lemma is elementary and certainly folklore; it is the multiplicative form of the observation that a finite arithmetic progression $P \subseteq \mathbb{Z}$ of length $n \geq 2$ satisfies $|P + P| = 2|P| - 1$ with exactly two sums of multiplicity one, the extremal case in Freiman's inverse theory of small doubling [1, 6]. Its use here is the immediate corollary.

Corollary 3.3. *Let $B \subseteq G$ and suppose B contains a power window of length $L \geq 2$. Then $u_B(n) \leq 2$ for every $2 \leq n \leq L$, since every sub-interval of an interval is an interval.*

So the reach of the method inside a ball is governed by one number: how long a run of consecutive powers of a single element fits in that ball. We compute it. Because a ball is inverse-closed and contains the identity, the natural normalisation is to intervals containing 0, and that is the invariant we compute:

$$L(r) \ = \ \max\bigl\{\, |I| \ : \ I \subseteq \mathbb{Z} \text{ an interval},\ 0 \in I,\ \{g^k : k \in I\} \subseteq B_r \text{ for some } g \in B_r \setminus \{1\} \,\bigr\}.$$

Theorem 3.4 (the windows of G_1)**.** *In G_1 with the generating set $\{x^{\pm 1}, y^{\pm 1}\}$,*

$$L(5) = 11 \qquad \textit{and} \qquad L(6) = 13,$$

and in both balls the maximum is attained exactly at the four generators $x^{\pm 1}, y^{\pm 1}$.

Proof (exhaustive computation). For each of the 400 elements $g \neq 1$ of B_5 and each of the 932 of B_6, the powers $g^0, g^1, g^2, \dots$ and $g^0, g^{-1}, g^{-2}, \dots$ are generated in the model and tested for membership in the ball until the first miss, with a cap of 20 steps in each direction; the window length is the sum of the two runs less one for the doubly counted identity. The cap is not active: the maxima are 11 and 13. Both sweeps are machine-checked; the identification of the maximisers is read off the same enumeration. □

Two remarks on the shape of this statement. It measures runs *containing the identity's exponent*; a run $g^a, \dots, g^{a+n-1}$ with $0 \notin \{a, \dots, a+n-1\}$ is also a power window, and nothing above excludes a longer one of that kind, so Theorem 3.4 is a statement about the normalised invariant and not about all runs. And the generators are undistorted, since $x^{-r}, \dots, x^r$ lies in B_r by inspection of the words; the content of the theorem is that nothing beats them.

4. The eleventh and the twelfth cell

Theorem 4.1. *The eleven consecutive powers $x^{-5}, x^{-4}, \dots, x^4, x^5$ are distinct, lie in $B_5 \subseteq G_1$, and satisfy*

$$u\big(\{x^k : |k| \leq 5\}\big) \;=\; 2.$$

Hence $u_{B_5}(11) \leq 2$, and with the lower bound of [11, Prop. 5.1],

$$u_{B_5}(11) = 2.$$

Proof. Membership is exhibited by the geodesic words x^k, $|k| \leq 5$, each of length at most 5; distinctness and the count $u = 2$ are computations in the model, both machine-checked (the count is a kernel evaluation over the 121 products, not a search). The value 2 is what Lemma 3.2 predicts with $g = x$, $a = -5$, $n = 11$: the unique products are x^{-10} and x^{10}. The equality then combines this upper bound with the lower bound 2 that [11, Prop. 5.1] reports for this cell. □

This closes the first of the two cells left open in [11, Prop. 5.1] and extends its flat row from $2 \leq n \leq 10$ to $2 \leq n \leq 11$. It is worth recording how far the flat row was from the solver's reach: the 11-element minimiser shipped with [11] has $u = 15$, which we re-verified (Section 2), so the window improves the published upper bound at this cell by thirteen. It was not found by search but by reading the coincidence signature of the source's own smaller minimisers.

The twelfth cell is different in kind, and Theorem 3.4 says why: $12 > L(5) = 11$, so no power window in B_5 reaches it and Corollary 3.3 is exhausted. What we can still do is add one element to a window.

Theorem 4.2. *The* 12*-element set*

$$A_{12} \;=\; \{\, y^k : |k| \leq 5 \,\} \;\cup\; \{\, xyx \,\} \;\subseteq\; B_5$$

consists of twelve distinct elements and satisfies $u(A_{12}) = 3$. Hence $u_{B_5}(12) \leq 3$, and with the lower bound of [11, Prop. 5.1],

$$u_{B_5}(12) \in [2, 3],$$

narrowing the interval $[2, 15]$ *reported there.*

Proof. Membership in B_5 is exhibited by the words y^k ($|k| \leq 5$) and xyx, all of length at most 5; distinctness and $u(A_{12}) = 3$ are machine-checked kernel evaluations over the 144 products. The set was produced by extending the 11-term power window of y by each of the remaining 390 elements of B_5 in turn and keeping a best result. □

Whether 2 is attained at $n = 12$ we do not know, and Theorem 3.4 shows the present method cannot decide it. Two exhaustive computations, run in the reference implementation and *not* part of the machine-checked development, say where such a witness cannot sit. The ball B_5 contains exactly two power windows of length 11, namely $\{x^k : |k| \leq 5\}$ and $\{y^k : |k| \leq 5\}$ (a window of g and of g^{-1} being the same set). Among the 2×390 twelve-element sets containing one of them entirely, the minimum of u is 3; among the $2 \times 834\,405$ twelve-element sets containing exactly ten elements of one of them, the minimum is 4. Consequently a 12-element subset of B_5 with $u = 2$, if one exists, meets each of the two length-11 windows in at most nine elements.

5. The staircases of the larger balls

No staircase outside B_5 is computed in [11]. Corollary 3.3 supplies one row of each of the next two, for free.

Theorem 5.1. *In G_1, $u_{B_6}(n) \le 2$ for every $2 \le n \le 13$.*

Proof. The thirteen consecutive powers $x^{-6}, \dots, x^{6}$ lie in B_6, exhibited by the words x^k, $|k| \le 6$. For each n with $2 \le n \le 13$ the initial segment $\{x^k : -6 \le k \le -7+n\}$ is a power window of length n inside B_6, and its unique products are counted directly in the model. All twelve counts are machine-checked in one sweep, and each is 2, as Lemma 3.2 predicts. □

Theorem 5.2. *In G_1, $u_{B_7}(n) \le 2$ for every $2 \le n \le 15$. At $n = 8$ the bound is not tight: $u_{B_7}(8) = 0$, attained by the Nielsen–Soelberg witness A_1.*

Proof. As for Theorem 5.1, with the fifteen consecutive powers $x^{-7}, \dots, x^{7}$ of B_7 and their fourteen initial segments; the counts are machine-checked in one sweep. The value at $n = 8$ is Proposition 2.2: $A_1 \subseteq B_7$ has $u(A_1) = 0$, and 0 is the least possible value. □

Remark 5.3. Part of these two rows is implicit in [11]. Since $B_5 \subseteq B_6 \subseteq B_7$, the staircase is non-increasing in the ball, so [11, Prop. 5.1] already gives $u_{B_6}(n) \le 2$ and $u_{B_7}(n) \le 2$ for $2 \le n \le 10$. What is new is the range $11 \le n \le 13$ in B_6 and $11 \le n \le 15$ in B_7; at $n = 11$ it rests on Theorem 4.1, and from $n = 12$ on it is out of reach of the published data altogether. In particular the cell $u_{B_6}(8) \le 2$, which our sweep also certifies, is not new: it follows from $u_{B_5}(8) = 2$ by the same monotonicity, and [11] states the interval $\{1, 2\}$ for it.

Remark 5.4. Theorem 3.4 bounds how far Corollary 3.3 can be pushed in B_5 and B_6; $L(7)$ is not among the machine-checked values, so the range $2 \le n \le 15$ of Theorem 5.2 is not claimed to be the limit of the method in B_7.

6. A parity obstruction for the remaining cell

The cell $u_{B_6}(8) \in \{1, 2\}$ of [11] asks for an 8-element subset of the 933-element ball whose square has *exactly one* uniquely represented element. The following observation, which we state as a remark because it is elementary and outside the machine-checked development in its general form, restricts where such a set can be.

Remark 6.1 (parity). Let G be torsion-free and let $A = A^{-1} \subseteq G$ be finite with $|A| \ge 2$. Then $u(A)$ is even. Indeed $(a, b) \mapsto (b^{-1}, a^{-1})$ is a bijection of $A \times A$ carrying a pair with product w to a pair with product w^{-1}, so w and w^{-1} have equal multiplicities and $\mathcal{U}(A)$ is closed under inversion. A fixed point $w = w^{-1}$ satisfies $w^2 = 1$, hence $w = 1$ as G is torsion-free; and 1 has the $|A| \ge 2$ representations $a \cdot a^{-1}$, so $1 \notin \mathcal{U}(A)$. Thus inversion is a fixed-point-free involution of $\mathcal{U}(A)$ and $|\mathcal{U}(A)|$ is even. Consequently *an n-set realising $u = 1$ is never inverse-closed*: the search for the missing cell may be restricted to sets with $A \ne A^{-1}$.

A window on an interval I with $I = -I$ is inverse-closed, so for such a window the remark already forbids an odd value; the full-width windows $\{x^k : |k| \le r\}$ and $\{y^k : |k| \le 5\}$ used above are of that shape, though their proper prefixes are not. Four instances are part of the machine-checked development — the windows of lengths 11 and 13 are inverse-closed with u even, while A_{12} and A_1 are not inverse-closed, A_{12} having odd u — but the general statement is a hand proof.

It does not explain the phenomenon [11] reports, that the value 1 is never attained in the studied balls of G_1, of Promislow's group P and of the Fibonacci group $H_4 = F(3, 4)$; of the last two that paper says that there "the two-unique-products property … and the unique-product property fail simultaneously, at the same critical size". The same paper exhibits a set with $u = 1$ in G_2 (its Proposition 4.2, at $n = 9$ in B_2), so no general law is available. Of the eleven B_5 minimisers shipped with [11], exactly one — the one at $n = 3$ — is inverse-closed, by a check in the reference implementation; the other ten are not, among them the $u = 2$ minimiser at $n = 8$ and both $u = 15$ sets at $n = 11, 12$. So the parity restriction applies to exactly one of them and excludes none. Remark 6.1 is a restriction on where a $u = 1$ witness can live, not a reason why none exists.

7. What remains open, and what it would cost

$u_{B_6}(8) \in \{1, 2\}$**.** This is the sharpest cell in [11], and it is a decision problem over the $\binom{933}{8} \approx 1.4{\cdot}10^{19}$ eight-subsets of B_6; the run reported there returned `UNKNOWN` after 21 633 seconds of eight-core constraint solving, and its ancillary log records the honest conclusion. Our contribution is negative evidence only, from the reference implementation and outside the machine-checked development. A steepest-descent local search with random restarts and sideways moves over 8-subsets of B_6, run to a hard stop on four independent seeds for a total of about thirteen minutes of CPU — with an earlier single-seed run of comparable length giving the same answer — never reached a value below 2; every seed bottomed out at 2 within the first tenth of a second. The optima found are isolated: a complete 1-swap exploration around the $u = 2$ minimiser of [11, Prop. 5.1], embedded in B_6, finds no neighbour with $u \leq 2$ among the 7 400 of them, the best neighbouring value being 12. Deciding the cell needs an infeasibility proof, not a search.

One natural decomposition is priced and does not help, and we record it so that it is not re-priced. Since H is normal of index 32, an 8-set $A \subseteq G_1$ has a coset profile $m : G_1/H \to \mathbb{Z}_{\geq 0}$ with $\sum m = 8$, and if a product of cosets is uniquely represented in the order-32 quotient then the corresponding product is uniquely represented in G_1; hence $\#\{w \in G_1/H : \sum_{qq'=w} m(q)m(q') = 1\} \leq u(A)$. Enumerating all $\binom{39}{8} = 61\,523\,748$ profiles leaves 5 778 052 of them (9.4%) admissible for $u \leq 1$ and 3 910 756 (6.4%) for $u = 0$: the finite quotient fails the unique product property so easily that it carries almost no information, and the search does not split.

$u_{B_5}(12)$**: is it** 2 **or** 3**?** By Theorem 4.2 the answer is one of the two. By Theorem 3.4 a witness for 2 would be a configuration of a genuinely new type, not a power window, and by the computation reported after Theorem 4.2 it would meet each of the two length-11 windows of B_5 in at most nine elements. Settling it is again an infeasibility question, over $\binom{401}{12} \approx 3 \cdot 10^{22}$ subsets.

Two questions of [11] **we do not touch.** The comparison table of [11] carries a "?" for the number of minimal 8-element non-UP sets in $B_7 \subseteq G_1$, which that paper calls beyond its solver budget; a solve-and-block enumeration over $\binom{1935}{8} \approx 5 \cdot 10^{21}$ subsets is well beyond ours. And the sharpness of the Nielsen–Soelberg profile bound — whether a two-sided witness with $(|A|, |B|) = (7, 9)$ exists in any torsion-free group — has no finite core to attack.

A lead for the companion paper. For Promislow's group P, [12] records $u(n) = 2$ for $2 \leq n \leq 9$ in $B_3 \subseteq$ P (41 elements), then 4 at $n = 10, 11$, then 2 at $n = 12, 13$ and 0 for $14 \leq n \leq 17$, and says verbatim: "We do not have a structural explanation for the bump at $n \in \{10, 11\}$, nor for its echo in the profile spike $\beta(15) = 15$". In an independent implementation of P as affine maps of $\mathbb{R}^3$ (reproducing its ball sizes $1, 5, 17, 41, 83, 147, 239$, whose $|B_3| = 41$ matches [12]) the longest power windows are 7 in B_3, 9 in B_4 and 11 in B_5, again normalised to contain the identity. So in P, unlike in G_1, the flat row runs two steps *past* the window: Lemma 3.2 accounts for $n \leq 7$ only, and whatever holds u at 2 for $n = 8, 9$ is also the thing that would have to break at $n = 10$. These computations are from the reference implementation and are not part of the machine-checked development; we record them as a lead, not a result.

8. Formal verification

Every count asserted in Propositions 2.1 and 2.2 and in Theorems 3.4, 4.1, 4.2, 5.1 and 5.2 is formalised and machine-checked in Lean 4 [5]. The exceptions are stated as such in the text: Lemma 3.2, Corollary 3.3 and Remark 6.1 are hand proofs, and so are the identification of the maximisers in Theorem 3.4 and every computation explicitly labelled as belonging to the reference implementation (the associativity fuzz of Section 2, the coincidence pattern and geodesic lengths of A_1, the re-evaluation of the eleven published minimisers, the inverse-closedness check on those minimisers in Section 6, the two exhaustive searches around the length-11 windows, the local search and coset-profile count of Section 7, and the data for Promislow's group). Note that the staircase values do not depend on the lemma: each is an independent evaluation of u on an explicit finite set. The development is four files — the model, the staircase witnesses, the balls and window lengths, and the controls — and uses core Lean only, with no mathematical library, so that kernel evaluation stays cheap; elements are pairs of a natural number and four integers, and u is computed by materialising the n^2 products. There is no incomplete proof. The axiom set of all 39 theorems was printed: none uses choice or quotient soundness, and the two headline results, Theorems 4.1 and 4.2, are kernel evaluations depending on propositional extensionality alone. Nine statements rest on compiled evaluation rather than on kernel reduction — the two coset-table sweeps over the order-32 quotient, the two staircase sweeps of Section 5, the two ball-size computations, the two window-length sweeps over a whole ball, and one control — and each carries its own named compiled-evaluation axiom. As controls, the development also refutes the neighbouring counts of both headline values ($u \neq 1$ and $u \neq 3$ for the eleventh cell, $u \neq 2$ and $u \neq 4$ for the twelfth), checks that the ball certificate can fail (the words of A_1 do not fit radius 6; a permuted or truncated word list is rejected), exhibits an 8-subset of B_5 with all 64 products distinct so that u is visibly not constant, and verifies that the model is neither trivial nor abelian, so that the relator checks are not statements about a degenerate quotient. One clean pass over the four modules took about five minutes of CPU time and peaked at three gigabytes of resident memory, the controls dominating both. The repository is private; repository reference to be added at submission.

Acknowledgement of sources

The models, tables, witnesses and open cells this paper builds on are those published by Tabei [11, 12]; the ancillary data of [11] is what made an independent implementation possible at all.

References

[1] G. A. Freiman, *Foundations of a Structural Theory of Set Addition*, Translations of Mathematical Monographs 37, American Mathematical Society, Providence, RI, 1973. Cited, together with [6], for the ancestry of Lemma 3.2: the extremal case $|P + P| = 2|P| - 1$ of small doubling in $\mathbb{Z}$ and its characterisation of arithmetic progressions. Neither citation is load-bearing — the lemma is proved in full above — and we attach no theorem number to either book.

[2] S. Gadgil and A. R. Tadipatri, *Formalizing Giles Gardam's Disproof of Kaplansky's Unit Conjecture*, in: Proceedings of the 13th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP '24), ACM, 2024, pp. 177–189; DOI `10.1145/3636501.3636947`. The associated Lean 4 development (`Polylean`) constructs Promislow's group, proves it torsion-free and exhibits Gardam's units; it is the closest formal-verification precedent to the present work. Its 33 Lean files define no unique-product or non-UP set and no Nielsen–Soelberg group; this was checked by search over the repository source at commit `b2379c5` (22 April 2025), the head of `github.com/siddhartha-gadgil/Polylean` on 3 September 2026, whose only occurrence of the string "Nielsen" is the imported Mathlib module `GroupTheory.FreeGroup.NielsenSchreier`.

[3] G. Gardam, *A counterexample to the unit conjecture for group rings*, Ann. of Math. (2) **194** (2021), no. 3, 967–979; arXiv:2102.11818.

[4] I. Kaplansky, *Problems in the theory of rings*, Report of a Conference on Linear Algebras, June 1956, Publ. 502, National Academy of Sciences, Washington, D.C., 1957, pp. 1–3. Cited through [11], whose bibliography carries this entry; not consulted directly.

[5] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science, Springer, 2021, pp. 625–635; DOI `10.1007/978-3-030-79876-5_37`.

[6] M. B. Nathanson, *Additive Number Theory: Inverse Problems and the Geometry of Sumsets*, Graduate Texts in Mathematics 165, Springer, New York, 1996.

[7] P. P. Nielsen and L. Soelberg, *Small sets without unique products in torsion-free groups*, J. Algebra Appl. **23** (2024), no. 8, 2550050; DOI `10.1142/S0219498825500501`. The source of the bound $|A| \geq 8$ (its Theorem 1.2), of the presentation (3.1) of G_1, of the subgroup H and its presentation (3.2), of the witness A_1 of (3.4), and of the coincidence fingerprint of §4. Section and item numbers are those of the first author's copy, posted at `mathdept.byu.edu/~pace/KaplanskyConjecture_web.pdf`, dated 22 September 2023 and read on 3 September 2026; the published version was not accessible to us and its numbering has not been compared, though the bibliographic data above were confirmed against Crossref on the same date.

[8] D. S. Passman, *The Algebraic Structure of Group Rings*, Wiley, New York, 1977. Cited through [11]; not consulted directly.

[9] S. D. Promislow, *A simple example of a torsion-free, non unique product group*, Bull. London Math. Soc. **20** (1988), no. 4, 302–304.

[10] E. Rips and Y. Segev, *Torsion-free group without unique product property*, J. Algebra **108** (1987), no. 1, 116–126. Cited through [7, 11]; not consulted directly.

[11] M. Tabei, *The quantitative non-unique-product landscape at the global minimum: the Nielsen–Soelberg groups*, arXiv:2607.19687, math.GR. All quotations, ball sizes, tables and open cells used here are from the source of v2 (30 August 2026), retrieved together with its ancillary data on 3 September 2026; the listing shows v1 of 22 July 2026 and v2 of 30 August 2026, and v2 is the current version. The internal numbering cited here — Theorem 3.1 (localization in G_1), Proposition 4.2 (staircase of G_2) and Proposition 5.1 (staircase of G_1 below the localization radius) — was read off the compiled PDF of that source, a 16-page document; it is pinned to v2 and may shift in a later version.

[12] M. Tabei, *Least sizes of non-unique-product sets: the Promislow group and a Heisenberg-type candidate*, arXiv: 2607.18346, math.GR. Read from the source of v1 (20 July 2026) on 3 September 2026; the listing shows no later version. On the same date a citation index reported exactly one work citing [11], namely this companion; so no third paper has yet touched the cells discussed here.