

Star and Support-Image Profiles for Class-Two p -Groups

Luca Blanchi

June 2026

Abstract

Let G be a finite group. The higher commuting probabilities

$$P_m(G) = \Pr(g_1, \dots, g_m \text{ commute pairwise}) = \frac{|\text{Hom}(\mathbb{Z}^m, G)|}{|G|^m}$$

measure the abundance of abelian m -tuples in G . We compare these probabilities with several other natural quadratic observables. The first family is the star commuting hierarchy

$$\Sigma_{r,t}(G) = \Pr([x_i, y_j] = 1 \ \forall i, j) = \frac{|\text{Hom}(F_r \times F_t, G)|}{|G|^{r+t}}.$$

We show that $\Sigma_{r,t}(G)$ is a moment of common-centralizer sizes. For groups of class 2 and exponent p , represented by alternating maps $\beta : V \times V \rightarrow W$, the star hierarchy becomes linear algebra: if $A_\beta : V \rightarrow W^r$ is the common contraction map and $N_r^\beta(a)$ counts r -tuples for which $\text{rank } A_\beta = a$, then $\Sigma_{r,t}(G_\beta) = |V|^{-r} \sum_a N_r^\beta(a) p^{-ta}$. By contrast, $P_m(G_\beta)$ is determined by the numbers of totally isotropic subspaces of V . We prove that these two kinds of information are genuinely distinct by constructing, for every odd prime p , two directly indecomposable special groups $G_A(p)$ and $G_B(p)$ of order p^7 , class 2, and exponent p , with identical P_m for all m but different N_1 . We then identify the star hierarchy with rank-support enumerators of the rank-metric code of contractions, prove a positive completeness theorem for complete multipartite graph blow-up groups, introduce canonical support-image profiles CSI_r , and close with the Centralizer Hom Transform.

Declaration of generative AI and AI-assisted technologies in the writing process

This article belongs to the Presentation theory section. Generative AI, including ChatGPT by OpenAI, was used to assist with mathematical drafting, formalization, review, and editing. The note may be preliminary, may have been only partially reviewed, and in some cases may be only AI-reviewed.

1 Introduction

The commuting probability of a finite group is

$$\text{cp}(G) = \Pr([x, y] = 1) = \frac{|\text{Hom}(\mathbb{Z}^2, G)|}{|G|^2}.$$

A natural hierarchy of refinements is given by

$$P_m(G) = \Pr(g_1, \dots, g_m \text{ commute pairwise}) = \frac{|\text{Hom}(\mathbb{Z}^m, G)|}{|G|^m}.$$

These are the higher commuting probabilities. They count homomorphisms from free abelian groups.

This paper studies a different refinement. Instead of asking for all elements in a tuple to commute pairwise, we ask for two tuples to commute across:

$$[x_i, y_j] = 1 \quad (1 \leq i \leq r, 1 \leq j \leq t).$$

The resulting probabilities are

$$\Sigma_{r,t}(G) = \Pr([x_i, y_j] = 1 \ \forall i, j) = \frac{|\text{Hom}(F_r \times F_t, G)|}{|G|^{r+t}}.$$

Thus P_m counts homomorphisms from $\mathbb{Z}^m$, while $\Sigma_{r,t}$ counts homomorphisms from direct products of free groups.

The main contribution of the paper is not a total ordering of all these invariants. Rather, we compare several natural families of quadratic observables and show that they retain different aspects of the alternating geometry of a class-two exponent- p group.

Such groups are governed by alternating bilinear maps

$$\beta : V \times V \rightarrow W.$$

For G_β , the higher commuting probabilities P_m are controlled by totally isotropic subspaces of V . The star probabilities $\Sigma_{r,t}$, on the other hand, are controlled by the ranks of common contraction maps

$$A_v : V \rightarrow W^r.$$

These two data sets are genuinely different. The uniform family constructed in Section 6 shows that P_m for all m does not determine N_1 , even on directly indecomposable groups.

The rank side has an independent algebraic meaning. The contraction maps

$$A_v : u \mapsto \beta(v, u)$$

form a rank-metric code

$$C_\beta \leq \text{Hom}(V, W),$$

and the profiles N_r are equivalent to the higher rank-support enumerators of this code. Thus star probabilities connect finite-group word observables to rank-metric support theory and q -polymatroidal invariants.

The canonical support-image profile CSI_r is introduced as a further codomain-localized refinement. It remembers the $\text{GL}(W)$ -orbit of the image subspace

$$\text{im } A_v \leq W^r,$$

rather than only its dimension or the dimensions of its coordinate projections.

The Centralizer Hom Transform is different in nature. It does not refine CSI in any formal sense asserted here; instead, it replaces centralizer-size moments by homomorphism counts into centralizers, and therefore records internal isomorphism types of common centralizers.

2 Star commuting probabilities

Let G be a finite group. For $r, t \geq 0$, define

$$\Sigma_{r,t}(G) = \Pr([x_i, y_j] = 1 \ \forall i, j),$$

where

$$x_1, \dots, x_r, y_1, \dots, y_t$$

are independent uniformly distributed elements of G . We use the convention

$$\Sigma_{r,0}(G) = \Sigma_{0,t}(G) = 1.$$

Let

$$\mathbf{x} = (x_1, \dots, x_r) \in G^r.$$

Write

$$C_G(\mathbf{x}) = \bigcap_{i=1}^r C_G(x_i).$$

Proposition 2.1 (Centralizer moment formula). *For every finite group G ,*

$$\Sigma_{r,t}(G) = \frac{1}{|G|^r} \sum_{\mathbf{x} \in G^r} \left(\frac{|C_G(\mathbf{x})|}{|G|} \right)^t.$$

Proof. Fix $\mathbf{x} = (x_1, \dots, x_r)$. The condition

$$[x_i, y_j] = 1 \quad \forall i, j$$

is equivalent to

$$y_j \in C_G(\mathbf{x}) \quad \forall j.$$

Thus, for fixed $\mathbf{x}$, the conditional probability is

$$\left(\frac{|C_G(\mathbf{x})|}{|G|} \right)^t.$$

Averaging over all $\mathbf{x} \in G^r$ gives the formula. □

A homomorphism

$$F_r \times F_t \rightarrow G$$

is precisely a pair of tuples whose entries commute across the two factors. Hence

$$\Sigma_{r,t}(G) = \frac{|\text{Hom}(F_r \times F_t, G)|}{|G|^{r+t}}.$$

The case $r = t = 1$ is the usual commuting probability.

3 Class-two exponent-(p) groups

Let p be an odd prime. Let V, W be finite-dimensional vector spaces over $\mathbb{F}_p$, and let

$$\beta : V \times V \rightarrow W$$

be an alternating bilinear map.

Define

$$G_\beta = V \oplus W$$

with multiplication

$$(v, z)(u, w) \left(v + u, z + w + \frac{1}{2}\beta(v, u) \right).$$

Then

$$[(v, z), (u, w)](0, \beta(v, u)).$$

Thus G_β has nilpotency class at most 2 and exponent p .

For

$$v \in V,$$

define the contraction

$$A_v : V \rightarrow W, \quad A_v(u) = \beta(v, u).$$

For an r -tuple

$$\mathbf{v} = (v_1, \dots, v_r) \in V^r,$$

define

$$A_{\mathbf{v}} : V \rightarrow W^r, \quad u \mapsto (A_{v_1}u, \dots, A_{v_r}u).$$

Let

$$N_r^\beta(a) = |\mathbf{v} \in V^r : \text{rank } A_{\mathbf{v}} = a|.$$

Proposition 3.1 (Star probabilities as rank moments). *For G_β ,*

$$\Sigma_{r,t}(G_\beta) = |V|^{-r} \sum_a N_r^\beta(a) p^{-ta}.$$

Proof. Let

$$x_i = (v_i, z_i) \in G_\beta.$$

The common centralizer of the tuple $x_1, \dots, x_r$ consists of all

$$(u, w) \in V \oplus W$$

such that

$$\beta(v_i, u) = 0 \quad \forall i.$$

Thus

$$|C_{G_\beta}(x_1, \dots, x_r)| = |W| \cdot |\ker A_{\mathbf{v}}| = |W| \cdot p^{\dim V - \text{rank } A_{\mathbf{v}}}.$$

Since

$$|G_\beta| = |V||W|,$$

we obtain

$$\frac{|C_{G_\beta}(x_1, \dots, x_r)|}{|G_\beta|} p^{-\text{rank } A_{\mathbf{v}}}.$$

The central coordinates z_i do not affect the rank. Substituting into the centralizer moment formula gives

$$\Sigma_{r,t}(G_\beta) = |V|^{-r} \sum_{\mathbf{v} \in V^r} p^{-t \text{rank } A_{\mathbf{v}}},$$

which is the stated formula. □

Thus the values $\Sigma_{r,t}$, as t varies, are precisely the moments of the rank distribution N_r^β .

4 Higher commuting probabilities and isotropic subspaces

A subspace

$$U \leq V$$

is totally isotropic for β if

$$\beta(U, U) = 0.$$

Let

$$I_d(\beta)$$

be the number of d -dimensional totally isotropic subspaces of V .

Proposition 4.1 (P_m and isotropic subspaces). *For G_β ,*

$$P_m(G_\beta) = |V|^{-m} \sum_d I_d(\beta) \prod_{i=0}^{d-1} (p^m - p^i).$$

Proof. An m -tuple in G_β commutes pairwise if and only if its image in V spans a totally isotropic subspace.

Fix a d -dimensional totally isotropic subspace $U \leq V$. The number of ordered m -tuples spanning U is the number of surjective linear maps

$$\mathbb{F}_p^m \twoheadrightarrow U,$$

namely

$$\prod_{i=0}^{d-1} (p^m - p^i).$$

The central coordinates in W are arbitrary and cancel in the normalization by $|G_\beta|^m$. Summing over all totally isotropic U gives the formula. $\square$

Therefore the entire hierarchy P_m depends only on the isotropic-subspace counts

$$I_d(\beta).$$

The star hierarchy sees different data: ranks of common contraction maps.

5 Rank-support enumerators

Assume in this section that the radical of β is zero, so that the map

$$V \rightarrow \text{Hom}(V, W), \quad v \mapsto A_v$$

is injective. Let

$$C_\beta = A_v : v \in V \leq \text{Hom}(V, W)$$

be the corresponding rank-metric code of contractions.

For a subcode

$$D \leq C_\beta,$$

define its row-support weight by

$$\text{wt}_R(D) = \dim V - \dim \bigcap_{A \in D} \ker A.$$

For $j, a \geq 0$, define

$$A_{j,a}(C_\beta) = |D \leq C_\beta : \dim D = j, \text{wt}_R(D) = a|.$$

Theorem 5.1 (Star profiles and rank-support enumerators). *For all r, a ,*

$$N_r^\beta(a) = \sum_{j=0}^{\min(r, \dim C_\beta)} \gamma_{r,j}(p) A_{j,a}(C_\beta),$$

where

$$\gamma_{r,j}(p) = \prod_{i=0}^{j-1} (p^r - p^i).$$

Consequently, the family

$$(N_r^\beta(a))_{r,a}$$

determines, and is determined by, the family

$$(A_{j,a}(C_\beta))_{j,a}.$$

Proof. For an r -tuple

$$\mathbf{v} = (v_1, \dots, v_r),$$

let

$$D(\mathbf{v}) = \langle A_{v_1}, \dots, A_{v_r} \rangle \leq C_\beta.$$

Then

$$\text{rank } A_{\mathbf{v}} = \dim V - \dim \bigcap_{i=1}^r \ker A_{v_i} = \text{wt}_R(D(\mathbf{v})).$$

Fix a j -dimensional subcode $D \leq C_\beta$. The number of ordered r -tuples generating D is the number of surjective linear maps

$$\mathbb{F}_p^r \twoheadrightarrow D,$$

namely

$$\gamma_{r,j}(p) = \prod_{i=0}^{j-1} (p^r - p^i).$$

Summing over all D of dimension j and row-support weight a gives the stated formula. The matrix $(\gamma_{r,j})$ is triangular in r, j , with nonzero diagonal

$$\gamma_{j,j}(p) = |\text{GL}_j(\mathbb{F}_p)|.$$

Hence the relation is invertible. □

This identifies the star hierarchy with higher rank-support enumerators of the contraction code.

6 An infinite separation between P_m and N_1

We now show that the higher commuting hierarchy P_m does not determine even the first star profile N_1 .

Let

$$V = \mathbb{F}_p^4$$

with basis

$$e_1, e_2, e_3, e_4.$$

Use the ordered basis

$$e_{12}, e_{13}, e_{14}, e_{23}, e_{24}, e_{34}$$

of $\Lambda^2 V$.

Define subspaces

$$K_A = \langle e_{12}, e_{34}, e_{13} + e_{24} \rangle$$

and

$$K_B = \langle e_{12}, e_{14}, e_{13} + e_{24} \rangle.$$

Let

$$W_A = \Lambda^2 V / K_A, \quad W_B = \Lambda^2 V / K_B,$$

and let

$$\beta_A, \beta_B : \Lambda^2 V \rightarrow W_A, W_B$$

be the quotient maps.

Lemma 6.1 (The K_A -isotropic planes are disjoint outside zero). *Let*

$$K_A = \langle e_{12}, e_{34}, e_{13} + e_{24} \rangle.$$

The $p + 1$ two-dimensional subspaces of V that are totally isotropic for β_A intersect pairwise only in 0.

Proof. A two-dimensional subspace $U \leq V$ is totally isotropic for β_A precisely when its Plucker point lies in

$$\mathbb{P}(K_A) \cap \mathcal{K},$$

where $\mathcal{K}$ is the Klein quadric.

For K_A , this intersection is the nondegenerate conic

$$ac - b^2 = 0$$

in $\mathbb{P}(K_A) \cong \mathbb{P}^2$. In particular, it contains no projective line.

Suppose two distinct isotropic planes $U_1, U_2 \leq V$ intersect in a nonzero vector v . Choose vectors u_1, u_2 such that

$$U_i = \langle v, u_i \rangle.$$

Then the Plucker points

$$[v \wedge u_1], \quad [v \wedge u_2]$$

are two distinct points of

$$\mathbb{P}(K_A) \cap \mathcal{K}.$$

Moreover, the projective line joining them is

$$\{[v \wedge (\lambda u_1 + \mu u_2)] : [\lambda : \mu] \in \mathbb{P}^1\}.$$

Every point on this line is decomposable, hence lies on the Klein quadric. Since the endpoints lie in $\mathbb{P}(K_A)$, the whole line lies in $\mathbb{P}(K_A) \cap \mathcal{K}$.

This contradicts the fact that $\mathbb{P}(K_A) \cap \mathcal{K}$ is a nondegenerate conic and contains no line. Therefore distinct isotropic planes intersect only in 0. $\square$

Theorem 6.2 (Infinite directly indecomposable separation). *For every odd prime p ,*

$$P_m(G_{\beta_A}) = P_m(G_{\beta_B}) \quad \forall m \geq 1,$$

but

$$N_1^{\beta_A} \neq N_1^{\beta_B}.$$

Moreover, both groups G_{β_A} and G_{β_B} are directly indecomposable.

Proof. The planes in V are parametrized by the Klein quadric in

$$\mathbb{P}(\Lambda^2 V),$$

given in Plucker coordinates by

$$x_{12}x_{34} - x_{13}x_{24} + x_{14}x_{23} = 0.$$

A plane is totally isotropic for β_X if and only if its Plucker point lies in

$$\mathbb{P}(K_X)$$

and on the Klein quadric.

For K_A , a vector has coordinates

$$(a, b, 0, 0, b, c),$$

and the Klein equation restricts to

$$ac - b^2 = 0.$$

This is a nondegenerate conic in $\mathbb{P}^2(\mathbb{F}_p)$, hence has $p + 1$ points.

For K_B , a vector has coordinates

$$(a, b, c, 0, b, 0),$$

and the Klein equation restricts to

$$-b^2 = 0.$$

This is a line in $\mathbb{P}^2(\mathbb{F}_p)$, again with $p + 1$ points.

Thus β_A and β_B have the same number of totally isotropic planes. Every one-dimensional subspace is isotropic because β is alternating. Since there are only $p + 1$ isotropic planes, no three-dimensional subspace can be totally isotropic, as any three-dimensional isotropic subspace would contain more than $p + 1$ planes.

Therefore the full sequence

$$I_d(\beta_A)$$

agrees with

$$I_d(\beta_B).$$

By Proposition 4.1,

$$P_m(G_{\beta_A}) = P_m(G_{\beta_B})$$

for all m .

It remains to compute N_1 . For $0 \neq v \in V$,

$$\text{rank } A_v = 3 - \dim(K_X \cap (v \wedge V)),$$

because $v \wedge V$ has dimension 3 and $W_X = \Lambda^2 V / K_X$.

By Lemma 6.2, the nonzero vectors lying in the isotropic planes for K_A are exactly $(p+1)(p^2-1)$. Equivalently, the intersection of $\mathbb{P}(K_A)$ with the Klein quadric is a nondegenerate conic, which contains no projective line. Hence no nonzero v gives a two-dimensional intersection

$$K_A \cap (v \wedge V).$$

Thus there are no rank-one contractions.

The $p+1$ isotropic planes are pairwise disjoint outside 0; otherwise the corresponding points of the conic would lie on a projective line contained in the Klein quadric and in $\mathbb{P}(K_A)$. Therefore the nonzero vectors lying in isotropic planes number

$$(p+1)(p^2-1).$$

For these vectors the contraction rank is 2, and for the remaining nonzero vectors it is 3. Hence

$$(c_0, c_1, c_2, c_3)_A \left(1, 0, (p+1)(p^2-1), p(p-1)^2(p+1) \right).$$

For K_B , the $p+1$ isotropic planes form a pencil through the line

$$\langle e_1 \rangle.$$

The $p-1$ nonzero vectors on this line have contraction rank 1. The vectors lying in one of the isotropic planes but not on the common line have contraction rank 2, and their number is

$$(p+1)(p^2-p) = p(p^2-1).$$

The remaining nonzero vectors have rank 3. Thus

$$(c_0, c_1, c_2, c_3)_B \left(1, p-1, p(p^2-1), p^3(p-1) \right).$$

Therefore $N_1^{\beta_A} \neq N_1^{\beta_B}$.

Finally, we prove direct indecomposability. Both maps are surjective and radical-free, so the corresponding groups are special. A direct product decomposition of a special class-two group induces an orthogonal decomposition

$$V = U_1 \oplus U_2$$

of the alternating map, with each factor contributing nontrivially to the commutator image. No U_i can be one-dimensional. Since $\dim V = 4$, any nontrivial such decomposition would have

$$\dim U_1 = \dim U_2 = 2.$$

But then each $\beta(\Lambda^2 U_i)$ has dimension at most 1, so the total commutator image has dimension at most 2, contradicting

$$\dim W_A = \dim W_B = 3.$$

Thus both groups are directly indecomposable. □

This proves a uniform separation

$$P_m < N_1$$

on directly indecomposable groups.

7 Complete multipartite blow-ups

We now give a positive classification result.

Let

$$V = V_1 \oplus \cdots \oplus V_k, \quad \dim V_i = n_i, \quad k \geq 2.$$

Let

$$W = \bigoplus_{1 \leq i < j \leq k} V_i \otimes V_j.$$

Define

$$\beta(v, u)_{ij} = v_i \otimes u_j - u_i \otimes v_j.$$

This is the alternating map associated with the complete multipartite blow-up.

Let

$$n = \sum_i n_i.$$

Lemma 7.1 (Rank formula). *For $v \in V$,*

$$\text{rank } A_v = \begin{cases} 0, & v = 0, \\ n - n_i, & \text{supp}(v) = \{i\}, \\ n - 1, & |\text{supp}(v)| \geq 2. \end{cases}$$

Proof. If $v = 0$, the contraction is zero.

Suppose v is supported only in V_i . Then

$$A_v(u) = 0$$

if and only if

$$u_j = 0 \quad (j \neq i).$$

Thus

$$\ker A_v = V_i$$

and

$$\text{rank } A_v = n - n_i.$$

Now suppose v has nonzero components in at least two blocks. If $j \notin \text{supp}(v)$, then for some $i \in \text{supp}(v)$, the edge component gives

$$v_i \otimes u_j = 0,$$

hence

$$u_j = 0.$$

If $i, j \in \text{supp}(v)$, the equation

$$v_i \otimes u_j = u_i \otimes v_j$$

forces

$$u_i = \lambda v_i, \quad u_j = \lambda v_j$$

with the same scalar λ . Therefore

$$u = \lambda v.$$

Hence

$$\ker A_v = \langle v \rangle$$

and

$$\text{rank } A_v = n - 1.$$

□

Theorem 7.2 (Completeness of N_1 for complete multipartite blow-ups). *For complete multipartite blow-ups, N_1 determines the multiset*

$$n_1, \dots, n_k.$$

Consequently, N_1 determines the corresponding group.

Proof. From

$$\sum_a N_1(a) = |V| = p^n$$

we recover n .

For $s \geq 2$, let

$$m_s = |\{i : n_i = s\}|.$$

By the rank formula, the number of nonzero vectors supported in a block of dimension s is

$$m_s(p^s - 1),$$

and these vectors have rank

$$n - s.$$

Thus

$$m_s = \frac{N_1(n-s)}{p^s - 1} \quad (s \geq 2).$$

Finally,

$$m_1 = n - \sum_{s \geq 2} s m_s.$$

Hence the multiset of block dimensions is recovered.

The alternating map is determined by this multiset up to equivalence, so the group is determined. $\square$

Proposition 7.3 (Direct indecomposability). *Every complete multipartite blow-up group with $k \geq 2$ is directly indecomposable.*

Proof. A direct product decomposition would induce an orthogonal decomposition

$$V = U_1 \oplus U_2.$$

If a nonzero vector of U_1 has support in at least two blocks, then by Lemma 7.1 its contraction kernel is one-dimensional. Orthogonality would force

$$U_2 \subseteq \ker A_v = \langle v \rangle,$$

which is impossible in a nontrivial decomposition.

Thus every nonzero vector in U_1 is supported in a single block. By linearity, U_1 is contained in some V_i . Similarly, U_2 is contained in some V_j . If $i \neq j$, then

$$\beta(U_1, U_2) \neq 0.$$

If $i = j$, then

$$U_1 + U_2$$

does not contain the other blocks. Since $k \geq 2$, this cannot be all of V . Contradiction. $\square$

Thus N_1 is complete on a natural infinite directly indecomposable class.

8 Flag profiles

The profile N_r records only the final rank of

$$A_{\mathbf{v}} : V \rightarrow W^r.$$

It forgets how ranks are distributed among subtuples.
For

$$\mathbf{v} = (v_1, \dots, v_r)$$

and nonempty

$$S \subseteq [r],$$

define

$$\rho_{\mathbf{v}}(S) = \text{rank } A_{\mathbf{v}_S}.$$

The aggregated flag profile of level r is the distribution of the functions

$$S \mapsto \rho_{\mathbf{v}}(S)$$

as $\mathbf{v}$ ranges over V^r .

For $r = 2$, this is the collection of numbers

$$M_{a,b,c} = |\{(v, w) : \text{rank } A_v = a, \text{ rank } A_w = b, \text{ rank } A_{v,w} = c\}|.$$

Flag profiles refine N_r , but they remain aggregated: they remember only dimensions of projected images, not the actual position of the image subspaces in W^r .

The next section introduces a compact codomain-localized refinement.

9 Canonical support-image profiles

For

$$\mathbf{v} = (v_1, \dots, v_r) \in V^r,$$

define

$$I_{\beta}(\mathbf{v}) = \text{im } A_{\mathbf{v}} \leq W^r.$$

For a subspace

$$T \leq W^r,$$

define

$$E_r^{\beta}(T) = |\{\mathbf{v} \in V^r : I_{\beta}(\mathbf{v}) = T\}|.$$

The group $\text{GL}(W)$ acts diagonally on W^r , hence on its subspaces.

Definition 9.1 (CSI). *The canonical support-image profile of level r , denoted*

$$\text{CSI}_r(\beta),$$

is the function on orbits

$$\mathcal{O} \in \text{Sub}(W^r)/\text{GL}(W)$$

given by

$$\text{CSI}_r(\beta)(\mathcal{O}) = \sum_{T \in \mathcal{O}} E_r^{\beta}(T).$$

For $R \geq 1$, define

$$\text{CSI}_{\leq R}(\beta)(\text{CSI}_1(\beta), \dots, \text{CSI}_R(\beta)).$$

Proposition 9.2 (Invariance). *If β and β' are equivalent under*

$$GL(V) \times GL(W),$$

then

$$\text{CSI}_r(\beta) = \text{CSI}_r(\beta')$$

for all r .

Proof. Suppose

$$\beta'(Pv, Pu) = Q\beta(v, u)$$

with

$$P \in GL(V), \quad Q \in GL(W).$$

Then

$$I_{\beta'}(P\mathbf{v}) = Q^{\oplus r} I_{\beta}(\mathbf{v}).$$

The map

$$\mathbf{v} \mapsto P\mathbf{v}$$

is a bijection of V^r , and $Q^{\oplus r}$ is precisely the diagonal action of $GL(W)$ on W^r . Thus the orbit distribution is preserved. $\square$

Proposition 9.3 (CSI dominates N_r and flag profiles). *The profile $\text{CSI}_r(\beta)$ determines N_r^β and the aggregated flag profile of level r .*

Proof. The rank profile N_r is obtained by summing CSI_r over orbits of subspaces $T \leq W^r$ of a fixed dimension:

$$N_r^\beta(a) = \sum_{\mathcal{O} \text{ dim } T=a} \text{CSI}_r(\beta)(\mathcal{O}).$$

For flag profiles, let

$$\pi_S : W^r \rightarrow W^S$$

be coordinate projection. For

$$T = I_{\beta}(\mathbf{v}),$$

we have

$$\dim \pi_S(T) = \rho_{\mathbf{v}}(S).$$

The function

$$T \mapsto (\dim \pi_S(T))_{\emptyset \neq S \subseteq [r]}$$

is invariant under diagonal $\text{GL}(W)$. Hence the aggregated distribution of flag functions is determined by the CSI orbit distribution. $\square$

Thus CSI is a codomain-localized refinement of both rank and flag profiles.

10 Computing CSI

Let

$$d = \dim W.$$

We now show that CSI_r is fixed-parameter computable in d and r .
For

$$S \leq W^r,$$

define

$$F_r^\beta(S) = |\{\mathbf{v} \in V^r : I_\beta(\mathbf{v}) \leq S\}|.$$

Lemma 10.1. *For every $S \leq W^r$, $F_r^\beta(S)$ is computable by linear algebra in time polynomial in $\dim V$, for fixed d, r .*

Proof. Let

$$\pi_S : W^r \rightarrow W^r/S$$

be the quotient map. The condition

$$I_\beta(\mathbf{v}) \leq S$$

is equivalent to

$$\pi_S \circ A_{\mathbf{v}} = 0.$$

The map

$$\Phi_S : V^r \rightarrow \text{Hom}(V, W^r/S), \quad \mathbf{v} \mapsto \pi_S \circ A_{\mathbf{v}}$$

is linear. Hence

$$F_r^\beta(S) = p^{\dim \ker \Phi_S},$$

and the dimension of the kernel is computed by Gaussian elimination. $\square$

Theorem 10.2 (CSI computation). *For fixed $d = \dim W$ and r , $\text{CSI}_r(\beta)$ is computable in time*

$$p^{O(d^2 r^2)} \text{poly}(\dim V, \log p).$$

Proof. The number of subspaces of W^r , where

$$\dim W^r = dr,$$

is

$$p^{O(d^2 r^2)}.$$

For each subspace $S \leq W^r$, compute

$$F_r^\beta(S)$$

using Lemma 10.1.

We have

$$F_r^\beta(S) = \sum_{T \leq S} E_r^\beta(T).$$

By Mobius inversion on the lattice of subspaces,

$$E_r^\beta(T) = \sum_{S \leq T} \mu(S, T) F_r^\beta(S),$$

where

$$\mu(S, T) = (-1)^c p^{\binom{c}{2}}, \quad c = \dim T - \dim S.$$

This computes all values $E_r^\beta(T)$. Finally, aggregate them over diagonal $\mathrm{GL}(W)$ -orbits on subspaces $T \leq W^r$. Since the relevant ambient dimension is dr , this orbit aggregation also costs

$$p^{O(d^2 r^2)}.$$

□

CSI is therefore a compact substitute for full localization on the domain side. The full profile over all subspaces of V has size roughly

$$p^{\Theta((\dim V)^2)},$$

whereas CSI_r is controlled by the small space W^r .

11 Centralizer Hom Transform

We close with a further refinement, which sees not only sizes of centralizers but their isomorphism types.

Let K be a finitely generated group. Define

$$\mathrm{CHT}_r(G; K) = |\mathrm{Hom}(F_r \times K, G)|.$$

Proposition 11.1 (CHT formula). *For every finite group G ,*

$$\mathrm{CHT}_r(G; K) = \sum_{\mathbf{x} \in G^r} |\mathrm{Hom}(K, C_G(\mathbf{x}))|.$$

Proof. A homomorphism

$$F_r \times K \rightarrow G$$

is determined by the images of the r free generators of F_r , say

$$\mathbf{x} = (x_1, \dots, x_r),$$

together with a homomorphism

$$K \rightarrow G$$

whose image commutes with every x_i . Such a homomorphism has image in

$$C_G(\mathbf{x}).$$

Summing over $\mathbf{x} \in G^r$ gives the formula. □

For $K = F_t$, one recovers

$$\text{CHT}_r(G; F_t) = \sum_{\mathbf{x} \in G^r} |C_G(\mathbf{x})|^t = |G|^{r+t} \Sigma_{r,t}(G).$$

Thus the star hierarchy is the free-group part of CHT.

Theorem 11.2 (CHT inversion). *Let G be finite of order n . For fixed r , the values*

$$\text{CHT}_r(G; K), \quad |K| \leq n,$$

with K finite, determine the multiset

$$C_G(\mathbf{x}) : \mathbf{x} \in G^r / \cong$$

with multiplicities.

Proof. Let

$$\mu_{r,G}(H) = |\{\mathbf{x} \in G^r : C_G(\mathbf{x}) \cong H\}|.$$

Then

$$\text{CHT}_r(G; K) = \sum_H \mu_{r,G}(H) |\text{Hom}(K, H)|.$$

For finite groups K, H , write

$$h_K(H) = |\text{Hom}(K, H)|,$$

and

$$i_K(H) = |\text{Inj}(K, H)|.$$

Every homomorphism $K \rightarrow H$ has kernel $N \trianglelefteq K$ and induces an embedding

$$K/N \hookrightarrow H.$$

Therefore

$$h_K(H) = \sum_{N \trianglelefteq K} i_{K/N}(H).$$

By induction on quotients of K , the CHT values determine

$$\sum_H \mu_{r,G}(H) i_K(H)$$

for every K .

Now order finite groups of order at most n by decreasing order. If

$$|K| > |H|,$$

then

$$i_K(H) = 0.$$

If

$$|K| = |H|,$$

then

$$i_K(H) > 0$$

if and only if

$$K \cong H,$$

and in that case

$$i_K(K) = |\text{Aut}(K)|.$$

Thus the resulting linear system is triangular and determines all multiplicities

$$\mu_{r,G}(H).$$

□

This shows that CHT strictly refines the centralizer-size moment hierarchy.

12 Comparison of the observable families

The observables studied in this paper should not be understood as a single totally ordered chain. Rather, they are several natural families of quadratic observables that retain different kinds of information.

For class-two exponent- p groups G_β , the higher commuting probabilities

$$P_m(G_\beta)$$

are determined by the numbers of totally isotropic subspaces of V . In contrast, the star probabilities

$$\Sigma_{r,t}(G_\beta)$$

are moments of the rank distribution

$$N_r^\beta(a) = |\{\mathbf{v} \in V^r : \text{rank } A_{\mathbf{v}} = a\}|.$$

Thus P_m and N_r measure different geometric aspects of the alternating map β . Theorem 6.1 shows that the P_m -hierarchy does not determine N_1 , even on directly indecomposable groups, uniformly for all odd primes p .

There are, however, genuine refinement relations among some of the later profiles. The aggregated flag profile of level r refines N_r , since it records the ranks of all coordinate subtuples, not only the rank of the full r -tuple. The canonical support-image profile CSI_r further refines the aggregated flag profile: it records the $\text{GL}(W)$ -orbit of the actual image subspace

$$I_\beta(\mathbf{v}) = \text{im } A_{\mathbf{v}} \leq W^r,$$

rather than only the dimensions of its coordinate projections.

The Centralizer Hom Transform is of a different nature. It refines the centralizer-size moments by replacing

$$|C_G(\mathbf{x})|^t$$

with

$$|\text{Hom}(K, C_G(\mathbf{x}))|.$$

For finite test groups K , the full transform recovers the isomorphism types of common centralizers with multiplicity. We do not claim that CHT is comparable, in general, with CSI. Rather, it records internal group structure of centralizers, while CSI records codomain-localized linear data in class-two exponent- p groups.

The resulting picture is therefore a ladder of increasingly structured tests in several directions:

$$\begin{aligned} &\text{isotropic-subspace data } (P_m), \\ &\text{rank-support data } (N_r, \Sigma_{r,t}), \\ &\text{incidence data } (\text{Flag}_r), \\ &\text{codomain-localized support-image data } (\text{CSI}_r), \\ &\text{centralizer-type data } (\text{CHT}). \end{aligned}$$

The main point of this paper is that the first two layers are already genuinely distinct, and that the rank-support side admits both a rank-metric-code interpretation and effective codomain-localized refinements.

The results of the paper establish three conclusions.

First, higher commuting probabilities and star profiles encode genuinely different information in class-two exponent- p groups. The former are governed by isotropic-subspace counts, while the latter are governed by rank-support data.

Second, the star rank profiles are not ad hoc invariants: they are precisely the higher rank-support enumerators of the contraction code

$$C_\beta = \{A_v : v \in V\}.$$

Third, codomain-localized support-image profiles provide a compact refinement of rank and flag data, computable in fixed-parameter time with respect to the derived rank $\dim W$ and the tuple level r .

13 Further directions

13.1 Computational benchmarks

The 5-generator class-two exponent- p databases studied by O’Brien and Stanojkovski provide natural testing grounds for N_r , flag profiles, and CSI. A separate computational paper will treat these benchmarks with reproducible scripts and certificates.

13.2 Uniform classification in small derived rank

The rank-support and CSI profiles suggest a possible classification program for alternating maps

$$\beta : V \times V \rightarrow W$$

with small $\dim W$. The main problem is to understand when codomain-localized profiles determine β up to equivalence.

13.3 CSI and canonization

CSI is an invariant, not a canonization procedure in general. However, if the CSI structure on W^r has small stabilizer in $\mathrm{PGL}(W)$, it can be used as a pre-canonization tool. This leads to the study of residual stabilizers of support-image profiles.

13.4 Centralizer structure beyond class two

The CHT definition applies to all finite groups. It may be useful for studying groups with identical commuting graphs or identical centralizer-size distributions but different centralizer isomorphism types.

References

- [1] W. H. Gustafson, What is the probability that two group elements commute?, *American Mathematical Monthly* 80 (1973), 1031–1034.
- [2] S. Eberhard, Commuting probabilities of finite groups, *Bulletin of the London Mathematical Society* 47 (2015), 796–808.
- [3] V. E. Levit and R. Shwartz, Higher Commutativity in Finite Groups: Exact Asymptotics and Finite Spectrum, arXiv:2605.02071.
- [4] E. Gorla, R. Jurrius, H. H. Lopez, and A. Ravagnani, Rank-Metric Codes and q -Polymatroids, *Journal of Algebraic Combinatorics* 52 (2020), 1–19.
- [5] X. Sun, Faster Isomorphism for p -Groups of Class 2 and Exponent p , arXiv:2303.15412.
- [6] E. A. O’Brien and M. Stanojkovski, Geometric invariants for p -groups of class 2 and exponent p , arXiv:2411.19555.

- [7] P. Lescot, Isoclinism classes and commutativity degrees of finite groups, *Journal of Algebra* 177 (1995), 847–869.
- [8] L. Lovasz, Operations with structures, *Acta Mathematica Academiae Scientiarum Hungaricae* 18 (1967), 321–328.
- [9] A. Atserias, P. G. Kolaitis, and W.-L. Wu, On the Expressive Power of Homomorphism Counts, LICS 2021; arXiv:2101.12733.
- [10] J. B. Wilson, The Threshold for Subgroup Profiles to Agree is Logarithmic, *Theory of Computing* 15 (2019), Article 19, 1–25.