Dimension and identifiability of three-sunlet models

Caius Folkerts

Abstract. The three-sunlet model of a finite abelian group of order $m \geq 5$ has dimension $5m - 4$ in every characteristic. We prove the formula using the Latin structure of its parameterization, which also identifies the model with a torus times a toric secant. Over $\mathbb{C}$, a general point of that secant has a unique unordered decomposition; the remaining parameter ambiguity is given exactly by character actions. For a normalized Latin system of order at least five whose loop abelianization has order c, this gives generic degree $2c^2$, an explicit description of the fibers, and rationality over $\mathbb{Q}$. For cyclic groups of order $m \geq 6$, a general model point is determined by $8m - 12$ coordinates. The same dimension method classifies the drops caused by automorphism-orbit identifications and completes the dimension formula for binary level-one networks in these orders.

1. Introduction

The three-sunlet is the smallest phylogenetic network with a cycle, and its Fourier parameterization already has two monomial terms in each coordinate. Its dimension is the missing case in the formula of Gross, Krone and Martin for triangle-free level-one networks [1], who conjectured that a group of order $m \geq 5$ gives dimension $5m - 4$. Cox, Gross and Martin proved this for groups of odd order [2] by a construction that pairs nonidentity elements with their inverses.

We prove the formula for every finite abelian group of order at least five, in every characteristic, by using the Latin structure of the permutations in the parameterization. This makes the group law largely dispensable and brings models associated with loops into the same argument. The Latin structure also determines the generic fibers; after continuous redundancies are removed, the remaining ambiguity consists of two character actions and the exchange of two toric summands. This gives an exact degree formula and proves rationality of the model.

Let G be a finite abelian group of order m, written additively. The Fourier parameterization of its three-sunlet model is

$$q_{a,b,c} = w_{1,a}w_{2,b}w_{3,c}\big(w_{4,a+b}w_{5,a} + w_{4,b}w_{6,a}\big), \qquad a + b + c = 0.$$

Write V_G for its image closure. Absorbing the first family into the last two and replacing c by $-a-b$ gives the equivalent map

$$Q_{a,b} = s_b t_{a+b}\big(u_a h_{a+b} + v_a h_b\big).$$

Every choice of u, v is obtained by setting $w_{1,a} = 1$, so these maps have the same image.

The argument uses the permutations $b \mapsto a + b$ through their Latin property. A *normalized Latin system* on a set Ω of size m is a family $\Pi = (\pi_0, \dots, \pi_{m-1})$ with $\pi_0 = \mathrm{id}$ such that the values $\pi_a(b)$ are distinct for every fixed b. Its nonidentity rows are derangements. Fix $o \in \Omega$ and label each color a by $\pi_a(o)$. Then $a * b = \pi_a(b)$ is a loop with identity o. We retain 0 for the identity color in parameter subscripts; under this identification it is the element o. The corresponding map is

$$Q_{a,b} = s_b t_{\pi_a(b)}\big(u_a h_{\pi_a(b)} + v_a h_b\big). \tag{1}$$

Theorem 1.1. *For every normalized Latin system of order $m \geq 5$, the image closure of* (1) *has dimension $5m - 4$ over every algebraically closed field. In particular,* $\dim V_G = 5m - 4$ *for every finite abelian group of order at least five.*

For a group-based model, Theorem 4.2 shows that the normalized parameter map has generic degree $2m^2$ over $\mathbb{C}$. More generally, the degree is $2c^2$ when the associated loop has abelianization of order c. The theorem describes every component of the general fiber and its function-field symmetries. For cyclic groups of order $m \geq 6$, Theorem 6.1 recovers a general model point from $8m-12$ coordinates, a linear number out of the m^2 available coordinates.

Identifying parameters along automorphism orbits introduces further dimension drops, which the same initial-monomial method classifies exactly (Section 7). Together with the cut-edge formula, these dimensions give the dimension of every binary level-one model for groups of order at least five, including networks with triangles.

The proof of Theorem 1.1 starts with two transitive families of Latin permutations. A finite-difference identity with one omitted equation computes the rank of a suitable initial monomial map. Exponent rank, rather than Jacobian rank over the ground field, makes the argument independent of the characteristic.

The upper bound is immediate after the four redundant parameters have been removed. In the identity row only $\gamma = u_0 + v_0$ occurs. The three scalings

$$s \mapsto \alpha s, \qquad t \mapsto \beta t, \qquad h \mapsto \eta h, \qquad (u, v, \gamma) \mapsto (\alpha\beta\eta)^{-1}(u, v, \gamma)$$

set $s_o = t_o = h_o = 1$ on a dense open set. Thus $5m - 4$ parameters suffice. We prove the matching lower bound by selecting initial monomials whose exponent relations are especially simple.

2. **Punctured differences and initial monomials**

The initial terms will agree with one monomial family except at one edge of each permutation. Their exponent relations therefore say that a function has constant finite difference along every edge but one. Closing a cycle would normally force that constant to vanish. The next lemma shows that two distinct permutations already recover the missing constraint.

Lemma 2.1 (One omitted edge)**.** *Let $\mathcal{P}$ contain at least two distinct nonidentity permutations of a finite set Ω. If $F : \Omega \to \mathbb{Q}$ satisfies*

$$F(\pi x) - F(x) = c_\pi \qquad (x \neq \pi^{-1}o,\ \pi \in \mathcal{P}),$$

then every c_π is zero and $F \circ \pi = F$ for every π. The conclusion also holds when the omitted argument is o for all permutations.

Proof. For a permutation with more than one cycle, summation on a cycle avoiding o gives $c_\pi = 0$; summation on the remaining cycle fills the missing equation. If a full cycle has nonzero constant, its values are

$$F(\pi^i o) = F(o) + ic_\pi, \qquad 0 \leq i < m,$$

since the only omitted edge returns to o. They are distinct. A second permutation with zero constant would then be the identity. Otherwise it too is a full cycle, and summing all values of F gives

$$\sum_{x\in\Omega} F(x) = mF(o) + \frac{m(m-1)}{2}c_\pi = mF(o) + \frac{m(m-1)}{2}c_\sigma.$$

Hence $c_\sigma = c_\pi$. The displayed ordering of the distinct values now forces the two cycles to coincide, a contradiction. Inverting the permutations proves the outgoing version. □

Lemma 2.2. *The nonidentity colors of a normalized Latin system of order $m \geq 5$ can be partitioned into two transitive families P, N, each of size at least two.*

Proof. A family of c colors has every orbit of size at least $c+1$; at a point x, the points $x, \pi_a(x)$ are distinct. For odd $m = 2r + 1$, any partition into two r-sets therefore

works. For $m = 2r$, start with $|P| = r - 1$ and $|N| = r$. The latter family is transitive. If P is not, it has two r-element orbits, and every color of N interchanges them. Removing one color from P leaves a transitive action on each orbit, since every new orbit there has size at least $r - 1 > r/2$. Add one color of N. The new P is transitive, and its r-element complement remains so. □

The rank calculation also applies when only some permutation colors are present.

Theorem 2.3. *Let $\pi_0 = \mathrm{id}$ and let $\pi_1, \dots, \pi_c$ be distinct derangements of a set of size m. Suppose their colors split into transitive families P, N of size at least two. The map (1), using these $c + 1$ colors, has image dimension $3m + 2c - 2$ in every characteristic.*

Proof. Selecting the u-term on P and the v-term on N will split the exponent relations into two finite-difference systems. Raising the weight of a single h-variable reverses exactly one selection in each color; those exceptional rows will recover the coefficients absent from the ordinary rows.

Assign weight zero to s, t, weight 2 to h_o and zero to the other h's. Give (u_a, v_a) weights $(0, 1)$ on P and $(1, 0)$ on $N \cup \{0\}$. The unique terms of least weight are

	u-term	v-term
$a = 0$	$\varnothing$	all b
$a \in P$	$b \neq \pi_a^{-1} o$	$b = \pi_a^{-1} o$
$a \in N$	$b = o$	$b \neq o$.

Let (S, T, H, U, V) be a rational relation among the columns of their exponent matrix. The identity row gives $T_b = -S_b - H_b - V_0$. Put $K = S + H$. The ordinary rows on P, N become

$$\begin{aligned} S_{\pi_a b} - S_b &= U_a - V_0 \qquad (b \neq \pi_a^{-1} o), \\ K_{\pi_a b} - K_b &= V_a - V_0 \qquad (b \neq o), \end{aligned}$$

respectively. Lemma 2.1 and transitivity make both S and K constant. The exceptional rows determine the remaining coefficients. Precisely four scalars remain

$$\begin{aligned} &S_b = s, \quad H_b = k - s, && T_b = -k - v, \\ &V_a = v, \quad U_a = v \quad (a \neq 0), && U_0 = u. \end{aligned} \tag{2}$$

The unused coefficient U_0 accounts for one scalar; the other three are the scaling directions already present in the model. These relations suffice, so the exponent rank is $3m + 2c - 2$.

If polynomials have unique initial monomials with rationally independent exponents, they are algebraically independent over every field. Indeed, in a polynomial relation their least-weight terms have distinct exponents and nonzero coefficients, and cannot cancel. Apply this to an independent set of rows of the exponent matrix. The upper bound comes from combining the identity coefficients and the three scalings, exactly as above. □

Lemma 2.2 now proves Theorem 1.1. The use of rational exponent rank is essential in positive characteristic. Appendix C computes the integral row lattice for group translations; its torsion explains why the selected monomial map can have smaller Jacobian rank.

For later use, the calculation gives an exact formula even without transitivity. If P, N have c_P, c_N orbits and both have at least two colors, the same weights have exponent rank

$$3m + 2c - c_P - c_N.$$

Indeed S is arbitrary on the P-orbits, K on the N-orbits, and U_0, V_0 are free; all other coefficients are then determined by the exceptional rows. For all $m - 1$ Latin colors this is $5m - c_P - c_N - 2$.

3. **Diagonal normalization and toric secants**

The dimension calculation removes continuous redundancies, but it does not determine how many parameter choices remain. The diagonal coordinates provide the useful normalization; dividing each column by its diagonal entry removes the common t-factor and converts both summands into the same ratio form. The remaining map is thus the sum of two points of a single toric cone.

Write the coordinates as $Q_{b,d}$, where $d = \pi_a(b)$ and the color a is determined by (b, d). On the dense open set where the diagonal factors are nonzero, put

$$\begin{aligned} z_d &= \gamma s_d t_d h_d, \qquad k_b = s_b, \qquad l_b = s_b h_b, \\ U_a &= u_a/\gamma, \qquad V_a = v_a/\gamma \quad (a \neq 0). \end{aligned}$$

Rescale k, l to make $k_o = l_o = 1$. Then

$$\begin{aligned} Q_{b,b} &= z_b, \\ Q_{b,d}/z_d &= U_a k_b/k_d + V_a l_b/l_d, \qquad d = \pi_a(b) \neq b. \end{aligned} \tag{3}$$

Conversely take $s = k$, $h = l/k$, $t = z/l$, $u_0 = 1$, $v_0 = 0$ and the indicated remaining coefficients. This realizes every normalized parameter choice.

With $k_o = 1$, define the toric cone $X_\Pi = \overline{\{F(U,k)\}}$, where

$$F(U,k)_{a,b} = U_a k_b / k_{\pi_a(b)} \qquad (a \neq 0).$$

For a transitive family of c colors it has dimension $c + m - 1$; an exponent relation $U_a + K_b - K_{\pi_a(b)} = 0$ gives $U_a = 0$ by cycle summation and then makes K constant. The normalization identifies a dense open part of the model with $\mathbb{G}_m^m \times \overline{X_\Pi + X_\Pi}$. For a full normalized Latin system of order at least five, this gives, in every characteristic,

$$\dim \overline{X_\Pi + X_\Pi} = 4m - 4, \qquad \dim \sigma_2(\mathbb{P}X_\Pi) = 4m - 5.$$

For group translations, X_Π consists of the off-diagonal parts of diagonally conjugated group-circulant matrices.

4. **Tangential contact**

The dimension calculation gives the expected secant dimension. We now ask whether a general point determines its two toric summands. Work over $\mathbb{C}$, and write L for the loop associated with Π.

Definition 4.1. The *abelianization* L_{ab} of this loop is the abelian group with generators $[b]$, $b \in \Omega$, and relations

$$[o] = 0, \qquad [a * b] = [a] + [b] \quad (a, b \in \Omega).$$

Write $c_\Pi = |L_{\mathrm{ab}}|$.

This group is finite and its order divides m. Indeed, the image of the finite set Ω in L_{ab} is closed under addition and contains zero, so is a subgroup and hence all of L_{ab}. Each loop translation bijects the fibers of this map, so all fibers have the same cardinality. For a group this is the usual abelianization.

Theorem 4.2. *For every normalized Latin system of order $m \geq 5$, a general point of $\overline{X_\Pi + X_\Pi}$ has a unique unordered decomposition into two points of X_Π. Consequently the normalized map* (3) *has generic degree $2c_\Pi^2$, with function-field Galois group*

$$\big(\mathrm{Hom}(L_{\mathrm{ab}}, \mathbb{C}^*)\big)^2 \rtimes S_2.$$

The general complex fiber of the five-family Latin model is a disjoint union of $2c_\Pi^2$ copies of $(\mathbb{C}^*)^3 \times \mathbb{C}$. Both its image closure and $\overline{X_\Pi + X_\Pi}$ are rational over $\mathbb{Q}$.

The proof uses a pair of points supported on disjoint color families. A finite-difference identity controls their tangent spaces; removing a color in turn controls the contact locus. The degree and fiber statements then follow in Section 5.

4.1. **A two-ratio rigidity lemma**

For dimension, a single omitted equation was enough. For tangent spaces, the second summand contributes a second coefficient in each color. We choose its torus parameters so that their ratio is 1 except on the two edges incident with o. Cycle sums then isolate these two coefficients, giving the analogue of the one-hole lemma.

Put $l_o = 1$ and $l_x = \lambda$ for $x \neq o$, where $\lambda = 2$ or $\lambda = \frac{1}{2}$. For a permutation π, write

$$\Delta_\pi F(x) = F(x) - F(\pi x), \qquad r_\pi(x) = l_x / l_{\pi x}.$$

Lemma 4.3. *Let $\mathcal{P}$ contain at least two distinct derangements of Ω. If $F : \Omega \to \mathbb{C}$ satisfies*

$$\Delta_\pi F = c_\pi + d_\pi r_\pi \qquad (\pi \in \mathcal{P}),$$

then all c_π, d_π vanish, and F is invariant under every $\pi \in \mathcal{P}$.

Proof. If π has more than one cycle, summing on a cycle avoiding o gives $c_\pi + d_\pi = 0$. The cycle containing o has the two exceptional ratios λ^{-1} and λ, so its sum gives $d_\pi(\lambda + \lambda^{-1} - 2) = 0$. Thus $d_\pi = c_\pi = 0$.

Suppose instead that π is a single m-cycle. Subtract a constant from F so that $F(o) = 0$, and put $A = (\lambda - 1)^2/\lambda$. The cycle sum gives $c_\pi + d_\pi = -d_\pi A/m$, whence

$$F(\pi^i o) = d_\pi \left(\frac{\lambda - 1}{\lambda} + \frac{iA}{m} \right), \qquad 1 \leq i < m.$$

For $\lambda = 2$ these values are $d_\pi(m + i)/(2m)$; for $\lambda = 1/2$ they are $d_\pi(-1 + i/(2m))$. If $d_\pi \neq 0$, all m values of F are distinct. A second permutation $\sigma \in \mathcal{P}$ must then also be an m-cycle with $d_\sigma \neq 0$, since invariance under a nonidentity permutation would contradict this injectivity. Summing the displayed values yields

$$\sum_{x \in \Omega} F(x) = d_\pi \frac{(m-1)(\lambda - \lambda^{-1})}{2} = d_\sigma \frac{(m-1)(\lambda - \lambda^{-1})}{2}.$$

Consequently $d_\pi = d_\sigma$. The displayed formula and injectivity now imply $\pi^i o = \sigma^i o$ for every i, so $\pi = \sigma$, a contradiction. Thus every d_π vanishes; cycle sums force every c_π to vanish as well. □

4.2. **A sparse tangent pair**

For a transitive family of c nonidentity colors, the normalized source of F is $\mathcal{M} = \mathbb{A}^c \times (\mathbb{C}^*)^{m-1}$. Choose disjoint transitive families A, B, each of size at least two,

and take $x = F(\mathbf{1}_A, 1)$, $y = F(\mathbf{1}_B, l)$, with $l_o = 1$, $l_b = 2$ for $b \neq o$. Their combined differential is injective. A kernel vector (p, K, q, M) with $K_o = M_o = 0$ satisfies

$$p_a + \mathbf{1}_A(a)\Delta_a K + r_a\big(q_a + \mathbf{1}_B(a)\Delta_a M\big) = 0.$$

On A, Lemma 4.3 gives $K = 0$; on B, division by r_a gives $M = 0$. The distinct values of r_a then give $p_a = q_a = 0$. This yields, in the original six-family map, a point of rank $5m - 4$ with entries in $\{0, 1, 2, \frac{1}{2}\}$.

The distinct coordinate monomials make the projectivization of X_Π nondegenerate, and the injective tangent pair gives its expected secant dimension. We use the tangential-contact criterion of [3, Proposition 4.6]; for an irreducible nondegenerate projective variety of dimension at least two which is not two-defective, failure of generic two-identifiability forces the contact locus at a general pair to contain a curve through each member.

It suffices to test contact on the parameter space. The map F is generically finite and hence generically etale. If $T = \operatorname{im} DF_x + \operatorname{im} DF_y$, a nonzero minor giving rank $\dim \mathcal{M} - 1$ for

$$v \longmapsto \big(w \longmapsto D^2F_x(v, w) \bmod T\big) \tag{4}$$

persists at a general pair. There the image points are smooth and the contact tangent is radial. Thus no smoothness assertion at the sparse image points is required. Logarithmic torus tangent fields may be used; changing the tangent frame adds a tangent vector, which disappears modulo T.

Put $E_a = \operatorname{span}\{1, r_a\}$, $s_a = r_a^{-1}$, and let D consist of the unused colors. The tangent span separates as

$$T_A = \bigoplus_{a \in A} E_a + \operatorname{im} \Delta_A, \qquad T_B = \bigoplus_{a \in B} E_a + r \operatorname{im} \Delta_B, \qquad T_D = \bigoplus_{a \in D} E_a.$$

Differentiating the coefficient and torus directions gives two conditions for (p, K) to lie in the kernel of (4). The *coefficient condition* requires, for every color a, that the vector supported in that color with value $\Delta_a K$ belong to T. The *torus condition* requires, for every $H : \Omega \to \mathbb{C}$, that the vector with color components

$$p_a \Delta_a H + \mathbf{1}_A(a)(\Delta_a K)(\Delta_a H)$$

belong to T.

Lemma 4.4 (Both endpoints omitted)**.** *Let $\mathcal{P}$ contain at least three distinct derangements of Ω. Suppose $F : \Omega \setminus \{o\} \to \mathbb{C}$ satisfies*

$$F(\pi x) - F(x) = c_\pi \quad \textit{whenever } x \neq o \textit{ and } \pi x \neq o.$$

Then all c_π vanish. If the underlying graph of these ordinary edges is connected, F is constant.

Proof. A permutation with more than one cycle has a cycle avoiding o, so its constant is zero. A full cycle with nonzero constant orders the $m-1$ values of F as an injective arithmetic progression. For any other derangement a zero constant would contradict injectivity on an ordinary edge; thus it too would be a full cycle with nonzero constant. The same finite arithmetic progression admits only the forward and backward constant-step orderings. Indeed, divide by the first step and translate its values to $\{0, 1, \dots, m-2\}$; the other step is an integer, and comparison of the minimum and maximum makes it 1 or -1. The other permutation is therefore the first permutation or its inverse. Three distinct permutations are impossible. All constants vanish, and connectedness gives the final assertion. □

4.3. **Redundant color families**

The tangent calculation must remain rigid when one color is removed; the coefficient condition isolates that color, leaving all the others to constrain the same auxiliary function. We therefore choose B with redundant transitivity. Once the coefficient condition has forced $K = 0$, the torus condition first makes the coefficient variation constant on each support and then kills the variation on B. The only surviving direction is rescaling the first point.

Proposition 4.5. *Suppose the colors are a subset of a normalized Latin system, and A, B are disjoint transitive families with $|A| \geq 2$. If each $B \setminus \{b\}$ is transitive with at least two colors, and the graph of B is connected after deleting o, then the contact kernel at the sparse pair is radial.*

Proof. The coefficient condition in color $b \in B$ supplies M with $\Delta_c M \in \operatorname{span}\{1, s_c\}$ for $c \in B \setminus \{b\}$. Lemma 4.3 makes M constant, so $\Delta_b K \in E_b$. Applying the lemma again on B gives $K = 0$.

We first show that p is constant separately on A and B. Identify $Q = \mathbb{C}^\Omega/\mathbb{C}\mathbf{1}$ with functions vanishing at o, put $e = \mathbf{1}_{\Omega\setminus\{o\}}$, and write

$$W_a^r = \{H \in Q : \Delta_a H \in \operatorname{span}\{1, r_a\}\},$$
$$W_a^s = \{H \in Q : \Delta_a H \in \operatorname{span}\{1, s_a\}\}.$$

The proof of Lemma 4.3 shows that these spaces have dimension one for a full cycle and dimension one less than the number of cycles otherwise. Since every cycle of a derangement has length at least two, for distinct a, b and $m \geq 5$,

$$\dim(W_a^r + W_b^r),\ \dim(W_a^s + W_b^s) < m - 1. \tag{5}$$

The A-color equations of the torus condition give $J - p_a H \in W_a^r$ for one common J. If $p_a \neq p_b$, subtracting these relations would put every $H \in Q$ in $W_a^r + W_b^r$, contrary to (5). Thus $p_a = p_b$.

For B, division by r_a weights a difference by s_a. Since $s_a = 1$ away from the two edges at o, this weighted difference can be corrected, up to an additive constant, by changing its value at o. In the normalization $H(o) = 0$, that correction takes the form

$$L_a(H) = H(\pi_a o) + H(\pi_a^{-1} o), \qquad R_a H = H - L_a(H) e. \tag{6}$$

Checking the two exceptional endpoints gives

$$s_a \Delta_a H - \Delta_a(R_a H) = -\big(2H(\pi_a o) + H(\pi_a^{-1} o)\big)(s_a - 1).$$

Consequently the torus condition, divided by r_a, says $M - p_a R_a H \in W_a^s$ for one common M. If $p_a \neq p_b$, the operator $p_a R_a - p_b R_b$ is invertible; write it as $\lambda I - e\mu$, where $\lambda = p_a - p_b$ and $\mu(e) = 2\lambda$. It acts by λ on $\ker \mu$ and by $-\lambda$ on e. Its image cannot lie in $W_a^s + W_b^s$, by (5). Thus $p_a = p_b$ on B too.

Write this common value as p. The B-equations are

$$p s_a \Delta_a H = \Delta_a M + c_a + d_a s_a. \tag{7}$$

On ordinary edges, Lemma 4.4 applied to $M - pH$ gives $c_a + d_a = 0$ and makes that difference constant off o. Normalize it to zero there and put $M(o) = pH(o) + t$. The endpoint equations in (7) give

$$t = p\big(H(\pi_a o) + H(\pi_a^{-1} o) - 2H(o)\big) \quad (a \in B).$$

At most two Latin colors have the same unordered neighbor multiset at o, because their outgoing neighbors are distinct. As $|B| \geq 3$, varying H forces $p = 0$. Finally, for $a \in D$, the torus condition would put $p_a \operatorname{im} \Delta_a$ in E_a; since $\operatorname{rank} \Delta_a \geq \lceil m/2 \rceil \geq 3$, it gives $p_a = 0$. The remaining direction is exactly radial on A. □

4.4. **Support choices at every order at least six**

We now choose supports satisfying the preceding proposition. A small transitive family can serve as A, leaving enough colors in B that deletion of any one color preserves transitivity. The Latin property supplies such a small family by repeatedly merging orbits. Only the smallest orders require separate choices.

Every normalized Latin system has a transitive set of at most $\lfloor \log_2 m \rfloor$ colors. To see this, take the current orbit partition with $k > 1$ blocks. At most $|C| - 1$ nonidentity colors preserve a block C, by Latin distinctness at a point of C. The total number of block–color incidences is at most $m - k < m - 1$. Some color therefore

preserves no block. Adjoining it merges every old orbit with at least one other, so there are at most $\lfloor k/2 \rfloor$ new orbits. Iteration proves the bound. Pad the transitive set to size at least two.

If $m = 7$ or $m \geq 9$, choose such a set A and let B be its nonidentity complement. Then $|B| > m/2$. Every orbit of $B \setminus \{b\}$ has at least $|B|$ elements, so that set is transitive. After deleting o, each graph component for B has at least $|B|$ vertices; each remaining vertex has at least $|B| - 1$ distinct outgoing neighbors. Thus the deleted-base graph is connected. All the hypotheses of the preceding calculation hold.

For $m = 6$, a nontransitive pair of colors has two orbits of size three. On each orbit the two permutations are the two distinct three-cycles, and hence are inverses. Thus the nontransitive pairs form a matching on the five nonidentity colors. Split these colors into a two-set A and a three-set B, separating the ends of each edge of the matching. Then A and every pair in B are transitive. Each component of the deleted-base B graph has at least three vertices among the five remaining vertices, so that graph is connected.

For $m = 8$, a transitive pair can be used as A, leaving $|B| = 5$. Suppose no such pair exists. No color is a full cycle. A color with exactly two cycles would force every other color to preserve both cycles, since joining them would give a transitive pair; this contradicts the Latin bound on the number of colors preserving the smaller cycle. A color of cycle type $4 + 2 + 2$ similarly forces every color to preserve its four-cycle; a pair cannot have a two-element orbit, because the first color already uses the unique nontrivial permutation on that pair and Latin distinctness forbids a second color from preserving it. A color of type $3 + 3 + 2$ forces each other color to preserve one of its three-cycles, but at most one other color can preserve each of them. Both possibilities are impossible.

All colors are therefore fixed-point-free involutions. The union of the two matchings for any pair consists of two four-cycles, since an eight-cycle would be transitive. The involutions consequently commute. A faithful transitive abelian permutation group is regular, so this system is the regular C_2^3 system. The following direct calculation handles this system.

The group C_2^3. Label elements by binary integers with addition given by exclusive-or. Take $A = \{1, 2, 4\}$, $B = \{3, 6, 7\}$ and $D = \{5\}$. Every nonzero translation has several cycles, so the single-permutation conclusion of Lemma 4.3 applies. The coefficient condition in color 5 makes K invariant under 5. The coefficient condition in color 1 supplies H with

$$\Delta_2 H \in E_2, \quad \Delta_4 H \in E_4, \quad \Delta_1(H - K) \in E_1.$$

All three differences vanish. Thus $\Delta_1 K$ is invariant under 4. But invariance of K under $5 = 1+4$ also makes $\Delta_1 K(b+4) = -\Delta_1 K(b)$; hence $\Delta_1 K = 0$. The normalized function K is now zero on $\{0,1,4,5\}$ and has a value t on $\{2,3,6,7\}$.

The coefficient condition in color 3 supplies M invariant under 6 and 7 with $\Delta_3 K - r_3 \Delta_3 M \in E_3$. After a constant is subtracted, M is zero on $\{0,1,6,7\}$ and has a value v on $\{2,3,4,5\}$. Outside $\{0,3\}$ the ratio r_3 is one, so the membership equation has constant right side. Evaluation at the ordinary pair $1,2 = 1+3$ shows that constant is zero. Evaluation at 1 and 4 then gives $t = v$ and $-t = v$, respectively. Thus $K = 0$. The graph of B remains connected after deleting zero. For example, its ordinary edges contain the path $3,4,7,1,6,5,2$. The preceding radial calculation now applies.

This proves generic two-identifiability for every normalized Latin system of order at least six. The remaining five-point calculation is given in Appendix A.

5. **Degree, fibers, and rationality**

The contact argument proves that the two toric points in Theorem 4.2 are uniquely determined, but each point can still have several parameter representations. Computing the kernel of the one-component torus map identifies this remaining ambiguity with a character group. Restoring the three scalings and the split identity coefficient then recovers the continuous part of the original fiber.

Two normalized preimages of a general point of X_Π differ by ratios $\rho_b = k'_b/k_b$, $\rho_o = 1$. Comparison at $b = o$ first gives $U'_a = \rho_a U_a$, and the other coordinates then give

$$\rho_{a*b} = \rho_a \rho_b.$$

Thus the ambiguity is exactly a character of L_{ab}, of which there are c_Π. A general two-component sum has both components in the toric open set; sums involving its boundary have dimension at most $2\dim X_\Pi - 1$, strictly below the secant dimension. The unique unordered pair therefore has $2c_\Pi^2$ ordered normalized preimages. The two character actions and interchange are generically free and give that many automorphisms of the function field. Their invariant field is the field of the image, proving the Galois assertion.

The continuous fibers can be read directly from the normalization. For each normalized preimage (z,k,l,U,V) choose $S,H,\Gamma \in \mathbb{C}^*$ and $T \in \mathbb{C}$, and put

$$s_b = Sk_b, \qquad h_b = Hl_b/k_b, \qquad t_b = z_b/(\Gamma S H l_b),$$
$$u_a = \Gamma U_a \ (a \neq 0), \quad v_a = \Gamma V_a \ (a \neq 0), \quad u_0 = T, \quad v_0 = \Gamma - T.$$

These are all the parameters in that component. Over a general output, nonzero diagonal coordinates force Γ, s_b, t_b, h_b to be nonzero, so the normalization is defined and recovers a unique one of the finitely many normalized preimages. This proves that the five-family fiber is a disjoint union of the claimed tori times an affine line.

For the original group model choose, in addition, arbitrary $w_{1,a} \in \mathbb{C}^*$, then set $w_{5,a} = u_a/w_{1,a}$, $w_{6,a} = v_a/w_{1,a}$, $w_2 = s$, $w_{3,c} = t_{-c}$ and $w_4 = h$. Every $w_{1,a}$ is necessarily nonzero at a general output, since all coordinates of its color would otherwise vanish. Hence the original six-family fiber is a disjoint union of $2m^2$ copies of

$$(\mathbb{C}^*)^{m+3} \times \mathbb{C}.$$

For a nonabelian translation model with the same redundant factorization, replace m^2 by $|G_{\mathrm{ab}}|^2$. These descriptions also give the geometric generic fibers; they do not require a splitting over the unextended function field.

Finally X_Π is a split toric variety over $\mathbb{Q}$ and is therefore rational over $\mathbb{Q}$. Unique unordered decomposition makes $\mathrm{Sym}^2 X_\Pi \dashrightarrow \overline{X_\Pi + X_\Pi}$ birational. The map is defined over $\mathbb{Q}$ and has degree one after extension to $\mathbb{C}$, hence already over $\mathbb{Q}$. The symmetric square of a rational d-fold is rational; on two affine coordinate copies use their means u_i and differences v_i. Interchange fixes u_i and negates all v_i, so the invariant field has the algebraically independent generators

$$u_1, \dots, u_d, \quad v_1^2, \quad v_2/v_1, \dots, v_d/v_1.$$

The independent diagonal torus in (3) proves rationality of the model as well. This completes Theorem 4.2.

Corollary 5.1. *For a nonassociative loop of prime order $m \geq 5$, the normalized map has generic degree two. Its only general ambiguity is interchange of the two components.*

Proof. The order of L_{ab} divides m. If it were m, the loop map onto its abelianization would be bijective and would make the loop a cyclic group. Thus $c_\Pi = 1$. □

6. **Sparse coordinate recovery**

For cyclic groups, the identifiability calculation also works with just five nonidentity colors. This gives a concrete coordinate-recovery result whose size is linear in the group order.

Theorem 6.1. *Let $G = \mathbb{Z}/m\mathbb{Z}$ with $m \geq 6$. A general point of the full three-sunlet model V_G is determined by all coordinates $q_{a,b,-a-b}$ with $a \in \{0, 1, 2, -1, 3, -2\}$, together with the two coordinates $b = 0$ and $b = -a$ for every remaining nonzero color a. These are $8m - 12$ coordinates in total, and the corresponding coordinate projection is birational onto its image over $\mathbb{Q}$.*

Proof. Five complete nonidentity colors determine the shared toric parameters of the two summands. Once those are known, each remaining color has only two unknown coefficients, which two further coordinates recover by a linear system.

Put $A = \{1, 2\}$ and $B = \{-1, 3, -2\}$. These five colors are distinct for $m \geq 6$. The set A generates the cyclic group. Each pair in B also generates; two pairs contain -1, and the remaining pair $\{3, -2\}$ generates because $\gcd(3, 2) = 1$. The underlying graph of B after deleting zero is connected, since the -1 cycle alone leaves a spanning path. Also B contains colors with different unordered neighbor pairs at zero.

Proposition 4.5 therefore applies to the model using only these five nonidentity colors. Its proof uses exactly the properties just verified; it does not require that every Latin color be present. Theorem 2.3 gives dimension $3m+8$, and the normalized toric component has dimension $m + 4$. The two-component tangent span is independent and the contact kernel is radial. Thus the observed five-color normalized matrix has a unique unordered toric decomposition generically.

Its one-component torus ambiguity is exactly a character of the cyclic group, since the color 1 equations already force $\rho_{b+1} = \rho_1\rho_b$ and $\rho_1^m = 1$. Consequently the six fully observed colors recover z, k, l and their five pairs of coefficients up to the two character actions and interchange of components.

For an unobserved color a, the two further coordinates, after division by their known diagonal coordinates, give

$$\begin{pmatrix} k_0/k_a & l_0/l_a \\ k_{-a}/k_0 & l_{-a}/l_0 \end{pmatrix} \begin{pmatrix} U_a \\ V_a \end{pmatrix} = \begin{pmatrix} Q_{0,a}/z_a \\ Q_{-a,0}/z_0 \end{pmatrix}.$$

With $k_0 = l_0 = 1$, its determinant is

$$\frac{l_{-a}}{k_a} - \frac{k_{-a}}{l_a} = \frac{l_a l_{-a} - k_a k_{-a}}{k_a l_a},$$

which is a nonzero rational function, including when $a = -a$. It is therefore invertible at a general parameter point. This recovers U_a, V_a and hence every coordinate of that color. The reconstructed full matrix is unchanged by the character actions or component swap. The general full model point is uniquely determined.

The number of retained coordinates is $6m + 2((m-1)-5) = 8m-12$. Generic one-to-one-ness gives birationality in characteristic zero. The projection is defined over $\mathbb{Q}$ and has degree one after extension to $\mathbb{C}$, so it is birational already over $\mathbb{Q}$. □

7. **Parameter identifications**

Let $B \leq \mathrm{Aut}(G)$ and identify parameters within every B-orbit, in each of the six families. Write r for the number of orbits, including $\{0\}$, and denote the resulting variety by V_G^B. We ask whether these identifications cause dimension loss beyond the four normalizations already present in the full model.

The initial-monomial method reduces this question to a much smaller one. A relation among the identified parameters is a function on the classes, constrained by differences between them. For an inverse-stable partition, we shall make the entire remaining kernel lie in two copies of a space of dimension at most one. The exceptional models occur exactly when that last relation cannot be removed.

Theorem 7.1. *For a finite abelian group G of order at least five, over any algebraically closed field,*

$$\dim V_G^B = \begin{cases} 5, & r = 2, \\ 10, & r = 3 \text{ in the index-two case}, \\ 5r-4, & \text{otherwise.} \end{cases}$$

Here the index-two case has orbit partition $\{0\}, H\setminus\{0\}, G\setminus H$ with $[G:H]=2$. The same classification holds for any *inverse-stable partition $\{0\}, O_1, \dots, O_{r-1}$, meaning $O_i = -O_i$ for every i. Such a partition need not arise from automorphisms.*

The diagonal normalization in (3) gives r independent diagonal coordinates and the off-diagonal map

$$F_{b,c} = U_{c-b}\frac{k_b}{k_c} + V_{c-b}\frac{l_b}{l_c}, \qquad b \neq c. \tag{8}$$

Here $k_0 = l_0 = 1$, and every parameter is constant on the chosen classes. The normalized parameter count is $4r-4$; thus the total dimension is at most $5r-4$.

Lemma 7.2 (Stable generating supports)**.** *Suppose two disjoint unions A, C of nonzero classes each generate G and each contain at least two elements. Then the dimension is $5r-4$.*

Proof. Extend A, C to a partition P, N of all nonzero elements into unions of classes. Use the weights of Section 2. The unrestricted exponent matrix has exactly the four

kernel directions in (2), all constant on classes. Identifying parameters restricts its kernel to the class-constant subspace, so the identified exponent matrix has rank $5r-4$. The initial-monomial argument in Theorem 2.3 proves the claim. □

7.1. **Inverse-stable classes and a one-dimensional height space**

First assume that each $O_i = -O_i$, and write $s = r-1$. We use an ordering of the classes to select opposite summands on oppositely directed cross-class coordinates. The resulting relations are additive along increasing heights. Small independent choices within a class then let us eliminate any residual additive relation. Order the nonzero classes and give them distinct positive integer heights h_i, with $h_0 = 0$. Assign weights to (8) so that the difference between its two summand weights is

$$\delta_i + h_b - h_c \qquad (c - b \in O_i),$$

where $0 < |\delta_i|$ is smaller than every height gap. Multiplying all weights by a common denominator makes them integral. If $h_b < h_c$, the first summand is uniquely selected; if $h_b > h_c$, the second is selected. If b, c are in the same class, the sign of δ_i makes the choice.

Compute the rational kernel of the selected Laurent exponent matrix. Write K_i, L_i for its coordinates on k, l. The two coordinates with one endpoint zero give

$$U_i = K_i, \qquad V_i = -L_i.$$

The cross-class rows and their reversals therefore say that both K and L belong to the same linear space

$$\mathcal{A}_h = \{x \in \mathbb{Q}^s : x_k = x_j + x_i \text{ if } h_j < h_k \text{ and } (O_k - O_j) \cap O_i \neq \varnothing\}. \tag{9}$$

Let

$$R = \{i : O_i \cap (O_j - O_j) \neq \varnothing \text{ for some } j \neq 0\}.$$

For each $i \in R$, the same-class rows impose either $K_i = 0$ or $L_i = 0$, as chosen by the sign of δ_i.

Lemma 7.3. *The space $\mathcal{A}_h$ has dimension at most one. A nonzero vector, after a scalar rescaling, has values in $\{0, 1, 2\}$ and is nondecreasing in the class order. Its zero set in G, including zero, is a subgroup. Every class at level two belongs to R.*

Proof. A vector with maximum $M > 0$ and minimum $m < 0$ would, by the relation between those two classes, have a coordinate equal to either $M - m$ or $m - M$. Both possibilities contradict its extrema. Thus every vector is nonnegative or nonpositive.

A linear subspace with this property has dimension at most one; two independent nonnegative vectors have a linear combination with both signs.

For a nonnegative vector, extend its values to a function on G, with value zero at zero. Equation (9) makes its values nondecreasing in the class order. Together with inverse stability, it gives the useful identity

$$x_{c-b} = |x_c - x_b| \qquad \text{when } b, c \text{ belong to distinct classes.}$$

For elements b, c with $0 < x_b < x_c$, put $a = c - b$. Then $x_a = x_c - x_b$. If a and $-b$ belonged to different classes, applying the same equation to that pair would give $x_c = |x_a - x_b| < x_c$. Hence they belong to the same class, $x_a = x_b$, and $x_c = 2x_b$. There can be no three distinct positive values. This proves the assertion about levels.

Let Z be the zero set. Every $z \in Z$ translates each positive level into itself. The subgroup generated by Z therefore preserves each positive level and cannot meet one; a stabilizing element in a positive level would translate itself to zero. Thus that subgroup is exactly Z.

Finally, for a level-two element c and a level-one element b, the previous argument puts $c - b$ in the same class as $-b$. Consequently c is a difference of two elements of that class, proving its class belongs to R. □

If $\mathcal{A}_h = 0$, the selected exponent kernel is zero. If it is one-dimensional and at least two R-coordinates of a nonzero vector are positive, choose opposite signs of δ on those two classes. This kills both copies of $\mathcal{A}_h$. We next show that an order avoiding the remaining possibilities always exists outside the stated exceptions.

7.2. **Choosing the class order**

Call an order bad if $\mathcal{A}_h$ contains a nonzero vector with at most one positive coordinate indexed by R. Write Z for its zero subgroup.

Possible obstructions. If the vector has only one positive level, and that level contains two or more classes, differences between distinct such classes belong to Z. Fixing an element of either class shows that their entire union, and then the whole positive level, lies in one Z-coset. The level is $G \setminus Z$, so Z has index two. If the positive level consists of one class D, then $D = G \setminus Z$ and D is last in the order.

Suppose instead that both positive levels occur. By Lemma 7.3, level two is a single class C, and no level-one class belongs to R. For each $c \in C$, the proof of that lemma shows that translation by c preserves every level-one class. Also $C + Z = C$, so $Z \subseteq \langle C \rangle$. Conversely, $\langle C \rangle$ preserves each level-one class and cannot meet it. Hence

$$H = \langle C \rangle = Z \cup C$$

is a subgroup. Differences of level-one elements lie in H; for different classes this is (9), and within one class a level-one difference would give a second positive R-class. Thus $G \setminus H$ is one H-coset, and H has index two. Since H stabilizes every level-one class and acts transitively on that coset, the entire level-one set is a single class D. In this case D is second-last and C is last.

Every bad order therefore ends with an index-two coset block, has a class equal to the complement of a subgroup last, or has such a complement class second-last followed by a single level-two class.

Avoiding the obstructions. There is at most one class D which is the complement of a proper subgroup, since each complement has at least $|G|/2$ elements. If such a class exists and $s \geq 3$, put it first. If its subgroup has index greater than two, $|D| > |G|/2$, so no index-two subgroup can be a union of classes. If its subgroup has index two, it is the unique index-two subgroup that is a union of classes; D cannot lie inside any such subgroup, since it already has half the elements and omits zero. Putting D first therefore excludes all three bad types.

If no complement class exists, only the index-two coset tails remain. An index-two subgroup is the kernel of a homomorphism to $\mathbb{Z}/2\mathbb{Z}$; it is a union of classes precisely when that homomorphism is constant on each class. All such homomorphisms factor through the quotient by

$$K = 2G + \langle x - y : x, y \text{ belong to the same class} \rangle.$$

Thus the relevant index-two subgroups are exactly those containing K. If $K \neq 0$, put any nonzero class contained in K last; this excludes every index-two coset tail. Such a class exists because K is a union of classes; it is the intersection of all index-two subgroups which are unions of classes, or G if there are none. If $K = 0$, the group has exponent two and every class is a singleton, so the unrestricted theorem applies directly.

For $s = 2$ outside the index-two exception, the same conclusion is even simpler. If there is a complement class D, its subgroup has index at least three. Put D first. There is no index-two subgroup that is a union of classes, and the two-level bad case would require D to be an index-two complement. If there is no complement class, a bad tail would have to consist of one class, again a contradiction. Thus a good order exists. This proves the full expected dimension for every nonexceptional inverse-stable partition.

The exceptional dimensions. For $s = 1$, the normalized coordinate functions are precisely $u + v, uK + vL, u/K + v/L$, giving an upper bound of three. The height

space has dimension one and its single coordinate lies in R, so the selected exponent rank is $4-1=3$. Adding the two diagonal coordinates gives five. In the index-two case with $s=2$, put the outside class first and the inside class last. Then (9) is $x_{\mathrm{inside}} = 2x_{\mathrm{outside}}$, and only the inside class belongs to R. The exponent rank is $8-1=7$. The seven distinct normalized coordinate functions are

$$U_I+V_I,\quad U_Ik_I+V_Il_I,\quad U_I/k_I+V_I/l_I,$$
$$U_Ck_C+V_Cl_C,\quad U_C/k_C+V_C/l_C,$$
$$U_Ck_I/k_C+V_Cl_I/l_C,\quad U_Ck_C/k_I+V_Cl_C/l_I,$$

where $I=H\setminus\{0\}$ and $C=G\setminus H$. They give the matching upper bound. Adding the three diagonal coordinates gives ten. These arguments use rational exponent ranks, not Jacobians in the ground field.

7.3. **Arbitrary automorphism orbits**

Negation gives a way to use the inverse-stable calculation even when the original orbits are not inverse-stable. Coarsen the partition by adjoining negation, and split a possible relation into its even and odd parts. The coarser calculation kills the even part. A nonzero odd part will itself produce the two disjoint generating supports needed in Lemma 7.2.

Now let B be arbitrary, and put $B^+=\langle B,-\mathrm{id}\rangle$. Every B^+-orbit is either one B-orbit or two opposite B-orbits. Suppose first that the coarser inverse-stable model is nonexceptional, and take a full-rank initial system as above, with weights constant on coarse orbits. The fine exponent kernel is invariant under negation. Over $\mathbb{Q}$ it splits into even and odd parts. Its even part is the zero kernel of the coarser normalized initial system.

In the odd part, the zero-endpoint rows give $U_a=K_a, V_a=L_a$. For b,c in distinct coarse orbits, the selected row and its reversed row give

$$K(c-b)=K(c)-K(b),\qquad L(c-b)=L(c)-L(b).$$

If this odd kernel is nonzero, one of K,L is a nonzero odd, B-invariant rational function f with this cross-orbit additivity. Let $M>0$ be its maximum. Every negative value must lie in the same coarse orbit D as an element attaining M; otherwise additivity produces an absolute value greater than M. Oddness then puts every nonzero value in D. Thus D splits into two fine classes D^+, D^- with values $M,-M$.

For $z\notin D$ and $d\in D^+$, additivity gives $f(d-z)=M$. Thus $H=G\setminus D$ stabilizes both fine classes under translation. The stabilizer of D^+ meets neither D^+ nor D^-,

since a stabilizing element in D^+ would translate itself to zero, and the negative of an element of D^- is in D^+. Hence the stabilizer is exactly H, and H is a subgroup.

Both D^+ and D^- generate G. Indeed, for $d \in D^+$ and $h \in H$, $h = (d+h) - d$ puts H in $\langle D^+ \rangle$. In G/H, the images of the two classes are negatives of each other and exhaust all nonzero elements, so either generates the quotient. The same argument applies to D^-. Each class has at least two elements; a singleton would force $H = 0$ and $|G| = 3$. Lemma 7.2 now gives the full dimension directly. If the odd kernel is zero instead, the coarse initial system already has full rank on the fine parameters.

It remains to consider the two coarse exceptions. If B^+ is transitive on $G \setminus \{0\}$, a nonreal refinement consists of opposite orbits $O, -O$, each generating G and of size at least two. The support lemma applies.

In the index-two three-orbit exception, transitivity on $H \setminus \{0\}$ makes H elementary abelian of some prime exponent p. If p is odd, the outside coset contains elements of orders 2 and $2p$, contradicting transitivity. If $p = 2$ and an outside element has order four, then $2G$ is a nontrivial characteristic subgroup of H of order at most two. Transitivity forces $|H| = 2$ and $|G| = 4$, excluded. Therefore G has exponent two. Negation is trivial, so $B^+ = B$ and there is no fine refinement. This completes the proof of Theorem 7.1.

Corollary 7.4 (Cyclic orbit models)**.** *For $G = \mathbb{Z}/n\mathbb{Z}$, $n \geq 5$, and any $B \leq \mathrm{Aut}(G)$, the dimension is $5r - 4$ except when n is prime and $B = \mathrm{Aut}(G)$, when it is five.*

Proof. The index-two exception cannot occur for a cyclic group of order at least five. A cyclic group has a single nonzero automorphism orbit exactly in the prime-order case with its full automorphism group. □

8. **Binary level-one networks**

The three-sunlet calculation supplies the last local block needed for a network with disjoint cycles. Cutting its internal cut edges reduces the dimension to a sum of tree and sunlet dimensions; the exceptional orbit types contribute one correction for each triangle.

For this section, a network is a reduced binary semi-directed level-1 network; its root has been suppressed, its leaves have degree one, all other vertices have degree three, and its cycles are vertex-disjoint and have length at least three. Let E be its number of edges, C its number of cycles, C_3 its number of triangles, and n its number of leaves. Edges here are counted *before* any contraction of a reticulation edge.

Let r be the number of B-orbits in G and define $\delta_B = 1$ for either exceptional orbit type in Theorem 7.1, and $\delta_B = 0$ otherwise.

Corollary 8.1. *For every finite abelian G of order at least five and every $B \leq \mathrm{Aut}(G)$, the associated complex level-1 model satisfies*

$$\dim V_{\mathcal{N}}^{(G,B)} = (r-1)(E-C) + 1 - \delta_B C_3. \tag{10}$$

For the unrestricted model this becomes

$$\dim V_{\mathcal{N}}^{G} = (|G|-1)(E-C)+1 = (|G|-1)(2n-3+2C)+1.$$

Proof. The established ingredients are the tree dimension, the dimension of an s-sunlet for $s \geq 4$, and the cut-edge formula

$$\dim V_{\mathcal{N}}^{(G,B)} = \dim V_{\mathcal{N}_1}^{(G,B)} + \dim V_{\mathcal{N}_2}^{(G,B)} - r.$$

They are, respectively, Lemma 11, Theorem 17, and Corollary 10 of [1]. The cut-edge result is the dimension formula for the codimension-zero toric fiber product, applied to the r orbit labels on the cut edge.

Cut all internal cut edges. The resulting blocks are three-leaf trees and sunlets. A three-leaf tree has dimension $3(r-1)+1$; an s-sunlet with $s \geq 4$ has dimension $(2s-1)(r-1)+1$; a three-sunlet has dimension $5(r-1)+1-\delta_B$ by Theorem 7.1. For one gluing, $E_1 + E_2 = E + 1$, $C_1 + C_2 = C$, and the triangle counts add. Substitution in the cut-edge formula preserves (10), so induction over the cut edges proves it. Finally, the degree sum and Euler relation for a binary level-1 network give $E = 2n - 3 + 3C$. □

The unrestricted conclusion removes the odd-order assumption in [2, Theorem 3.8]. It also permits the same removal in that paper's restriction-based distinguishability criterion; if two networks with the same number of cycles have a leaf restriction with different cycle counts, or have distinct tree restrictions, the corresponding varieties are distinguishable.

Appendix A. The five-point contact calculation

We prove generic two-identifiability at order five, completing the contact argument in Section 4. Each support has two colors. Use the spaces W_π^r, W_π^s from the proof of Proposition 4.5 and the reflections in (6); the calculation now concerns lines in a four-dimensional space. For a derangement π, the line W_π^r is spanned by the indicator of the cycle avoiding 0 if π has type $2+3$. For a full cycle its values, in cyclic order after 0, are $6, 7, 8, 9$. The corresponding line W_π^s has values $-9, -8, -7, -6$. These descriptions follow directly from Lemma 4.3.

We need one observation about two Latin colors a, b. If

$$e \in W_a^s + W_b^s \quad \text{or} \quad L_a = L_b, \tag{11}$$

then their permutations are inverse full cycles. For the first assertion, two indicator lines could span e only if their supports were disjoint two-sets covering the four nonzero points. The three-cycle through zero for either permutation would then share an edge with the other's two-cycle, contrary to the Latin condition. An indicator and a full-cycle line cannot span e, because the latter has distinct values on the indicator's support. Finally, two full-cycle lines span e only when their four arithmetic progressions are affinely related. The orderings are then equal or reversed; distinct colors leave only the reversed, inverse cycles. For the second assertion, the two permutations have opposite incoming and outgoing neighbors at zero. A three-cycle through zero leaves a two-cycle on the remaining points; the other permutation would have to use one of its two edges. Both must be full cycles, and their two middle points must occur in reverse order, since equal order would give a common edge.

A Latin system containing a full cycle π and its inverse is the cyclic system; every remaining entry must be $\pi^2(x)$ or $\pi^{-2}(x)$. If one row uses $\pi^2(x)$ at x, it cannot use $\pi^{-2}(\pi^{-1}x) = \pi^2(x)$ at $\pi^{-1}x$; this forces the same choice around the five-cycle. Thus, outside the cyclic system, neither condition in (11) can occur.

The distinguished vertex can be chosen freely in this calculation. Indeed, $F(U, k)$ depends only on the ratios $k_b/k_{\pi_a(b)}$. Simultaneously relabeling the vertices and conjugating the permutations, with an independent reindexing of the colors, merely permutes its output coordinates. A common rescaling of k restores the chosen normalization. We may therefore choose the vertex called 0 and then label each color by its value there.

This freedom leaves only the following three cases. Each word lists the values at $0, 1, 2, 3, 4$; the identity row is omitted.

Type	π_1	π_2	π_3	π_4
1	10342	23401	34120	42013
2	10342	23410	34021	42103
3	12340	23401	34012	40123

To see this, suppose first that some row has type $2 + 3$. Choose 0 in its two-cycle and relabel the row as 10342. Filling unused entries in the columns gives exactly four choices for the second row; $23401, 23410, 24013, 24103$; each forces the remaining two rows. Conjugation by (0 1) interchanges the first and fourth completions, and the second and third, so the first two give types 1 and 2 above. If all rows are full cycles, normalize the first to 12340. A full second row can be $20413, 23401$ or 24103; the first

and last force a $2+3$ row on completion. The middle choice forces the cyclic system, type 3.

Take $A=\{1,2\}$ and $B=\{3,4\}$. Every pair of colors is transitive, since each orbit has at least three elements. The coefficient condition, expressed through (6), is

$$K \in W_1^r + W_2^r, \qquad R_3K, R_4K \in W_3^s + W_4^s. \tag{12}$$

In every displayed type π_3, π_4 are not inverse full cycles. Their plane does not contain e, so (12) also gives $L_3(K) = L_4(K)$. Write vectors by their four nonzero coordinates.

For type 1, write $K = x(0,1,1,1) + yv$, where v is the full-cycle vector for π_2. Then $L_3(K) - L_4(K) = 3y$, so $y = 0$ and $R_3K = x(-2,-1,-1,-1)$. The two full-cycle vectors spanning $W_3^s + W_4^s$ have coordinate sum -30 and first coordinate -7. Thus every vector in that plane has total sum $30/7$ times its first coordinate. The displayed R_3K satisfies this relation only if $x = 0$.

For type 2 the two planes give

$$K = (y, x, x+y, x),$$
$$R_3K = (-2x, -x-y, -x, -x-y) \in \{(a+b, b, 0, a) : a, b \in \mathbb{C}\}.$$

The third coordinate gives $x = 0$, and the other three then give $y = 0$.

Only the cyclic type remains. Here

$$K = x(6,7,8,9) + y(8,6,9,7).$$

The functional $(1,-6,0,4)$ annihilates $W_3^s + W_4^s$. Its value on R_3K is $15(x+y)$, so $K = \lambda(-2,1,-1,2)$. Write $d = p_1 - p_2$ and use addition modulo five. Commuting the two A-color equations of the torus condition gives

$$W_H = d\Delta_1\Delta_2 H + \Delta_2(\Delta_1 K\, \Delta_1 H) - \Delta_1(\Delta_2 K\, \Delta_2 H) \in \operatorname{span}\{\Delta_2 r_1, \Delta_1 r_2\}.$$

Every vector in the latter space vanishes at 1 and satisfies $2W(0) - W(2) = 0$. The equation at 1 gives $d = 2\lambda$. Substitution then gives

$$W_H(0) = 5\lambda\Delta_1\Delta_2 H(0), \qquad W_H(2) = 5\lambda\Delta_1\Delta_2 H(2).$$

Taking $H = \mathbf{1}_{\{0\}}$ in the remaining relation gives $5\lambda = 0$. Hence $K = 0$ in the cyclic case as well.

Once $K = 0$, the dimension and reflection arguments in Proposition 4.5 make p constant on each of A, B. If its value on B were nonzero, the difference of the two B-equations would put $(R_3 - R_4)Q = \mathbb{C}e$ in $W_3^s + W_4^s$. This contradicts (11), since these two colors are not inverse cycles. Thus the contact kernel is radial also at order five.

Appendix B. Latin systems without an identity row

The shared h-family prevents arbitrary Latin-square isotopy from normalizing the polynomial map. A second use of the one-hole lemma removes the identity-row hypothesis for every order at least six.

Theorem B.1. *Let $\pi_1, \dots, \pi_m$ be the rows of any Latin square on Ω, where $m = |\Omega| \geq 6$. Put $\iota = 1$ if one row is the identity permutation, and $\iota = 0$ otherwise. Over every algebraically closed field, the image closure of*

$$f_{a,b} = s_b t_{\pi_a(b)}\big(u_a h_{\pi_a(b)} + v_a h_b\big)$$

has dimension

$$5m - 3 - \iota.$$

Lemma B.2. *Let $\{\pi_a\}$ be a Latin system on Ω, with $m \geq 6$. Fix $o \in \Omega$, and omit the unique color a_* satisfying $\pi_{a_*}(o) = o$. The other colors admit a partition $P \sqcup N$, with $|P| \geq 3$ and $|N| \geq 2$, such that both*

$$\big\langle \pi_a^{-1}\pi_b : a, b \in P \big\rangle \quad \text{and} \quad \big\langle \pi_a : a \in N \big\rangle$$

act transitively on Ω.

Proof. View the colors as the perfect matchings of a factorization of $K_{m,m}$. A color family has connected union precisely when its relative permutations act transitively.

First suppose $m \geq 10$. Choose a reference color $\sigma \neq a_*$ and normalize by $\rho_a = \pi_a^{-1}\pi_\sigma$. Among the resulting nonidentity colors, forbid a_*. The $m-2$ available colors together act transitively, since every orbit has at least $m-1$ points. If the colors chosen so far have $k > 1$ orbits, the total number of incidences between available colors and preserved orbits is at most $m - k$; an orbit C is preserved by at most $|C| - 1$ nonidentity colors. For $k \geq 3$, some available color therefore preserves none of the orbits; for $k = 2$, transitivity supplies a color crossing between them. Adjoining such a color at least halves the number of orbits. Thus at most $\lfloor \log_2 m \rfloor$ colors suffice. Adjoin σ and pad to at least three colors to obtain P, with

$$|P| \leq \lfloor \log_2 m \rfloor + 1 \leq m/2 - 1.$$

Its complement N among the admissible colors has $|N| \geq m/2$. Under the original N-permutations, the orbit of o has at least $|N| + 1$ points, and any other orbit has at least $|N|$ points. There is consequently only one orbit.

For $m = 6$, label the five admissible colors by $a = \pi_a(o)$. At most two triples of these colors have disconnected matching union. Indeed, such a union consists of two

$K_{3,3}$ components. Any two of its colors already connect each component, and the third matching is forced as the unused edge at each vertex. Thus two bad triples cannot share two colors. Their complementary pairs in the five-set are consequently disjoint, so there are at most two.

At least three pairs of admissible original permutations act transitively. For a nontransitive pair a, b, the orbit of o contains o, a, b, and every orbit has at least two points. Its orbit partition is therefore $3+3$ or $4+2$. The pairs of the first kind form a matching; their root orbit is $\{o, a, b\}$, so if $\pi_a(a) \neq o$ then $b = \pi_a(a)$, while if $\pi_a(a) = o$ then $\pi_b(a) = a$, which again determines b uniquely. For a pair of the second kind, one permutation fixes the complementary two-set pointwise and the other exchanges its points. Assign the pair to the fixing permutation. A permutation with f fixed points receives at most $\min\{4, \binom{f}{2}\} \leq f$ pairs; each two-set can be exchanged by at most one other color. Here $f \leq 4$, and the fixed-point sets of the five admissible permutations are disjoint subsets of $\Omega \setminus \{o\}$. Hence at most $2+5=7$ of the ten pairs are nontransitive. Choose a transitive pair N whose complementary triple P is not bad; the required partition follows.

For $m = 2r+1 \in \{7, 9\}$, divide the $2r$ admissible colors into two r-sets. Call an r-set bad if its matching union is disconnected. Its components then have sizes r and $r+1$ on each side, so the smaller component is a Latin subsquare of order r using exactly those colors. Two bad color sets cannot share $r-1$ colors. Indeed, if their row sets or column sets meet, the common colors force both intersections to have size $r-1$; the intersecting blocks give a subsquare of that order inside one of order r. A proper subsquare has at most half the order of its containing square; in a row outside it, its symbols must occur outside its columns. This is impossible for $r \geq 3$. If instead the two row sets and the two column sets are disjoint, the cross $r \times r$ block uses the r colors outside their union. In each row of that block, the one remaining column must then carry the same color, again impossible. Counting $(r-1)$-subsets gives

$$\#\{\text{bad sets}\} \leq \frac{1}{r}\binom{2r}{r-1} = \frac{1}{r+1}\binom{2r}{r} < \frac{1}{2}\binom{2r}{r}.$$

Some complementary pair therefore has both matching unions connected. Take it for P, N; relative transitivity of N implies transitivity of its original permutations as well.

It remains to treat $m = 8$. If three admissible colors have connected union, use them for P and apply the orbit-size argument to the four colors in N. Otherwise normalize relative to an admissible reference color. The other six admissible colors are derangements, no two of which generate a transitive group. Each is a fixed-point-free involution. To see this, a permutation with two cycles would force all six colors to preserve its smaller cycle, contradicting the bound $|C| - 1$. With three cycles, every

color preserves at least one of them, whereas the total number of such incidences is at most $8-3=5$. An eight-cycle is already transitive, leaving only four two-cycles. The matchings of any two involutions form two alternating four-cycles, so the involutions commute. Their group has an orbit containing a point and its six distinct images; as an elementary abelian 2-group, it is therefore transitive and hence regular. The omitted row is forced pointwise to be the eighth translation.

Undoing the normalization expresses the original system as a coset of this regular group. Conjugating the group if necessary, label Ω by $G=(\mathbb{Z}/2)^3$ with $o=0$. Then the original rows have the form $\pi_a=\tau_a\sigma'$, where τ_a is translation by $a\in G$ and $\sigma'=\pi_{a_*}$ fixes 0. Choose a plane $H<G$ and $c\notin H$, and put

$$P=(H\setminus\{0\})\cup\{c\},\qquad N=(G\setminus H)\setminus\{c\}.$$

The differences of P span G, so its relative action is transitive. For $a,b\in N$, the products $\pi_a\pi_b^{-1}=\tau_{a+b}$ generate the translations by H. Since each N-row also takes 0 outside H, its original action is transitive. □

Proof of Theorem B.1. Fix $o\in\Omega$, let a_* be the unique color fixing it, and take P,N from Lemma B.2. Fix $\sigma\in P$. All anchor colors in $P\cup N$ move o.

Give h_o weight 2, and give the other h_b pairwise distinct rational weights in $(0,1/4)$. Give s,t weight zero and use $(\mathrm{wt}(u_a),\mathrm{wt}(v_a))=(0,1)$ on P and $(1,0)$ on N. Their selected terms are exactly the P,N rows of the table in Section 2; P selects u except at the incoming edge to o, while N selects v except at the outgoing edge from o.

If π_{a_*} is not the identity, its restriction to $\Omega\setminus\{o\}$ is nontrivial. The differences $\mathrm{wt}(h_{\pi_{a_*}b})-\mathrm{wt}(h_b)$ on that set have both positive and negative values. Choose the difference $\mathrm{wt}(u_{a_*})-\mathrm{wt}(v_{a_*})$ sufficiently small and nonzero, avoiding all finitely many ties, so that both summands are selected among these off-root coordinates. If π_{a_*} is the identity, choose any nonzero difference, so precisely one of its two coefficients is used. Multiplying all weights by a positive integer makes them integral.

Compute the selected exponent kernel over $\mathbb{Q}$, with coordinates (S,T,H,U,V). For $a\in P$, write the ordinary u-equations at a common output $y\neq o$

$$S_{\pi_a^{-1}y}+T_y+H_y+U_a=0.$$

Subtract the equation for σ. With $x=\pi_\sigma^{-1}y$ this becomes a constant finite difference for S under $\rho_a=\pi_a^{-1}\pi_\sigma$, with the single omitted argument $x=\pi_\sigma^{-1}o$. There are at least two distinct nonidentity relative permutations. The outgoing-hole lemma and transitivity imply that S is constant, say $S=s$. The same equations then make $T+H$ constant, say k, on $\Omega\setminus\{o\}$.

For $a \in N$, the ordinary v-equations at $b \neq o$ now read

$$T_{\pi_a b} - T_b = -s - k - V_a.$$

The outgoing-hole lemma applies directly to the original N-permutations. Their transitivity makes T constant, say t, and hence H constant away from o, say h. Every selected monomial of an anchor color in $P \cup N$ uses an h-index different from o. Each anchor uses both of its coefficients, so all its U, V coordinates equal $-s - t - h$.

The off-root equations of color a_* fix both of its coefficient coordinates to $-s-t-h$ if its permutation is nonidentity; in the identity case they fix the single selected coefficient and leave the other unused. Finally its coordinate at $b = o$ gives $H_o = h$. Consequently the exponent kernel has dimension exactly $3 + \iota$; the three constants s, t, h, and the unused coefficient precisely when an identity row exists. Hence its rational rank is $5m - 3 - \iota$, and initial-monomial independence gives that lower bound in every characteristic.

The three independent scalings of Section 1 give the matching upper bound, with the additional combination $u_a + v_a$ when there is an identity row. □

Appendix C. The exponent lattice for groups

The dimension proof uses the rational rank of an exponent matrix. Its integral row lattice also records finite torus kernels and possible inseparability. For group translations, the calculation below exhibits the same pair of character groups that occurs in the parameter fibers.

For a finite group G, use its translation Latin system and a partition P, N into generating sets of size at least two. Let L_E be the row lattice of the selected exponent matrix in $\mathbb{Z}^{5m}$.

Proposition C.1. *Writing G_{ab} for the abelianization, one has*

$$\mathbb{Z}^{5m}/L_E \cong \mathbb{Z}^4 \oplus G_{\mathrm{ab}} \oplus G_{\mathrm{ab}}.$$

In particular, for abelian G, the row lattice has saturation index m^2. If G has nontrivial invariant factors $d_1 \mid \cdots \mid d_s$, the nonunit Smith factors are $d_1, d_1, \ldots, d_s, d_s$.

Proof. We compute the kernel of the selected monomial map on the complex torus; its character group is the required cokernel of the exponent lattice. Use multiplicative group notation and the translations $b \mapsto ab$. The identity-color equations give $t_b = (s_b h_b v_e)^{-1}$. Put $k_b = s_b h_b$. The ordinary selected rows give

$$s_{ab}/s_b = u_a/v_e \quad (a \in P,\ b \neq a^{-1}), \qquad k_{ab}/k_b = v_a/v_e \quad (a \in N,\ b \neq e).$$

The omitted ratios are equal to the displayed constants as well. For a noncyclic group, multiplication around a translation orbit outside $\langle a \rangle$ shows that the constant r_a satisfies $r_a^{\mathrm{ord}(a)} = 1$; multiplication around the exceptional cycle then fills its omitted edge. For a cyclic group, multiplying all ratios gives a single defect $z_a = r_a^{-m}$ at the omitted edge. For the incoming-hole equations, the ratio under translation by a is $r_a z_a^{\mathbf{1}_{\{a^{-1}\}}(b)}$. Commuting two distinct nonidentity translations a, d and evaluating at $b = a^{-1}$ and $b = d^{-1}$ gives $z_a = z_d = 1$; the common translated defect location $a^{-1}d^{-1}$ is distinct from both. For the outgoing-hole equations, the same argument uses defect locations e, a^{-1}, d^{-1}. Thus the omitted edges are filled in either convention.

Because each color set generates G, it follows that $s_b = A\chi(b)$ and $k_b = B\psi(b)$ for two characters $\chi, \psi : G \to \mathbb{C}^*$. The exceptional equations determine the other parameters. Every kernel point has the unique form

$$\begin{aligned} &s_b = A\chi(b), && h_b = (B/A)\psi(b)/\chi(b), \quad t_b = (BC)^{-1}\psi(b)^{-1}, \\ &v_a = C\psi(a) \quad \text{for all } a, && u_a = C\chi(a) \quad (a \neq e), \quad u_e = D, \end{aligned}$$

where $A, B, C, D \in \mathbb{C}^*$ are free. These formulas also suffice and give an algebraic-group isomorphism

$$\ker M_E \cong (\mathbb{C}^*)^4 \times \widehat{G_{\mathrm{ab}}} \times \widehat{G_{\mathrm{ab}}}. \qquad \square$$

This explains why a Jacobian may lose rank in a characteristic dividing $|G|$ even though the image dimension remains unchanged; a finite torus kernel can become nonreduced. The initial-monomial proof uses rational rank of the exponent lattice and avoids this issue entirely.

References

[1] E. Gross, R. Krone, and S. Martin, *Dimensions of Level-1 Group-Based Phylogenetic Networks*, Bulletin of Mathematical Biology **86** (2024), article 90. DOI 10.1007/s11538-024-01314-z.

[2] S. Cox, E. Gross, and S. Martin, *Group-based phylogenetic models on 3-sunlet networks*, Bulletin of Mathematical Biology **87** (2025), article 132. DOI 10.1007/s11538-025-01506-1.

[3] A. T. Blomenhofer, A. Casarotti, M. Michałek, and A. Oneto, *Identifiability for mixtures of centered Gaussians and sums of powers of quadratics*, Bulletin of the London Mathematical Society **55** (2023), 2407–2424. DOI 10.1112/blms.12871.