A Logarithmic Universality Principle for Cycle Counts in Wreath Products $\Gamma^n \rtimes H_n$

Liz Lemma   Future Detective

January 17, 2026

Abstract

Diaconis–Tung prove that for fixed k and $\Gamma \leq S_k$, the small-cycle counts of a uniform permutation in the imprimitive wreath product $\Gamma^n \rtimes S_n \leq S_{kn}$ converge to an explicit (generally dependent) compound Poisson vector, with a coupling bound of order $O(B/n)$ for the first B cycle counts. This work abstracts the role played by S_n and replaces it by a general sequence of conjugacy-invariant measures μ_n on S_n (including uniform measures on subgroups $H_n \leq S_n$) whose Poissonized cycle index is asymptotically logarithmic in the sense of Arratia–Barbour–Tavaré. Under a verifiable analytic hypothesis on the Poissonized cycle index—expressed as a uniform control on the logarithm of the truncated cycle generating function—we prove a transfer theorem: the small-cycle vector of $\Gamma^n \rtimes \mu_n$ converges to a marked Poisson pushforward law obtained by placing i.i.d. Γ-cycle-type marks on the limiting Poisson process of block-cycles from μ_n. The result provides a black-box method: once cycle asymptotics for μ_n are known (e.g. A_n, Ewens(θ), derangements, restricted-cycle measures), the wreath-product cycle limits follow immediately with explicit dependence structure and quantitative approximation.

Table of Contents

4. 4. Exact Poissonized transfer: prove that Poissonization + wreath composition turns the (α_ℓ)-logarithmic structure of (μ_n) into a marked Poisson process with marks in partitions of k.
5. 5. De-Poissonization with uniform error: an analytic Tauberian lemma tailored to truncated cycle generating functions, yielding convergence for fixed n and a quantitative TV bound.
6. 6. Examples and computations: (i) $\mu_n = \mathrm{Unif}(S_n)$ $(\alpha_\ell \equiv 1)$, (ii) $\mu_n = \mathrm{Unif}(A_n)$, (iii) Ewens$(\theta)$, (iv) derangements, (v) restricted cycle-length measures. Explicitly compute the resulting (A_i) formulas in selected small k cases.
7. 7. Extensions and limitations: allowing μ_n on other permutation sets, $H_n \leq S_n$ that are not logarithmic (e.g. C_n), and what breaks; discussion of how to incorporate H_n acting on a different ground set size (e.g. B_n acting on $2n$) with notational changes.
8. 8. Open problems and next steps: mesoscopic windows $B = B(n)$, sharper constants, and classification of which families (H_n) satisfy the hypothesis without case-by-case cycle-index computation.

1 Introduction and statement of the universality principle

We study the cycle structure of random permutations on a set of size kn which preserve a fixed block system of n blocks of size k. Concretely, we identify $[kn]$ with $[n] \times [k]$ and let $\Gamma^n \rtimes S_n$ act imprimitively by

$$(\gamma_1, \dots, \gamma_n; h) \cdot (i, r) := (h(i),\, \gamma_{h(i)}(r)),$$

so that h permutes the blocks while the $\gamma_i \in \Gamma \leq S_k$ act within blocks. The random input consists of a conjugacy-invariant law μ_n on S_n for the block permutation h, and i.i.d. internal permutations $\gamma_i \sim \mathrm{Unif}(\Gamma)$ independent of h. The induced random permutation on $[kn]$ will be denoted

$$\sigma = (\gamma_1, \dots, \gamma_n; h) \in S_{kn}.$$

Our objective is to understand, for fixed B, the joint distribution of the first B cycle counts

$$(a_1(\sigma), \dots, a_B(\sigma)),$$

where $a_i(\sigma)$ is the number of i-cycles of σ as a permutation of $[kn]$. We regard this as a local statistic: as $n \to \infty$ with k fixed, cycles of length $O(1)$ only probe $O(1)$ blocks and should therefore admit a limit description under broad conditions on the block-level randomness μ_n.

The starting point is the observation that the cycle structure of σ is obtained by a deterministic transformation of two ingredients: (i) the cycle structure of the block permutation $h \in S_n$; (ii) the cycle types of certain Γ-valued "monodromies" along the cycles of h. More precisely, if h has an ℓ-cycle on blocks,

$$(i_1\, i_2\, \cdots\, i_\ell),$$

then σ restricted to the union of these ℓ blocks acts as a permutation of $[k\ell]$ obtained by composing the within-block permutations along the cycle. The relevant Γ-element is the product

$$g := \gamma_{i_\ell} \gamma_{i_{\ell-1}} \cdots \gamma_{i_1} \in \Gamma,$$

whose cycle type in S_k we denote by a partition $\lambda \vdash k$. If λ has $a_j(\lambda)$ parts of size j, then this block-cycle contributes exactly $a_j(\lambda)$ cycles of length $j\ell$ to σ. Thus, for each $i \geq 1$, the count $a_i(\sigma)$ can be expressed as a sum over block-cycle lengths $\ell \mid i$ of contributions coming from ℓ-cycles in h, each contribution depending on the cycle type of the associated product $g \in \Gamma$.

This mechanism makes it plausible that the small-cycle statistics of σ depend on μ_n only through the small-cycle statistics of h. The point of the present work is that this dependence is universal once one assumes a weak Poisson approximation for the block permutation, formulated in terms of a

Poissonized cycle generating function. Under this hypothesis, the block-level cycle counts behave, after Poissonization and truncation, as if the numbers of ℓ-cycles were asymptotically independent Poisson with means α_ℓ/ℓ for a fixed sequence $(\alpha_\ell)_{\ell\geq 1}$. Our main theorem shows that the same Poisson parameters α_ℓ govern the limit law of $(a_1(\sigma),\ldots,a_B(\sigma))$, with the internal group Γ entering only through the distribution of cycle type of a uniform element of Γ.

We describe the limit object informally. Fix $\ell \geq 1$. In the block permutation h, an ℓ-cycle on blocks will, in the limit model, occur with asymptotically Poisson frequency α_ℓ/ℓ. Each such ℓ-cycle carries an independent mark $\lambda \vdash k$, interpreted as the cycle type of an independent uniform element of Γ, with probability

$$P_\Gamma(\lambda) := \mathbb{P}_{\gamma\sim\mathrm{Unif}(\Gamma)}\big(\mathrm{type}(\gamma)=\lambda\big).$$

Hence the limiting block-cycle data are a collection of independent Poisson random variables

$$Z_{\ell,\lambda} \sim \mathrm{Poisson}\Big(\frac{\alpha_\ell P_\Gamma(\lambda)}{\ell}\Big), \qquad \ell\geq 1,\ \lambda\vdash k,$$

where $Z_{\ell,\lambda}$ counts block-cycles of length ℓ whose internal mark equals λ. The corresponding induced cycle counts in S_{kn} are then obtained by the deterministic pushforward

$$A_i := \sum_{\ell\geq 1}\sum_{\lambda\vdash k}\sum_{\substack{j\geq 1\\ j\ell=i}} a_j(\lambda)\, Z_{\ell,\lambda}, \qquad i\geq 1.$$

Since the variables $Z_{\ell,\lambda}$ are independent and A_i is a finite sum for each fixed i, the vector $(A_1,\ldots,A_B)$ is a compound Poisson (more precisely, marked-Poisson pushforward) vector. The theorem asserts that $(a_1(\sigma),\ldots,a_B(\sigma))$ converges in distribution to $(A_1,\ldots,A_B)$, with an explicit quantitative bound in total variation under the assumed analytic control of the Poissonized generating function.

In this formulation, the universality principle is the following: *the small-cycle law of the imprimitive wreath-product permutation σ depends on the sequence (μ_n) only through the asymptotic Poisson parameters (α_ℓ) governing the small-cycle law of the block permutation h, and depends on Γ only through the distribution of cycle types of a uniform element of Γ.* In particular, any two sequences (μ_n) and (μ'_n) which share the same Poissonized cycle-index asymptotics (equivalently, the same parameters α_ℓ at the level of fixed truncations) yield the same limiting joint law for $(a_1(\sigma),\ldots,a_B(\sigma))$ for every fixed B.

One can read off from the limit model several concrete consequences that are stable across the entire class of admissible block measures. For example, taking expectations in the definition of A_i gives

$$\mathbb{E}A_i = \sum_{\ell | i} \sum_{\lambda \vdash k} a_{i/\ell}(\lambda) \, \frac{\alpha_\ell P_\Gamma(\lambda)}{\ell} = \sum_{\ell | i} \frac{\alpha_\ell}{\ell} \, \mathbb{E}_{\gamma \sim \mathrm{Unif}(\Gamma)} \big[a_{i/\ell}(\gamma) \big].$$

Thus the mean number of i-cycles in σ is a divisor sum of the block-level intensities α_ℓ/ℓ, weighted by the expected number of (i/ℓ)-cycles in a uniform element of Γ. More generally, joint factorial moments of $(A_1, \dots, A_B)$ are polynomial expressions in the parameters α_ℓ and the cycle index of Γ. The theorem shows that the same formulas asymptotically govern the corresponding moments of $(a_1(\sigma), \dots, a_B(\sigma))$.

The theorem includes the special case in which μ_n is uniform on a subgroup $H_n \le S_n$ (with H_n conjugacy-invariant as a set). In that setting, the parameters α_ℓ encode the asymptotic ℓ-cycle intensities in a random element of H_n, as detected by the Poissonized truncated cycle index. Thus our result transfers information from a sequence of block-level groups (H_n) to the induced imprimitive permutation model on $[kn]$. The transfer is explicit: once the values of $\alpha_1, \dots, \alpha_B$ are known (or bounded) for H_n, the theorem yields the limiting law of $(a_1(\sigma), \dots, a_B(\sigma))$ in the wreath-product model with internal group Γ.

The comparison point is the classical uniform case $\mu_n = \mathrm{Unif}(S_n)$. In that case the cycle counts $(b_1(h), \dots, b_L(h))$ converge to independent Poisson with means $1/\ell$, and the present hypothesis holds with $\alpha_\ell \equiv 1$. The resulting limit law for σ is then described by

$$Z_{\ell,\lambda} \sim \mathrm{Poisson}\Big(\frac{P_\Gamma(\lambda)}{\ell} \Big), \qquad A_i = \sum_{\ell | i} \sum_{\lambda \vdash k} a_{i/\ell}(\lambda) \, Z_{\ell,\lambda}.$$

This recovers, in our notation and level of generality, the asymptotic small-cycle description obtained by Diaconis–Tung for random elements in wreath products under the imprimitive action (in particular, for $\Gamma \wr S_n$ sampled from the product of uniform measures). The new point here is not the computation of this limit in the uniform case, but rather the stability of the same structural conclusion for a broad class of non-uniform block measures μ_n, including measures supported on proper subgroups H_n and measures with nontrivial global dependencies, provided that their Poissonized cycle indices satisfy the stated asymptotic expansion with controlled remainder.

It is useful to emphasize what is and is not assumed about (μ_n). We do not assume that the cycle counts $b_\ell(h)$ are independent, nor do we assume an Ewens-type product structure, nor even that μ_n arises from a consistent family of measures under restriction maps. The hypothesis is analytic: it controls the logarithm of a Poissonized truncated cycle generating function and asserts that it is close, uniformly on a complex polydisc, to the logarithm of the generating function of independent Poisson variables with intensities

α_ℓ/ℓ. This is exactly the level of information needed to deduce Poisson limits for fixed collections of small cycle counts, together with a quantitative de-Poissonization error. The conclusion then transports these Poisson limits through the wreath-product construction.

At a conceptual level, the argument proceeds in two steps. First, we show that for fixed L, the joint law of the marked block-cycle counts (block cycles of length $\ell \le L$ together with the Γ-cycle-type mark obtained from the monodromy product along the block cycle) is asymptotically that of independent Poisson variables with the indicated means. This step uses only conjugacy-invariance of μ_n and the Poissonized cycle-index hypothesis. Second, we express $(a_1(\sigma), \dots, a_B(\sigma))$ as a deterministic function of these marked block-cycle counts, up to an error coming from block cycles longer than B (which cannot contribute to cycles of length at most B). Since k is fixed, there is no additional approximation at this stage: an ℓ-cycle on blocks contributes only to cycle lengths that are multiples of ℓ, and hence for the first B cycle counts we need only consider block cycles with $\ell \le B$. The limiting marked-Poisson description is therefore pushed forward exactly to the stated limit law for $(a_1, \dots, a_B)$.

The quantitative total-variation estimate supplied by the theorem should be interpreted similarly: the error is controlled uniformly in the choice of bounded test functions depending only on $(a_1, \dots, a_B)$, and it is inherited from the analytic remainder bound in the Poissonized generating function for μ_n via an analytic Tauberian de-Poissonization lemma. In particular, once a family (μ_n) is shown to satisfy the cycle-index hypothesis with an explicit rate, the same rate (up to constants depending on k and Γ and linear growth in B) propagates automatically to the wreath-product model.

We regard this as a universality statement because the wreath-product small-cycle statistics do not depend on any finer details of μ_n than those encoded in the parameters (α_ℓ), and the dependence on Γ is mediated entirely by the cycle-type law P_Γ. In subsequent sections we develop the minimal toolkit needed to implement this program: cycle index generating functions and their composition under wreath products, and the marked-Poisson (compound Poisson) formalism that captures the limiting cycle-count vector.

2 Background: cycle indices, wreath-product composition, and marked Poisson vectors

We collect the basic algebraic and probabilistic devices used throughout. The overarching theme is that small-cycle data are encoded by multivariate generating functions, and that the wreath-product construction corresponds to an explicit substitution rule for these generating functions. The limiting objects are naturally expressed as marked (multi-type) Poisson families and their deterministic pushforwards.

2.1 Cycle index generating functions for conjugacy-invariant laws

For $n \geq 1$ and $h \in S_n$, recall that $b_\ell(h)$ denotes the number of ℓ-cycles of h (so that $\sum_{\ell \geq 1} \ell\, b_\ell(h) = n$). For any conjugacy-invariant probability measure μ_n on S_n, it is convenient to package the joint law of the cycle counts $(b_1, \dots, b_n)$ into the *cycle index* (or cycle-count generating function)

$$Z_{\mu_n}(x_1, \dots, x_n) := \mathbb{E}_{h \sim \mu_n}\Big[\prod_{\ell=1}^{n} x_\ell^{b_\ell(h)}\Big], \qquad (x_1, \dots, x_n) \in \mathbb{C}^n. \tag{1}$$

When $\mu_n = \mathrm{Unif}(G_n)$ is uniform on a conjugacy-invariant subset $G_n \subseteq S_n$ (e.g. a subgroup), this reduces to the normalized cycle index sum

$$Z_{\mu_n}(x_1, \dots, x_n) = \frac{1}{|G_n|} \sum_{h \in G_n} \prod_{\ell=1}^{n} x_\ell^{b_\ell(h)}.$$

We will also use the truncated version

$$Z_{\mu_n}^{(\leq L)}(x_1, \dots, x_L) := \mathbb{E}_{h \sim \mu_n}\Big[\prod_{\ell=1}^{L} x_\ell^{b_\ell(h)}\Big], \tag{2}$$

which is obtained from (1) by setting $x_\ell \equiv 1$ for $\ell > L$. This truncation is natural because for local questions involving only cycles up to some fixed size, only finitely many $b_\ell(h)$ can contribute.

For the internal group $\Gamma \leq S_k$ we similarly define its cycle index polynomial

$$Z_\Gamma(y_1, \dots, y_k) := \mathbb{E}_{\gamma \sim \mathrm{Unif}(\Gamma)}\Big[\prod_{j=1}^{k} y_j^{a_j(\gamma)}\Big] = \sum_{\lambda \vdash k} P_\Gamma(\lambda) \prod_{j=1}^{k} y_j^{a_j(\lambda)}. \tag{3}$$

Here $a_j(\gamma)$ denotes the number of j-cycles of γ (as a permutation of $[k]$), and for a partition $\lambda \vdash k$ we write $a_j(\lambda)$ for the number of parts of size j in λ. The last identity in (3) is simply the decomposition of $\mathrm{Unif}(\Gamma)$ by cycle type.

Two elementary observations will be used repeatedly.

- The maps $x \mapsto Z_{\mu_n}(x)$ and $y \mapsto Z_\Gamma(y)$ are *polynomials* with nonnegative coefficients when restricted to $x_\ell, y_j \geq 0$, and they are analytic on all of $\mathbb{C}^n$ and $\mathbb{C}^k$ respectively.
- Factorial moments of the vector $(b_1, \dots, b_L)$ can be read off from derivatives of $Z_{\mu_n}^{(\leq L)}$ at $x_\ell = 1$, and similarly for $(a_1(\gamma), \dots, a_k(\gamma))$ from Z_Γ.

In particular, for fixed indices $\ell_1, \dots, \ell_m \leq L$,

$$\mathbb{E}\Big[\prod_{r=1}^{m} (b_{\ell_r})_{q_r}\Big] = \Big(\prod_{r=1}^{m} \partial_{x_{\ell_r}}^{q_r}\Big) Z_{\mu_n}^{(\leq L)}(x)\Big|_{x \equiv 1},$$

where $(u)_q = u(u-1)\cdots(u-q+1)$ is the falling factorial.

2.2 Cycle structure under the imprimitive wreath-product action

We next record the precise substitution rule that describes the induced permutation $\sigma \in S_{kn}$ in terms of the block permutation $h \in S_n$ and the internal permutations $\gamma_1, \dots, \gamma_n \in \Gamma$.

Fix n. Let $\sigma = (\gamma_1, \dots, \gamma_n; h)$ act on $[n] \times [k]$ by

$$(\gamma_1, \dots, \gamma_n; h) \cdot (i, r) = (h(i), \gamma_{h(i)}(r)).$$

Consider a single ℓ-cycle $c = (i_1\, i_2\, \cdots\, i_\ell)$ of h. Restricting σ to the union of the corresponding ℓ blocks, we see that after ℓ applications of σ we return to the same block, and the action on the internal coordinate is given by the product

$$g_c := \gamma_{i_\ell}\gamma_{i_{\ell-1}} \cdots \gamma_{i_1} \in \Gamma, \tag{4}$$

which we refer to as the *monodromy* along the block cycle c. The cycle structure of σ on these $k\ell$ points is determined by the cycle structure of g_c on $[k]$: each j-cycle of g_c lifts to a $(j\ell)$-cycle of σ. Consequently, writing $\mathrm{type}(g_c) = \lambda \vdash k$, the block cycle c contributes $a_j(\lambda)$ cycles of length $j\ell$ to σ.

This description immediately yields a generating function identity. For formal variables $(x_i)_{i\geq 1}$ and fixed B, define the truncated cycle monomial for σ by

$$M_B(\sigma; x) := \prod_{i=1}^{B} x_i^{a_i(\sigma)}.$$

Similarly, define for $\gamma \in \Gamma$ and $\ell \geq 1$ the *ℓ-lift monomial*

$$\Phi_{\ell,B}(\gamma; x) := \prod_{\substack{j\geq 1 \\ j\ell \leq B}} x_{j\ell}^{a_j(\gamma)}. \tag{5}$$

Then by multiplicativity over disjoint block cycles we have the pointwise identity

$$M_B(\sigma; x) = \prod_{\substack{\text{block cycles } c \text{ of } h \\ \mathrm{len}(c) \leq B}} \Phi_{\mathrm{len}(c),B}(g_c; x), \tag{6}$$

since block cycles of length $> B$ cannot contribute to $a_i(\sigma)$ for $i \leq B$.

Taking conditional expectation given h simplifies further because the monodromies g_c along distinct cycles c are independent and uniform on Γ. Indeed, the cycles of h partition $[n]$, and each g_c is a product of i.i.d. uniform Γ-elements along that cycle; since Haar-uniform on a finite group is invariant under convolution, we have $g_c \sim \mathrm{Unif}(\Gamma)$ for each c, and independence follows because the underlying γ_i are independent across disjoint cycles. Therefore, defining the *ℓ-lift cycle index* by

$$G_{\ell,B}(x_1,\ldots,x_B) := \mathbb{E}_{\gamma\sim\mathrm{Unif}(\Gamma)}\big[\Phi_{\ell,B}(\gamma;x)\big] = \sum_{\lambda\vdash k} P_\Gamma(\lambda) \prod_{\substack{j\geq 1\\ j\ell\leq B}} x_{j\ell}^{a_j(\lambda)}, \quad (7)$$

we obtain from (6) that

$$\mathbb{E}\big[M_B(\sigma;x)\,\big|\,h\big] = \prod_{\ell=1}^{B} G_{\ell,B}(x)^{b_\ell(h)}. \quad (8)$$

Averaging over $h\sim\mu_n$ yields the fundamental substitution identity

$$\mathbb{E}\big[M_B(\sigma;x)\big] = \mathbb{E}_{h\sim\mu_n}\Big[\prod_{\ell=1}^{B} G_{\ell,B}(x)^{b_\ell(h)}\Big] = Z_{\mu_n}^{(\leq B)}\big(G_{1,B}(x),\, G_{2,B}(x),\, \ldots,\, G_{B,B}(x)\big). \quad (9)$$

In words: the truncated cycle index of σ is obtained from the truncated cycle index of h by substituting, for each block-cycle length ℓ, the effective weight $G_{\ell,B}(x)$ coming from the internal group Γ.

It is useful to connect $G_{\ell,B}$ with the internal cycle index Z_Γ. If we momentarily ignore truncation in B, the formal identity corresponding to (7) is

$$G_\ell(x_1,x_2,\ldots) = \mathbb{E}_{\gamma\sim\mathrm{Unif}(\Gamma)}\Big[\prod_{j=1}^{k} x_{j\ell}^{a_j(\gamma)}\Big] = Z_\Gamma\big(x_\ell, x_{2\ell},\ldots,x_{k\ell}\big), \quad (10)$$

where x_m for $m>B$ may be set to 1 when working with M_B. Thus (9) is the wreath-product cycle index composition rule specialized to our random model and truncated at small cycle lengths.

2.3 Marked block cycles and the multi-type Poisson viewpoint

The substitution formula (9) can be refined by keeping track of internal cycle types explicitly. For $\ell\geq 1$ and $\lambda\vdash k$, let

$$C_{\ell,\lambda} := \#\{\ell\text{-cycles } c \text{ of } h \text{ such that } \mathrm{type}(g_c)=\lambda\}.$$

Then $\sum_{\lambda\vdash k} C_{\ell,\lambda} = b_\ell(h)$, and the cycle counts of σ can be written deterministically as

$$a_i(\sigma) = \sum_{\ell\mid i}\sum_{\lambda\vdash k} a_{i/\ell}(\lambda)\, C_{\ell,\lambda}, \qquad i\geq 1. \quad (11)$$

Moreover, conditional on h, the random vector $(C_{\ell,\lambda})_{\lambda\vdash k}$ is multinomial with total count $b_\ell(h)$ and cell probabilities $(P_\Gamma(\lambda))_{\lambda\vdash k}$, and these multinomial vectors are independent over distinct ℓ's. Equivalently, one may view the family $(C_{\ell,\lambda})_{\ell,\lambda}$ as obtained from the unmarked counts $(b_\ell(h))_\ell$ by independent marking: each block cycle of length ℓ receives an independent mark λ distributed as the cycle type of a uniform element of Γ.

This marked formulation is the natural bridge to the limiting model stated in the introduction: if in some asymptotic regime the unmarked counts $b_\ell(h)$ are close to independent Poisson variables, then independent marking yields *multi-type* independent Poisson variables. Concretely, suppose $b_\ell(h) \approx \text{Poisson}(\alpha_\ell/\ell)$, independently over ℓ. Conditional on $b_\ell(h) = m$, marking splits these m items into types λ with probabilities $P_\Gamma(\lambda)$, so by the Poisson splitting property the marked counts are approximately independent Poisson with means $(\alpha_\ell/\ell)P_\Gamma(\lambda)$. This is exactly the family $(Z_{\ell,\lambda})$ used to define the limit vector $(A_1, \dots, A_B)$.

2.4 Compound Poisson vectors and their generating functions

We now isolate the general probabilistic structure underlying the limiting vector. Let $\mathcal{T}$ be a finite or countable type set; in our application $\mathcal{T} = \{(\ell, \lambda) : \ell \geq 1,\ \lambda \vdash k\}$. Let $(Z_\tau)_{\tau\in\mathcal{T}}$ be independent Poisson variables with means $(\nu_\tau)_{\tau\in\mathcal{T}}$. Given a deterministic map $\psi : \mathbb{Z}_{\geq 0}^{\mathcal{T}} \to \mathbb{Z}_{\geq 0}^{B}$, we call $\psi(Z)$ a *(multi-type) compound Poisson pushforward.* In our setting,

$$\psi\big((z_{\ell,\lambda})\big)_i = \sum_{\ell | i} \sum_{\lambda \vdash k} a_{i/\ell}(\lambda)\, z_{\ell,\lambda}, \qquad 1 \leq i \leq B,$$

so that $(A_1, \dots, A_B) = \psi((Z_{\ell,\lambda}))$.

The principal computational advantage is that exponential generating functions factor. For $x = (x_1, \dots, x_B) \in \mathbb{C}^B$, set

$$\mathcal{G}_B(x) := \mathbb{E}\Big[\prod_{i=1}^{B} x_i^{A_i}\Big].$$

Using the independence of $(Z_{\ell,\lambda})$ and the identity $\mathbb{E}[u^{\text{Poisson}(\nu)}] = \exp(\nu(u - 1))$, we obtain

$$\begin{aligned} \mathcal{G}_B(x) &= \prod_{\ell\geq 1} \prod_{\lambda\vdash k} \mathbb{E}\Big[\Big(\prod_{\substack{j\geq 1\\ j\ell\leq B}} x_{j\ell}^{a_j(\lambda)}\Big)^{Z_{\ell,\lambda}}\Big] \\ &= \exp\Big(\sum_{\ell\geq 1} \sum_{\lambda\vdash k} \nu_{\ell,\lambda}\Big(\prod_{\substack{j\geq 1\\ j\ell\leq B}} x_{j\ell}^{a_j(\lambda)} - 1\Big)\Big). \end{aligned} \tag{12}$$

Specializing to $\nu_{\ell,\lambda} = \alpha_\ell P_\Gamma(\lambda)/\ell$ and observing that for $1 \leq i \leq B$ only $\ell \leq B$ can contribute, (12) becomes

$$\mathcal{G}_B(x) = \exp\Big(\sum_{\ell=1}^{B} \frac{\alpha_\ell}{\ell}\Big(G_{\ell,B}(x) - 1\Big)\Big), \tag{13}$$

with $G_{\ell,B}$ as in (7). Thus the limiting vector $(A_1, \dots, A_B)$ is characterized by an explicit log-generating function of the same form as the Poissonized cycle-index asymptotics imposed on μ_n, after composition with the internal substitution $x \mapsto G_{\ell,B}(x)$.

We also record the corresponding factorial-moment formulas, which follow by differentiating (13). For example, for $1 \le i \le B$,

$$\begin{aligned}\mathbb{E}A_i &= \sum_{\ell | i} \sum_{\lambda \vdash k} a_{i/\ell}(\lambda) \, \frac{\alpha_\ell P_\Gamma(\lambda)}{\ell} \\ &= \sum_{\ell | i} \frac{\alpha_\ell}{\ell} \, \mathbb{E}_{\gamma \sim \mathrm{Unif}(\Gamma)} \big[a_{i/\ell}(\gamma) \big],\end{aligned} \tag{14}$$

and for joint factorial moments one obtains finite sums of products of the parameters α_ℓ/ℓ weighted by mixed factorial moments of the internal cycle counts $(a_j(\gamma))_{1 \le j \le k}$. The important point for us is structural: once we know that marked block-cycle counts converge to independent Poisson variables, the law of $(A_1, \dots, A_B)$ is fixed by the deterministic lift rule (11).

2.5 How the toolkit interfaces with Poissonization

The identities above explain why the Poissonized cycle-index hypothesis in the next section is the appropriate analytic input. Indeed, for fixed B, (9) expresses the truncated cycle index of σ as the truncated cycle index of h evaluated at the points $x_\ell = G_{\ell,B}(x)$:

$$\mathbb{E}\Big[\prod_{i=1}^{B} x_i^{a_i(\sigma)} \Big] = Z_{\mu_n}^{(\le B)} \big(G_{1,B}(x), \dots, G_{B,B}(x) \big).$$

Thus any uniform control on $\log Z_{\mu_n}^{(\le B)}$ (or on its Poissonized analogue) in a complex polydisc transfers directly to σ, provided we can bound the image of the map $x \mapsto (G_{1,B}(x), \dots, G_{B,B}(x))$ inside that polydisc. Since k and Γ are fixed, and each $G_{\ell,B}$ is a polynomial with nonnegative coefficients satisfying $G_{\ell,B}(1, \dots, 1) = 1$, such bounds are straightforward once the variables x_i are restricted to a fixed polydisc.

Finally, the marked viewpoint clarifies the limit mechanism. The cycle-index hypothesis will yield, after Poissonization and de-Poissonization, that for fixed L the unmarked block-cycle counts $(b_1, \dots, b_L)$ behave as if independent Poisson with means α_ℓ/ℓ. Independent marking by internal cycle type then produces independent Poisson variables $(Z_{\ell,\lambda})$ with means $\alpha_\ell P_\Gamma(\lambda)/\ell$, and the induced small-cycle counts of σ are obtained by the deterministic lifting rule (11). This is the only place where Γ enters: through its cycle index (3), equivalently through the mark distribution P_Γ.

In the next section we formalize the analytic hypothesis on (μ_n) in a way tailored to the Poissonized generating function F_L, and we discuss several equivalent formulations that will be convenient for verification in examples.

3 The Poissonized cycle-index hypothesis for (μ_n)

3.1 Poissonization and an ABT-type logarithmic condition

For fixed L we consider the Poissonized (more precisely, geometrically mixed) law obtained by first sampling a random size

$$N_t \sim \mathrm{Geom}(1-t), \qquad \mathbb{P}(N_t = n) = (1-t)t^n, \qquad n \geq 0,$$

and then sampling $h \sim \mu_{N_t}$. The corresponding truncated cycle index generating function is exactly

$$F_L(t; x_1, \dots, x_L) = \mathbb{E}\Big[\prod_{\ell=1}^{L} x_\ell^{b_\ell(h)}\Big],$$

with the understanding that the outer expectation includes the randomness of N_t. The point of this Poissonization is that it transforms coefficient extraction at a fixed n into analytic control of a single holomorphic function in the parameter t; this is the same mechanism that underlies the Arratia–Barbour–Tavaré "logarithmic combinatorial structures" regime.

The hypothesis imposed on (μ_n) is a logarithmic asymptotic for $\log F_L$ as $t \uparrow 1$, uniform on a complex polydisc in the variables $x_1, \dots, x_L$. Concretely, we assume there exist $(\alpha_\ell)_{\ell \geq 1} \subset [0, \infty)$, a radius $R > 1$, and a constant $C_0 < \infty$ such that, for each fixed L,

$$\log F_L(t; x_1, \dots, x_L) = \sum_{\ell=1}^{L} \frac{\alpha_\ell t^\ell}{\ell}(x_\ell - 1) + \mathcal{E}_L(t; x_1, \dots, x_L), \qquad |\mathcal{E}_L(t; \cdot)| \leq C_0 L^2 (1-t), \tag{15}$$

uniformly over $t \in (0,1)$ and $|x_\ell| \leq R$. The leading term in (15) is precisely the logarithm of the probability generating function of independent Poisson variables with means $(\alpha_\ell t^\ell / \ell)_{1 \leq \ell \leq L}$. Thus the hypothesis says: *under Poissonization, the small cycle counts behave as if they were asymptotically independent Poisson, with an error that is small on the natural scale* $1-t$.

Two remarks clarify what is and is not encoded by (15).

- The constants α_ℓ are allowed to be 0, which covers hard constraints such as the absence of ℓ-cycles. (For example, in the uniform derangement measure one expects $\alpha_1 = 0$ and $\alpha_\ell = 1$ for $\ell \geq 2$.)
- The complex-uniform bound on $\mathcal{E}_L$ is stronger than convergence of moments at real x_ℓ, but it is the natural input for analytic de-Poissonization and total variation estimates; we will use Cauchy estimates on derivatives and Tauberian bounds that require such a holomorphic control.

3.2 Equivalent formulations: Poisson approximation under Poissonization

We record several consequences of (15) that are, in practice, often taken as alternative formulations.

(i) Approximate independent Poisson laws for $(b_1, \dots, b_L)$ under N_t. Let $(Y_1^{(t)}, \dots, Y_L^{(t)})$ be independent Poisson with means $\mathbb{E}Y_\ell^{(t)} = \alpha_\ell t^\ell/\ell$. Then the multivariate probability generating function of $Y^{(t)}$ is

$$\mathbb{E}\Big[\prod_{\ell=1}^{L} x_\ell^{Y_\ell^{(t)}}\Big] = \exp\Big(\sum_{\ell=1}^{L} \frac{\alpha_\ell t^\ell}{\ell}(x_\ell - 1)\Big).$$

Comparing with (15) yields

$$\frac{F_L(t;x)}{\mathbb{E}\big[\prod_{\ell\le L} x_\ell^{Y_\ell^{(t)}}\big]} = \exp(\mathcal{E}_L(t;x)), \qquad |x_\ell| \le R. \tag{16}$$

In particular, $|\log(\cdot)| \le C_0 L^2(1-t)$ implies that, for t close to 1, the two generating functions are uniformly close on the polydisc. By standard inversion bounds for probability generating functions (for instance, bounding coefficients by Cauchy estimates on a circle $|x_\ell| = \rho \in (1, R)$ and summing), one obtains that for each fixed L,

$$\big\|\mathcal{L}\big((b_1, \dots, b_L) \text{ under } h \sim \mu_{N_t}\big) - \mathcal{L}\big(Y_1^{(t)}, \dots, Y_L^{(t)}\big)\big\|_{\mathrm{TV}} \;\le\; C(L, R, C_0)\,(1-t), \tag{17}$$

uniformly for t in a neighborhood of 1. We emphasize that (17) is an *intrinsically Poissonized* statement: the left-hand side refers to the mixture over n induced by N_t.

(ii) Asymptotics of cumulants and factorial moments. Differentiating (15) at $x \equiv 1$ gives uniform control of joint cumulants. For example,

$$\partial_{x_\ell} \log F_L(t;x)\big|_{x\equiv 1} = \mathbb{E}\big[b_\ell(h)\big] = \frac{\alpha_\ell t^\ell}{\ell} + O(L^2(1-t)).$$

More generally, mixed derivatives of $\log F_L$ at $x \equiv 1$ are mixed cumulants of $(b_1, \dots, b_L)$ under μ_{N_t}. Since the main term in (15) is linear in $(x_\ell - 1)$, all cumulants of order ≥ 2 arise only from $\mathcal{E}_L$. Cauchy bounds on derivatives on the polydisc $|x_\ell| \le R$ yield, for each fixed multi-index $(q_1, \dots, q_L)$ with $\sum q_\ell \ge 2$,

$$\Big|\kappa\big(\underbrace{b_1, \dots, b_1}_{q_1 \text{ times}},\ \dots,\ \underbrace{b_L, \dots, b_L}_{q_L \text{ times}}\big)\Big| \;\le\; C'(L, R, C_0)\,(1-t). \tag{18}$$

Equivalently, mixed factorial moments of $(b_1, \ldots, b_L)$ under μ_{N_t} approximately factor as in the independent Poisson case, with an error $O(1-t)$.

The same differentiation applied to F_L itself (rather than $\log F_L$) provides explicit approximations for factorial moments:

$$\mathbb{E}\Big[\prod_{\ell=1}^{L}(b_\ell)_{q_\ell}\Big] = \prod_{\ell=1}^{L}\Big(\frac{\alpha_\ell t^\ell}{\ell}\Big)^{q_\ell} + O_{L,q,R,C_0}(1-t), \tag{19}$$

uniformly for t near 1. Statements of the form (18)–(19) are often the most direct route in applications where one can control correlations among cycle counts but does not have a closed form for F_L.

(iii) De-Poissonized consequences at fixed n. The purpose of working with F_L rather than $Z^{(\le L)}_{\mu_n}$ is that analytic control in t can be turned into coefficient control in n via de-Poissonization. Under the bound $|\mathcal{E}_L| \le C_0 L^2(1-t)$, an analytic Tauberian lemma implies that the coefficient sequence

$$Z^{(\le L)}_{\mu_n}(x_1, \ldots, x_L) = \mathbb{E}_{h\sim\mu_n}\Big[\prod_{\ell=1}^{L} x_\ell^{b_\ell(h)}\Big]$$

is close, for large n, to the coefficient sequence one would obtain from the formal exponential model

$$\exp\Big(\sum_{\ell=1}^{L}\frac{\alpha_\ell}{\ell}(x_\ell - 1)\Big)$$

after conditioning on total size n. In the simplest case $\alpha_\ell \equiv 1$ this is the classical statement that the small cycles of a uniform random permutation are asymptotically independent Poisson$(1/\ell)$; the hypothesis (15) is designed to allow the same inference for general conjugacy-invariant families.

We will not commit here to a single de-Poissonization statement, since later arguments require a quantitative version tailored to total variation bounds after wreath-product composition. What matters structurally is that the error term in (15) is of order $(1-t)$, which corresponds to the natural coefficient scale $1/n$ after choosing $t = t_n \uparrow 1$ with $(1-t_n) \asymp 1/n$.

3.3 Discussion: how one verifies the hypothesis

We now summarize typical verification strategies for (15), emphasizing what requires explicit cycle-index information and what can be deduced from more probabilistic input.

(a) Direct cycle-index computation (explicit generating functions). In some families (μ_n) the truncated cycle index $Z_{\mu_n}^{(\le L)}$ is known in closed form (or admits a manageable recursion), and Poissonization can be analyzed explicitly. Examples include:

- $\mu_n = \mathrm{Unif}(S_n)$, where F_L can be computed from the exponential formula and one has $\alpha_\ell \equiv 1$.
- Ewens(θ) measures, where $\alpha_\ell \equiv \theta$.
- Measures obtained by imposing finitely many local constraints on cycle counts, such as forbidding ℓ-cycles for ℓ in a fixed set: this typically replaces α_ℓ by 0 on the forbidden set while leaving the remaining α_ℓ unchanged (up to lower-order perturbations), and the analytic control can be obtained by comparing the constrained and unconstrained cycle indices.

In these settings one often proves (15) by writing F_L as $\exp(\sum_{\ell \le L} c_\ell(t)(x_\ell - 1))$ times a remainder term coming from cycles longer than L, and then checking that the remainder contributes $O(L^2(1-t))$ uniformly in the polydisc.

(b) Conditioning relations (logarithmic assemblies in the ABT sense). A broad class of examples arises from a conditioning paradigm: one starts with independent random variables $(Y_\ell)_{\ell \ge 1}$ with $Y_\ell \sim \mathrm{Poisson}(\alpha_\ell/\ell)$ (or with a mild t-tilt giving $\mathrm{Poisson}(\alpha_\ell t^\ell/\ell)$), and then conditions on the size constraint $\sum_{\ell \ge 1} \ell Y_\ell = n$. If μ_n is defined as the law of the random cycle count vector under this conditioning, then (15) is essentially automatic: Poissonization removes the conditioning, and the cycle index factorizes. In this sense, (15) can be viewed as an analytic surrogate for the existence of an approximate conditioning relation.

In practice, one may not have an exact conditioning representation, but one can often show that the Radon–Nikodym derivative of μ_n with respect to a reference logarithmic measure (e.g. Ewens) depends weakly on the small cycle counts; such "smooth tilts" preserve the form (15) with at most a perturbation in α_ℓ and an acceptable remainder.

(c) Probabilistic control of correlations (moment/cumulant method). When no explicit cycle index is available, a useful route is to verify (15) indirectly by proving bounds of the form (18)–(19) for $h \sim \mu_{N_t}$. Indeed, if one can show that:

- the means satisfy $\mathbb{E} b_\ell(h) = \alpha_\ell t^\ell/\ell + O(1-t)$ for $\ell \le L$;
- mixed cumulants of order ≥ 2 are $O(1-t)$ uniformly for $\ell \le L$;

then Taylor expansion of $\log F_L(t;x)$ around $x \equiv 1$ and Cauchy bounds imply (15) on a sufficiently small polydisc, and a standard bootstrapping argument enlarges the domain to $|x_\ell| \le R$ for fixed $R > 1$ because F_L is a polynomial in each x_ℓ with nonnegative coefficients.

We stress that the hypothesis is *conjugacy-invariant* but otherwise flexible: it does not demand that μ_n be uniform on a group, or that μ_n have any spatial construction. What must be controlled is the dependence structure among the indicators of small cycles, as seen through the cumulants. This is often accessible via couplings to a reference measure, exchangeable pair methods, or explicit character bounds when μ_n is uniform on a subgroup with strong mixing properties in the conjugacy class space.

(d) What information is genuinely required. To check (15) for a given family (μ_n), one needs input at the level of *small-cycle statistics*:

- Identification of the parameters α_ℓ typically comes from the leading asymptotic $\ell\, \mathbb{E}_{h\sim\mu_n}[b_\ell(h)] \to \alpha_\ell$ as $n \to \infty$ (or from the analogous Poissonized limit as $t \uparrow 1$).
- The error term $O(L^2(1-t))$ is, in effect, a uniform quantitative bound on the aggregate influence of global constraints (such as the identity $\sum_\ell \ell b_\ell = n$) on the local cycle counts. Verifying such a bound usually requires either an explicit generating function or a robust probabilistic approximation argument.

In subsequent sections we will treat (15) as the basic analytic input and refrain from re-proving it in examples. The role of the present discussion is only to indicate that the hypothesis is not ad hoc: it is the natural quantitative form of "logarithmic structure" adapted to conjugacy-invariant measures on S_n, and it is tailored to remain stable under the wreath-product substitution identities developed earlier.

4 Exact Poissonized transfer through wreath composition

Fix $B \in \mathbb{N}$. In the Poissonized model we sample $N_t \sim \mathrm{Geom}(1-t)$, then $h \sim \mu_{N_t}$, and independently $\gamma_1, \ldots, \gamma_{N_t} \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}(\Gamma)$, and we set $\sigma = (\gamma_1, \ldots, \gamma_{N_t}; h) \in \Gamma^{N_t} \rtimes S_{N_t} \le S_{kN_t}$. Our aim in this section is to identify, at the level of probability generating functions, the exact transformation that carries the Poissonized small-cycle information of h into a marked Poisson model for the small cycles of σ.

4.1 Cycle structure along a single block-cycle

We begin with the deterministic description of the action of σ on the union of blocks belonging to one cycle of h. Let $c = (i_1\, i_2\, \cdots\, i_\ell)$ be an ℓ-cycle of h in S_n. Let B_{i_r} denote the r-th block (of size k) in this cycle. In the standard imprimitive action, σ maps B_{i_r} onto $B_{i_{r+1}}$ (indices modulo ℓ), applying the internal permutation γ_{i_r} before moving to the next block.

Define the *cycle product* along c by

$$\tau(c) := \gamma_{i_\ell}\gamma_{i_{\ell-1}}\cdots\gamma_{i_1} \in \Gamma.$$

This element controls the return map on internal coordinates after one traversal of the block-cycle.

Lemma 4.1 (Inflation of internal cycles)**.** *Let c be an ℓ-cycle of h, and let $\tau = \tau(c) \in \Gamma$ be the corresponding cycle product. Then the restriction of σ to the $k\ell$ points in $\bigcup_{r=1}^{\ell} B_{i_r}$ has cycle lengths equal to ℓ times the cycle lengths of τ. Equivalently, for each $j \geq 1$, the number of $(j\ell)$-cycles of σ supported on these $k\ell$ points is exactly $a_j(\tau)$.*

Proof. Label points in the union of blocks by pairs (r, u) where $r \in \{1, \ldots, \ell\}$ indexes the block B_{i_r} and $u \in [k]$ is the internal coordinate in that block. By the wreath action, σ maps

$$(r, u) \longmapsto (r + 1, \gamma_{i_r}(u)),$$

with the block index taken modulo ℓ. Iterating ℓ times gives

$$(r, u) \overset{\sigma^\ell}{\longmapsto} (r, \gamma_{i_\ell} \cdots \gamma_{i_1}(u)) = (r, \tau(u)).$$

Thus σ^ℓ acts on each fiber $\{r\} \times [k]$ as τ, and the orbit of (r, u) under σ returns to the same block after ℓ steps while applying τ once to the internal coordinate. If u lies in a j-cycle of τ, then $\sigma^{\ell j}(r, u) = (r, u)$ and no smaller positive multiple of ℓ achieves this, so the cycle length of (r, u) under σ is ℓj. Moreover, each j-cycle of τ yields exactly one (ℓj)-cycle of σ (it threads once through the ℓ blocks), so the count of (ℓj)-cycles is $a_j(\tau)$. $\square$

A key simplification is that, under our sampling scheme, the random mark $\tau(c)$ is uniform on Γ and does not depend on ℓ.

Lemma 4.2 (Uniformity of cycle products)**.** *Let $\gamma_1, \ldots, \gamma_\ell$ be independent uniform random elements of a finite group Γ. Then the product $\gamma_\ell \cdots \gamma_1$ is uniform on Γ. In particular, for every $\ell \geq 1$, for each $\lambda \vdash k$,*

$$\mathbb{P}\big(\mathrm{cyc}(\tau(c)) = \lambda\big) = P_\Gamma(\lambda),$$

where $\mathrm{cyc}(\cdot)$ *denotes cycle type in* S_k*.*

Proof. For any fixed $g \in \Gamma$,

$$\mathbb{P}(\gamma_\ell \cdots \gamma_1 = g) = \sum_{h \in \Gamma} \mathbb{P}(\gamma_\ell = h)\, \mathbb{P}(\gamma_{\ell-1} \cdots \gamma_1 = h^{-1}g) = \frac{1}{|\Gamma|} \sum_{h \in \Gamma} \mathbb{P}(\gamma_{\ell-1} \cdots \gamma_1 = h^{-1}g).$$

By induction on ℓ (with $\ell = 1$ trivial), the inner probability is $1/|\Gamma|$ for each h, hence the sum equals $1/|\Gamma|$. □

4.2 An exact substitution identity for Poissonized generating functions

We now encode Lemma 4.1 at the level of truncated probability generating functions for cycle counts. For $B \in \mathbb{N}$ and complex variables $z_1, \ldots, z_B$, define the Poissonized truncated cycle generating function of σ by

$$G_B(t; z_1, \ldots, z_B) := \sum_{n \geq 0} (1-t) t^n\, \mathbb{E}\Big[\prod_{i=1}^{B} z_i^{a_i(\sigma)} \;\Big|\; N_t = n \Big],$$

where the expectation is over $h \sim \mu_n$ and $\gamma_1, \ldots, \gamma_n \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}(\Gamma)$.

For each $\ell \geq 1$ we introduce the *internal mark generating polynomial*

$$\Theta_\ell(z_1, \ldots, z_B) := \sum_{\lambda \vdash k} P_\Gamma(\lambda) \prod_{j=1}^{k} z_{j\ell}^{a_j(\lambda)}, \tag{20}$$

with the convention $z_m := 1$ for $m > B$. Thus Θ_ℓ is the expected contribution, in the z-weights, of an ℓ-cycle of h after random internal permutations are applied and the induced cycle lengths in S_{kn} are recorded up to level B.

Proposition 4.3 (Exact wreath-product substitution)**.** *For every $B \in \mathbb{N}$ and every $t \in (0,1)$,*

$$G_B(t; z_1, \ldots, z_B) = F_B\big(t; \Theta_1(z), \ldots, \Theta_B(z)\big), \tag{21}$$

where F_B is the Poissonized truncated cycle index generating function of (μ_n).

Proof. Condition on $N_t = n$ and on $h \in S_n$. The permutation h decomposes into disjoint cycles; cycles of h act on disjoint collections of blocks, hence on disjoint subsets of $[kn]$. Therefore, conditional on h, the random variables counting the contributions from distinct cycles of h multiply in the generating function.

Fix an ℓ-cycle c of h. By Lemma 4.1, the contribution of c to the exponent vector $(a_1(\sigma), \ldots, a_B(\sigma))$ is determined by $\tau(c) \in \Gamma$: it adds $a_j(\tau(c))$ to $a_{\ell j}(\sigma)$ for each $j \geq 1$. Consequently, conditional on h, the z-weight contributed by c is

$$\prod_{i=1}^{B} z_i^{\Delta a_i(c)} = \prod_{j=1}^{k} z_{\ell j}^{a_j(\tau(c))},$$

again with $z_{\ell j} = 1$ if $\ell j > B$.

Now average over the internal permutations. The cycles of h use disjoint sets of indices, so the corresponding products $\tau(c)$ are independent; by Lemma 4.2, each $\tau(c)$ is uniform on Γ, hence has cycle type distribution $P_\Gamma(\cdot)$. Therefore the expected weight of a single ℓ-cycle is precisely $\Theta_\ell(z)$ defined in (20). If h has $b_\ell(h)$ cycles of length ℓ, then the conditional expectation over internal permutations yields

$$\mathbb{E}\Big[\prod_{i=1}^{B} z_i^{a_i(\sigma)} \;\Big|\; h\Big] = \prod_{\ell\geq 1} \Theta_\ell(z)^{b_\ell(h)}.$$

Since only $\ell \leq B$ can affect $a_1, \ldots, a_B$ (because all induced cycle lengths are multiples of ℓ), we may truncate to $\ell \leq B$ without changing the value. Finally, averaging over $h \sim \mu_n$ and then mixing over n with weights $(1-t)t^n$ gives

$$G_B(t;z) = \sum_{n\geq 0}(1-t)t^n\, \mathbb{E}_{h\sim\mu_n}\Big[\prod_{\ell=1}^{B} \Theta_\ell(z)^{b_\ell(h)}\Big] = F_B\big(t; \Theta_1(z), \ldots, \Theta_B(z)\big),$$

which is (21). □

Proposition 4.3 is the fundamental transfer identity: it shows that *all* information needed about the small cycles of σ under Poissonization is obtained from F_B by a deterministic substitution map $x_\ell \mapsto \Theta_\ell(z)$ depending only on (k, Γ).

4.3 Marked Poisson structure after Poissonization

We now combine (21) with the Poissonized cycle-index hypothesis for F_B. To apply the hypothesis uniformly on a polydisc, we choose $r > 1$ so that $r^k \leq R$. If $|z_i| \leq r$ for $1 \leq i \leq B$, then, since Θ_ℓ has nonnegative coefficients and total degree at most k,

$$|\Theta_\ell(z)| \leq \sum_{\lambda\vdash k} P_\Gamma(\lambda) \prod_{j=1}^{k} |z_{j\ell}|^{a_j(\lambda)} \leq \sum_{\lambda\vdash k} P_\Gamma(\lambda)\, r^{\sum_j a_j(\lambda)} \leq r^k \leq R,$$

so the hypothesis applies to $F_B(t; \Theta_1(z), \ldots, \Theta_B(z))$.

Proposition 4.4 (Poissonized marked-Poisson approximation)**.** *Fix $B \in \mathbb{N}$ and let $|z_i| \leq r$ with $r^k \leq R$. Then, uniformly over such z and over $t \in (0,1)$,*

$$\log G_B(t;z) = \sum_{\ell=1}^{B} \frac{\alpha_\ell t^\ell}{\ell}\big(\Theta_\ell(z) - 1\big) + \widetilde{\mathcal{E}}_B(t;z), \qquad |\widetilde{\mathcal{E}}_B(t;z)| \leq C_0 B^2 (1-t). \tag{22}$$

Moreover, the leading term admits the expansion

$$\sum_{\ell=1}^{B} \frac{\alpha_\ell t^\ell}{\ell}\big(\Theta_\ell(z) - 1\big) = \sum_{\ell=1}^{B} \sum_{\lambda \vdash k} \frac{\alpha_\ell t^\ell}{\ell} P_\Gamma(\lambda) \left(\prod_{j=1}^{k} z_{j\ell}^{a_j(\lambda)} - 1 \right). \tag{23}$$

Proof. Insert (21) into the hypothesis (15) with $L = B$ and $x_\ell = \Theta_\ell(z)$. The bound $|\Theta_\ell(z)| \le R$ shown above yields (22) with $\widetilde{\mathcal{E}}_B(t; z) = \mathcal{E}_B(t; \Theta_1(z), \ldots, \Theta_B(z))$. The identity (23) is immediate from the definition of Θ_ℓ. $\square$

The representation (23) is the logarithm of the joint probability generating function of a *marked Poisson process* indexed by (ℓ, λ), and the wreath-product map is exactly the pushforward described by Lemma 4.1. Concretely, define independent random variables

$$Z_{\ell,\lambda}^{(t)} \sim \text{Poisson}\left(\frac{\alpha_\ell t^\ell}{\ell} P_\Gamma(\lambda) \right), \qquad 1 \le \ell \le B,\ \lambda \vdash k,$$

and define a random vector $(A_1^{(t)}, \ldots, A_B^{(t)})$ by

$$A_i^{(t)} := \sum_{\ell=1}^{B} \sum_{\lambda \vdash k} \sum_{j \ge 1:\ j\ell = i} a_j(\lambda)\, Z_{\ell,\lambda}^{(t)}, \qquad 1 \le i \le B. \tag{24}$$

Then the multivariate pgf of $(A_1^{(t)}, \ldots, A_B^{(t)})$ is

$$\begin{aligned}
\mathbb{E}\Big[\prod_{i=1}^{B} z_i^{A_i^{(t)}} \Big] &= \prod_{\ell=1}^{B} \prod_{\lambda \vdash k} \exp\left(\frac{\alpha_\ell t^\ell}{\ell} P_\Gamma(\lambda) \left(\prod_{j=1}^{k} z_{j\ell}^{a_j(\lambda)} - 1 \right) \right) \\
&= \exp\left(\sum_{\ell=1}^{B} \sum_{\lambda \vdash k} \frac{\alpha_\ell t^\ell}{\ell} P_\Gamma(\lambda) \left(\prod_{j=1}^{k} z_{j\ell}^{a_j(\lambda)} - 1 \right) \right),
\end{aligned}$$

which matches the main term in (23). In this sense, Proposition 4.4 shows that, *under Poissonization*, the cycle counts of σ up to level B behave as if they were obtained by:

- sampling an (approximately) independent Poisson number of ℓ-cycles in the block permutation, with mean $\alpha_\ell t^\ell / \ell$;
- independently marking each such ℓ-cycle by a cycle type $\lambda \vdash k$ with law P_Γ (equivalently, by a uniform internal product in Γ);
- inflating the internal j-cycles to cycles of length $j\ell$ in S_{kN_t}.

The error term $\widetilde{\mathcal{E}}_B(t; z)$ is inherited *without amplification* from the hypothesis on F_B: the substitution $x_\ell = \Theta_\ell(z)$ preserves the $O(B^2(1-t))$ scale.

At the purely Poissonized level (i.e. with N_t random), one can pass from (22) to quantitative approximation of distributions by the same generating-function inversion used for (17): since both $G_B(t;z)$ and the pgf of $A^{(t)}$ are holomorphic on $|z_i| \le r$ and have nonnegative coefficients, Cauchy bounds on the torus $|z_i| = \rho \in (1, r)$ show that the coefficient arrays (hence the laws) are close, with discrepancy of order $B^2(1-t)$ uniformly for t near 1. We do not isolate this Poissonized total variation statement here, because the subsequent de-Poissonization step will require a version that is stable as $t = t_n \uparrow 1$ with $(1 - t_n) \asymp 1/n$ and that interacts well with truncation in B.

Finally, letting $t \uparrow 1$ in (24) yields the limiting marked Poisson model: if $Z_{\ell,\lambda} \sim \text{Poisson}(\alpha_\ell P_\Gamma(\lambda)/\ell)$ are independent for all $\ell \ge 1$ and $\lambda \vdash k$, then the limiting cycle count vector $(A_1, \dots, A_B)$ defined by

$$A_i := \sum_{\ell \ge 1} \sum_{\lambda \vdash k} \sum_{j \ge 1:\ j\ell = i} a_j(\lambda)\, Z_{\ell,\lambda}, \qquad 1 \le i \le B,$$

is exactly the $t \uparrow 1$ limit of $(A_1^{(t)}, \dots, A_B^{(t)})$ at the level of finite-dimensional distributions. Thus the wreath-product composition does not merely preserve the "logarithmic" form of the Poissonized cycle index: it refines it into an explicit marked Poisson process, with marks taking values in partitions of k according to the cycle-type law induced by Γ.

The remaining task is to convert these Poissonized conclusions into statements at fixed size n, uniformly in the truncation parameter B at the level required for total variation bounds. This is the role of the de-Poissonization argument developed next.

5 De-Poissonization with uniform error

For fixed $B \in \mathbb{N}$, let

$$g_{n,B}(z_1, \dots, z_B) := \mathbb{E}\Big[\prod_{i=1}^{B} z_i^{a_i(\sigma)} \;\Big|\; N_t = n\Big] = \mathbb{E}\Big[\prod_{i=1}^{B} z_i^{a_i(\sigma_n)}\Big],$$

where $\sigma_n = (\gamma_1, \dots, \gamma_n; h) \in S_{kn}$ is the wreath-product element at deterministic size n. Thus

$$G_B(t;z) = \sum_{n \ge 0} (1-t)t^n\, g_{n,B}(z)$$

is the geometric (Poissonized) transform of the sequence $(g_{n,B})_{n \ge 0}$.

On the other hand, the marked-Poisson model from the end of the previous section defines, for each $t \in (0,1)$, a random vector $A^{(t)} = (A_1^{(t)}, \dots, A_B^{(t)})$ with probability generating function

$$\widehat{G}_B(t;z) := \mathbb{E}\Big[\prod_{i=1}^{B} z_i^{A_i^{(t)}}\Big] = \exp\left(\sum_{\ell=1}^{B} \sum_{\lambda \vdash k} \frac{\alpha_\ell t^\ell}{\ell} P_\Gamma(\lambda) \left(\prod_{j=1}^{k} z_{j\ell}^{a_j(\lambda)} - 1\right)\right),$$

and the limiting vector $A = (A_1, \dots, A_B)$ has pgf

$$\widehat{g}_B(z) := \mathbb{E}\Big[\prod_{i=1}^{B} z_i^{A_i}\Big] = \exp\left(\sum_{\ell=1}^{B}\sum_{\lambda\vdash k}\frac{\alpha_\ell}{\ell}P_\Gamma(\lambda)\left(\prod_{j=1}^{k} z_{j\ell}^{a_j(\lambda)} - 1\right)\right), \tag{25}$$

which is $\widehat{G}_B(t;z)$ evaluated at $t = 1$.

The goal of this section is to pass from the Poissonized approximation of $G_B(t;z)$ (uniform in z on a polydisc) to approximation of the fixed-size pgf $g_{n,B}(z)$, and then to translate this into total variation bounds for the corresponding laws on $\mathbb{Z}_{\geq 0}^B$.

5.1 From logarithmic error to an additive Poissonized error

We first record a convenient consequence of Proposition 4.4. Fix $r > 1$ with $r^k \leq R$ and then choose $\rho \in (1, r)$. For $|z_i| \leq \rho$, the bound (22) gives

$$\log G_B(t;z) = \log \widehat{G}_B(t;z) + \widetilde{\mathcal{E}}_B(t;z), \qquad |\widetilde{\mathcal{E}}_B(t;z)| \leq C_0 B^2 (1-t),$$

where we view $\log \widehat{G}_B(t;z)$ as the explicit main term in (23). Exponentiating yields a uniform multiplicative control:

$$\frac{G_B(t;z)}{\widehat{G}_B(t;z)} = \exp(\widetilde{\mathcal{E}}_B(t;z)), \qquad \left|\frac{G_B(t;z)}{\widehat{G}_B(t;z)} - 1\right| \leq e^{C_0 B^2 (1-t)} - 1. \tag{26}$$

In particular, for t sufficiently close to 1 (depending on B) the right-hand side is $O(B^2(1-t))$, uniformly for $|z_i| \leq \rho$.

We also need a uniform bound controlling the approach $\widehat{G}_B(t;z) \to \widehat{g}_B(z)$ as $t \uparrow 1$. Since $|t^\ell - 1| \leq \ell|1-t|/(1-|1-t|)$ for t near 1, and since $\Theta_\ell(z)$ has degree at most k and $|\Theta_\ell(z)| \leq r^k$ on $|z_i| \leq r$, we have

$$\big|\log \widehat{G}_B(t;z) - \log \widehat{g}_B(z)\big| \leq |1-t| \cdot C_1(B,\rho) \sum_{\ell=1}^{B} \alpha_\ell, \tag{27}$$

for a constant $C_1(B,\rho)$ depending only on the fixed truncation and the polydisc radius. (When (α_ℓ) is bounded, the sum is $O(B)$.) Combining (26) and (27) we obtain, for t close to 1 and $|z_i| \leq \rho$,

$$\big|G_B(t;z) - \widehat{g}_B(z)\big| \leq C_2(B,\rho)\,(1-t), \tag{28}$$

with an explicit C_2 growing at most polynomially in B (under mild control of (α_ℓ)). The de-Poissonization step amounts to transferring (28) from the Abel means $G_B(t;\cdot)$ to the coefficients $g_{n,B}(\cdot)$.

5.2 A quantitative analytic Tauberian lemma

We use an Abelian-to-coefficient estimate adapted to the geometric transform $(1-t)\sum_n t^n(\cdot)$. We state it in a form suited to families of holomorphic functions in auxiliary variables (here the z_i).

Lemma 5.1 (Analytic de-Poissonization for geometric means)**.** *Fix $\rho > 1$. Let $(u_n(\,\cdot\,))_{n\geq 0}$ be a sequence of functions holomorphic on the polydisc $\{z \in \mathbb{C}^B : |z_i| \leq \rho\}$ and uniformly bounded there by M_0. Define*

$$U(t;z) := \sum_{n\geq 0}(1-t)t^n u_n(z), \qquad |t| < 1.$$

Assume there exist a function $u_\infty(z)$ holomorphic on $|z_i| \leq \rho$ and constants $\eta \in (0,1)$, $M_1 < \infty$ such that $U(t;z)$ extends holomorphically to the Stolz region

$$\mathcal{S}_\eta := \{t \in \mathbb{C} : |t| < 1+\eta,\ |1-t| < \eta(1-|t|)\},$$

and satisfies the uniform bound

$$\sup_{t\in\mathcal{S}_\eta}\ \sup_{|z_i|\leq\rho} |U(t;z) - u_\infty(z)| \leq M_1\,|1-t|. \tag{29}$$

Then there is a constant $C = C(\eta)$ such that for all $n \geq 1$,

$$\sup_{|z_i|\leq\rho} |u_n(z) - u_\infty(z)| \leq \frac{CM_1}{n}. \tag{30}$$

Proof. Write $H(t;z) := \sum_{n\geq 0} u_n(z)t^n = U(t;z)/(1-t)$, so that $u_n(z) = [t^n]H(t;z)$. By Cauchy's coefficient formula, for any $r \in (0,1)$,

$$u_n(z) = \frac{1}{2\pi i}\int_{|t|=r} \frac{H(t;z)}{t^{n+1}}\,dt.$$

We choose $r = r_n := 1 - \frac{1}{n}$ and deform the contour in the standard way for de-Poissonization (see, e.g., (**?**, Ch. VI)): the circle $|t| = r_n$ is split into an arc contained in $\mathcal{S}_\eta$ where (29) applies, and a complementary arc at positive distance from $t = 1$ where the factor $t^{-(n+1)}$ yields exponential decay in n. On the Stolz arc we insert

$$H(t;z) = \frac{u_\infty(z)}{1-t} + \frac{U(t;z) - u_\infty(z)}{1-t}$$

and estimate the second term using (29) as $O(M_1)$ uniformly. The first term contributes exactly $u_\infty(z)$ after coefficient extraction (since $[t^n](1-t)^{-1} = 1$). The complementary arc contributes $O(r_n^{-n})$ times a uniform bound on H, which is $O(1)$ because U is bounded and $|1-t|$ is bounded below there. Since $r_n^{-n} = (1-\frac{1}{n})^{-n} = O(1)$ and the arc length is $O(1/n)$, the total contribution of the error term is $O(M_1/n)$, uniformly for $|z_i| \leq \rho$. Collecting these estimates yields (30). □

In applications, the analyticity condition in $\mathcal{S}_\eta$ is used only to justify the contour deformation and to control $U(t;z)$ on non-real t near 1. For our $G_B(t;z)$, holomorphy for $|t|<1$ is automatic from the defining series; the additional extension to a Stolz region is obtained from the assumed uniform control of $\log F_L$ on a complex polydisc together with the substitution $x_\ell = \Theta_\ell(z)$ (and the fact that Θ_ℓ is a polynomial), by an application of the maximum principle to the error term.

5.3 Application to $g_{n,B}$ and convergence of generating functions

We apply Lemma 5.1 with $u_n(z) = g_{n,B}(z)$, $U(t;z) = G_B(t;z)$, and $u_\infty(z) = \widehat{g}_B(z)$ from (25). The bound (28) provides (29) with $M_1 = C_2(B,\rho)$, hence

$$\sup_{|z_i|\le\rho} \left|g_{n,B}(z) - \widehat{g}_B(z)\right| \le \frac{C(\eta)\, C_2(B,\rho)}{n}. \tag{31}$$

In particular, for each fixed B, $g_{n,B}(z) \to \widehat{g}_B(z)$ uniformly on $|z_i| \le \rho$, and therefore the joint law of $(a_1(\sigma_n),\ldots,a_B(\sigma_n))$ converges to that of $(A_1,\ldots,A_B)$.

5.4 From uniform pgf control to total variation

It remains to convert (31) into a quantitative bound in total variation. For a probability measure ν on $\mathbb{Z}_{\ge 0}^B$, write its pgf as

$$\mathcal{P}_\nu(z) = \sum_{m\in\mathbb{Z}_{\ge 0}^B} \nu(m)\, z_1^{m_1}\cdots z_B^{m_B}.$$

If $\mathcal{P}_\nu$ and $\mathcal{P}_{\nu'}$ are holomorphic on $|z_i| \le \rho$ for some $\rho > 1$, then multidimensional Cauchy estimates on the torus $|z_i| = \rho$ give

$$|\nu(m) - \nu'(m)| \le \rho^{-\|m\|_1} \sup_{|z_i|=\rho} |\mathcal{P}_\nu(z) - \mathcal{P}_{\nu'}(z)|.$$

Summing over $m \in \mathbb{Z}_{\ge 0}^B$ yields the geometric series bound

$$\|\nu - \nu'\|_{\mathrm{TV}} \le \frac{1}{2}\left(\frac{\rho}{\rho-1}\right)^B \sup_{|z_i|=\rho} \left|\mathcal{P}_\nu(z) - \mathcal{P}_{\nu'}(z)\right|. \tag{32}$$

Applying (32) with $\nu = \mathcal{L}(a_1(\sigma_n),\ldots,a_B(\sigma_n))$ and $\nu' = \mathcal{L}(A_1,\ldots,A_B)$, and using (31) evaluated on $|z_i| = \rho$, we obtain

$$\left\|\mathcal{L}(a_1,\ldots,a_B) - \mathcal{L}(A_1,\ldots,A_B)\right\|_{\mathrm{TV}} \le \frac{C_3(B,\rho)}{n}, \tag{33}$$

where $C_3(B,\rho)$ is explicit in terms of $C_2(B,\rho)$ and the factor $(\rho/(\rho-1))^B$.

For later use (e.g. allowing B to grow moderately with n), it is convenient to isolate the dependence on n in a single parameter. We therefore define a de-Poissonization error term

$$\varepsilon_n := \frac{1}{n},$$

and absorb all dependence on (k, Γ) and on the analyticity constants (including ρ and the constants in the Poissonized cycle-index hypothesis) into a single multiplicative constant. With this notation, (33) implies the quantitative statement announced in the introduction: for each fixed B,

$$\big\|\mathcal{L}(a_1, \ldots, a_B) - \mathcal{L}(A_1, \ldots, A_B)\big\|_{\mathrm{TV}} \le C\, B\, \varepsilon_n,$$

after enlarging C if necessary (since B is fixed and all polynomial dependence on B may be upper bounded by a constant multiple of B).

We emphasize that the argument is stable under truncation: the same contour method applies uniformly in z on a polydisc, and the only input from the underlying measures (μ_n) is the Poissonized hypothesis on $\log F_L$ together with the wreath substitution. This completes the passage from Poissonized marked-Poisson structure to fixed-size convergence with an explicit total variation estimate. The next section specializes these general formulas in concrete examples of (μ_n) and computes the resulting limiting cycle counts (A_i) for selected (k, Γ).

6 Examples and computations

In this section we record the coefficients $(\alpha_\ell)_{\ell \ge 1}$ for several standard conjugacy-invariant measures (μ_n), and we spell out the corresponding limiting cycle counts $(A_i)_{i \ge 1}$ in a few small block sizes k. Throughout, once (α_ℓ) is identified for (μ_n), the marked-Poisson variables

$$Z_{\ell,\lambda} \;\sim\; \mathrm{Poisson}\Big(\frac{\alpha_\ell}{\ell} P_\Gamma(\lambda)\Big), \qquad \ell \ge 1,\ \lambda \vdash k,$$

are independent, and the limiting cycle counts are given by

$$A_i = \sum_{\ell \ge 1} \sum_{\lambda \vdash k} \sum_{j \ge 1:\ j\ell = i} a_j(\lambda)\, Z_{\ell,\lambda}.$$

We will often use the equivalent "divisor form"

$$A_i = \sum_{d \mid i} \sum_{\lambda \vdash k} a_{i/d}(\lambda)\, Z_{d,\lambda}, \tag{34}$$

since $j\ell = i$ is the same as $\ell = d$ with $d \mid i$ and $j = i/d$. In particular,

$$\mathbb{E}[A_i] = \sum_{d \mid i} \frac{\alpha_d}{d} \sum_{\lambda \vdash k} a_{i/d}(\lambda)\, P_\Gamma(\lambda). \tag{35}$$

6.1 Identifying (α_ℓ) for common measures on S_n

(i) Uniform measure on S_n. For $\mu_n = \mathrm{Unif}(S_n)$, the Poissonized cycle-index hypothesis holds with

$$\alpha_\ell \equiv 1, \qquad \ell \geq 1.$$

Equivalently, for each fixed L, the Poissonized truncated cycle counts behave like independent Poisson variables with means t^ℓ/ℓ, and hence the limiting (de-Poissonized) small-cycle intensities are $1/\ell$.

(ii) Uniform measure on A_n. For $\mu_n = \mathrm{Unif}(A_n)$, the same Poissonized hypothesis holds with

$$\alpha_\ell \equiv 1, \qquad \ell \geq 1.$$

At the level of finitely many cycle counts, the even-parity constraint is carried predominantly by the long cycles; consequently the finite-dimensional limits for $(b_1, \dots, b_L)$ coincide with those under $\mathrm{Unif}(S_n)$. In particular, our limiting vector $(A_1, \dots, A_B)$ is the same as in the uniform S_n case for each fixed B.

(iii) Ewens(θ) measure. Let μ_n be the Ewens measure with parameter $\theta > 0$, i.e. $\mu_n(h) \propto \theta^{\#\mathrm{cycles}(h)}$. Then the Poissonized hypothesis holds with

$$\alpha_\ell \equiv \theta, \qquad \ell \geq 1,$$

so that the limiting block-cycle counts have intensities θ/ℓ, and hence

$$Z_{\ell,\lambda} \sim \mathrm{Poisson}\left(\frac{\theta}{\ell} P_\Gamma(\lambda)\right).$$

(iv) Derangements. Let μ_n be uniform on the derangements in S_n. This is conjugacy-invariant and forces $b_1(h) = 0$ almost surely. The Poissonized hypothesis holds with

$$\alpha_1 = 0, \qquad \alpha_\ell = 1 \ \ (\ell \geq 2).$$

Accordingly, all marked variables $Z_{1,\lambda}$ vanish almost surely, and only $\ell \geq 2$ block-cycles contribute to (A_i).

(v) Restricted cycle-length measures. A broad class of logarithmic examples is obtained by forbidding (or down-weighting) certain cycle lengths. For instance, fix parameters $(\theta_\ell)_{\ell \geq 1}$ with $\theta_\ell \geq 0$ and consider generalized Ewens-type weights (restricted to those $h \in S_n$ for which $\theta_\ell > 0$ whenever $b_\ell(h) > 0$), so that heuristically the small cycles behave like independent Poisson with means θ_ℓ/ℓ. In our notation this corresponds to

$$\alpha_\ell = \theta_\ell, \qquad \ell \geq 1,$$

with the special case of forbidding a length ℓ given by $\theta_\ell = 0$ (so $Z_{\ell,\lambda} \equiv 0$ for all λ). The derangement model is exactly the choice $\theta_1 = 0$ and $\theta_\ell = 1$ for $\ell \geq 2$.

6.2 Selected small block sizes

We now compute (A_i) more explicitly for a few (k, Γ). The input from Γ is the cycle-type distribution $P_\Gamma(\lambda)$; once this is known, (34) gives an explicit representation of each A_i as a finite sum of independent Poisson variables (with integer coefficients coming from $a_j(\lambda)$).

6.2.1 Trivial internal action: $\Gamma = \{e\} \leq S_k$

When Γ is trivial, the only cycle type that occurs is $\lambda = (1^k)$, for which $a_1(\lambda) = k$ and $a_j(\lambda) = 0$ for $j \geq 2$. Thus $P_\Gamma((1^k)) = 1$ and

$$Z_{\ell,(1^k)} \sim \text{Poisson}\Big(\frac{\alpha_\ell}{\ell}\Big), \qquad A_i = k\, Z_{i,(1^k)}.$$

In particular, $(A_i)_{i\geq 1}$ is simply a rescaled version of the limiting cycle counts for h, with the deterministic multiplicity factor k reflecting that σ consists of k disjoint copies of the block permutation when there is no internal motion.

Under $\mu_n = \text{Unif}(S_n)$, this yields $A_i = k\,\text{Poisson}(1/i)$ (supported on multiples of k). Under Ewens(θ), we obtain $A_i = k\,\text{Poisson}(\theta/i)$. Under derangements, we have $A_1 \equiv 0$ and $A_i = k\,\text{Poisson}(1/i)$ for $i \geq 2$.

6.2.2 Case $k = 2$

For $k = 2$ there are two partitions: $\lambda = (1, 1)$ and $\lambda = (2)$. We record

$$a_1(1,1) = 2, \quad a_2(1,1) = 0; \qquad a_1(2) = 0, \quad a_2(2) = 1.$$

Thus a block-cycle of length ℓ marked by $(1, 1)$ produces two ℓ-cycles in $[2n]$, while a block-cycle of length ℓ marked by (2) produces one (2ℓ)-cycle.

(a) $\Gamma = S_2$ (equivalently $\Gamma = C_2$). A uniform element of S_2 is the identity with probability 1/2 and the transposition with probability 1/2, hence

$$P_{S_2}(1,1) = \frac{1}{2}, \qquad P_{S_2}(2) = \frac{1}{2}.$$

Define independent Poisson variables

$$X_\ell := Z_{\ell,(1,1)} \sim \text{Poisson}\Big(\frac{\alpha_\ell}{2\ell}\Big), \qquad Y_\ell := Z_{\ell,(2)} \sim \text{Poisson}\Big(\frac{\alpha_\ell}{2\ell}\Big).$$

Then (34) becomes

$$A_i = 2X_i + \mathbf{1}_{2|i}\, Y_{i/2}, \qquad i \geq 1, \tag{36}$$

where $\mathbf{1}_{2|i}$ is the indicator of i even. In particular,

$$A_1 = 2X_1, \qquad A_2 = 2X_2 + Y_1, \qquad A_3 = 2X_3, \qquad A_4 = 2X_4 + Y_2,$$

and so on.

For $\mu_n = \mathrm{Unif}(S_n)$ (or $\mathrm{Unif}(A_n)$) we have $\alpha_\ell \equiv 1$, so $X_\ell, Y_\ell \sim \mathrm{Poisson}(1/(2\ell))$ and

$$\mathbb{E}[A_i] = \frac{1}{i} + \mathbf{1}_{2|i}\frac{1}{i} = \frac{1 + \mathbf{1}_{2|i}}{i}.$$

For Ewens(θ), $X_\ell, Y_\ell \sim \mathrm{Poisson}(\theta/(2\ell))$ and $\mathbb{E}[A_i] = \theta(1 + \mathbf{1}_{2|i})/i$. For derangements, $\alpha_1 = 0$ forces $X_1 = Y_1 = 0$ a.s., hence $A_1 \equiv 0$ and $A_2 = 2X_2$ has no contribution from Y_1 (reflecting the absence of fixed blocks).

(b) $\Gamma = \{e\}$ inside S_2. This is the trivial case discussed above: $A_i = 2\,\mathrm{Poisson}(\alpha_i/i)$.

6.2.3 Case $k = 3$

For $k = 3$, the partitions are $\lambda = (1, 1, 1)$, $\lambda = (2, 1)$, and $\lambda = (3)$. The part-counts are

λ	$a_1(\lambda)$	$a_2(\lambda)$	$a_3(\lambda)$
$(1,1,1)$	3	0	0
$(2,1)$	1	1	0
(3)	0	0	1

Accordingly, a block-cycle of length ℓ marked by $(1, 1, 1)$ yields three ℓ-cycles; a mark $(2, 1)$ yields one ℓ-cycle and one (2ℓ)-cycle; and a mark (3) yields one (3ℓ)-cycle.

(a) $\Gamma = S_3$. A uniform element of S_3 is of type $(1, 1, 1)$ with probability $1/6$, of type $(2, 1)$ with probability $1/2$, and of type (3) with probability $1/3$. Define independent Poisson variables

$$X_\ell := Z_{\ell,(1,1,1)} \sim \mathrm{Poisson}\Big(\frac{\alpha_\ell}{6\ell}\Big), \quad Y_\ell := Z_{\ell,(2,1)} \sim \mathrm{Poisson}\Big(\frac{\alpha_\ell}{2\ell}\Big), \quad W_\ell := Z_{\ell,(3)} \sim \mathrm{Poisson}\Big(\frac{\alpha_\ell}{3\ell}\Big).$$

Then, for each $i \geq 1$, (34) gives

$$A_i = 3X_i + Y_i + \mathbf{1}_{2|i}\, Y_{i/2} + \mathbf{1}_{3|i}\, W_{i/3}. \tag{37}$$

Thus, for example,

$$A_1 = 3X_1 + Y_1, \qquad A_2 = 3X_2 + Y_2 + Y_1, \qquad A_3 = 3X_3 + Y_3 + W_1,$$

and similarly at higher lengths. Under $\alpha_\ell \equiv 1$, the expectations are

$$\mathbb{E}[A_i] = \frac{1}{i} + \mathbf{1}_{2|i}\frac{1}{i} + \mathbf{1}_{3|i}\frac{1}{i} = \frac{1 + \mathbf{1}_{2|i} + \mathbf{1}_{3|i}}{i},$$

while under Ewens(θ) the right-hand side is multiplied by θ. Under derangements, all terms involving α_1 vanish, so in particular $W_1 \equiv 0$ and A_3 has no contribution from fixed block-cycles marked by a 3-cycle.

(b) $\Gamma = C_3$ (or $\Gamma = A_3$). Here $\Gamma = \{e, (123), (132)\}$, so the cycle type is $(1, 1, 1)$ with probability $1/3$ and (3) with probability $2/3$, while $(2, 1)$ does not occur. Define independent Poisson variables

$$X_\ell := Z_{\ell,(1,1,1)} \sim \text{Poisson}\Big(\frac{\alpha_\ell}{3\ell}\Big), \qquad W_\ell := Z_{\ell,(3)} \sim \text{Poisson}\Big(\frac{2\alpha_\ell}{3\ell}\Big).$$

Then

$$A_i = 3X_i + \mathbf{1}_{3|i}\, W_{i/3}, \qquad i \geq 1. \tag{38}$$

Under $\alpha_\ell \equiv 1$, this yields $\mathbb{E}[A_i] = (1 + 2\mathbf{1}_{3|i})/i$, and under Ewens($\theta$) the same expression with an overall factor θ.

6.3 Restricted cycle lengths and the disappearance of marked terms

It is useful to make explicit how forbidden block-cycle lengths propagate to the $[kn]$-cycle counts. Suppose that $\alpha_\ell = 0$ for all ℓ in a set $\mathcal{F} \subset \mathbb{N}$ (e.g. $\mathcal{F} = \{1\}$ for derangements, or a finite set of prohibited short lengths). Then $Z_{\ell,\lambda} \equiv 0$ for all $\ell \in \mathcal{F}$ and all $\lambda \vdash k$. Consequently, in the divisor representation (34), only divisors $d \mid i$ with $d \notin \mathcal{F}$ contribute:

$$A_i = \sum_{\substack{d|i \\ d\notin\mathcal{F}}} \sum_{\lambda\vdash k} a_{i/d}(\lambda)\, Z_{d,\lambda}.$$

In particular, any contribution to A_i that would have come from fixed block-cycles ($d = 1$) is removed when $\alpha_1 = 0$. This is exactly what we saw concretely in (36) and (37): for derangements, the terms $Y_{i/2}$ at $i = 2$ and $W_{i/3}$ at $i = 3$ disappear because they arise from $\ell = 1$ block-cycles decorated by nontrivial internal permutations.

6.4 Summary: how to compute (A_i) in practice

Given k, $\Gamma \leq S_k$, and a measure (μ_n) satisfying the Poissonized cycle-index hypothesis, the computation of the limit law proceeds in two steps.

1. Identify (α_ℓ) from the small-cycle asymptotics of (μ_n); in the examples above this yields $\alpha_\ell \equiv 1$ (uniform on S_n or A_n), $\alpha_\ell \equiv \theta$ (Ewens(θ)), $\alpha_1 = 0$ and $\alpha_\ell = 1$ for $\ell \geq 2$ (derangements), or $\alpha_\ell = \theta_\ell$ (restricted/generalized Ewens weights).

2. Compute $P_\Gamma(\lambda)$, i.e. the cycle-type distribution of a uniform element of Γ, and then apply (34). For small k one can write each A_i explicitly as a finite sum of independent Poisson variables, as in (36), (37), and (38).

These explicit decompositions are often sufficient for extracting further information (moments, support properties such as parity obstructions, and limiting probabilities of no short cycles) without additional analytic input.

7 Extensions and limitations

The examples in Section 6 all fall into the "logarithmic" regime encoded by the Poissonized cycle-index hypothesis: for each fixed truncation level L, the joint law of $(b_1(h), \dots, b_L(h))$ behaves (after Poissonization and de-Poissonization) as though the counts were approximately independent Poisson with bounded means α_ℓ/ℓ. In this section we delineate two directions in which one might try to go beyond that regime, and we explain what, in our argument, must be modified and what may fail outright.

7.1 Allowing μ_n on other permutation sets

Our formal input from the base measure μ_n is the analytic asymptotic for the Poissonized truncated cycle generating function F_L. In particular, we never use that μ_n is supported on all of S_n, nor even that it is uniform on a subgroup, except insofar as these properties make the hypothesis verifiable. Consequently, one can replace μ_n by any conjugacy-invariant probability measure supported on a subset $\mathcal{C}_n \subseteq S_n$ (for example, a union of conjugacy classes), provided the same Poissonized hypothesis holds.

A useful sanity check is the extreme case where μ_n is supported on permutations with *no* short cycles. For instance, let μ_n be uniform on the conjugacy class of n-cycles. Then for every fixed L and all $n > L$ we have $b_\ell(h) = 0$ for $\ell \le L$ almost surely, hence

$$\mathbb{E}\Big[\prod_{\ell=1}^{L} x_\ell^{b_\ell(h)}\Big] = 1 \qquad (n > L),$$

so that the Poissonized function is identically $F_L(t; x_1, \dots, x_L) \equiv 1$ and $\log F_L \equiv 0$. The hypothesis holds with $\alpha_\ell \equiv 0$, and the resulting limit vector satisfies $A_i \equiv 0$ for each fixed i. This is consistent with the elementary observation that, for such a base measure, the induced wreath-product permutation σ typically has macroscopic cycles (of lengths comparable to n), and no bounded cycles survive in the limit.

More generally, one may consider measures supported on permutations whose allowed cycle lengths come from a prescribed set $\mathcal{S} \subseteq \mathbb{N}$ (possibly depending on n), or on permutations conditioned on rare events involving only long cycles. As long as the small-cycle sector admits a Poissonized expansion of the prescribed form, the same marked-Poisson pushforward description for $(a_1, \dots, a_B)$ remains valid. The content of the hypothesis is precisely that conditioning and support restrictions do not introduce long-range correlations among the *bounded* cycle counts at the level of fixed L.

7.2 Uniform measures on subgroups: what is automatic and what is not

A recurring special case is $\mu_n = \mathrm{Unif}(H_n)$ for a subgroup $H_n \leq S_n$. Two points should be separated.

Conjugacy invariance. Uniform measure on H_n is conjugacy-invariant as a measure on S_n if and only if H_n is a union of conjugacy classes of S_n, equivalently a normal subgroup. Apart from $H_n = S_n$ and $H_n = A_n$ (for $n \geq 5$), there are no other nontrivial normal subgroups, so "uniform-on-subgroup" and "conjugacy-invariant" are typically incompatible if one insists on viewing μ_n as a measure on S_n.

From the perspective of our proof, conjugacy invariance is a convenient sufficient condition ensuring that the law of h is determined by its cycle counts and that the cycle index formalism applies cleanly. If one is willing to assume the Poissonized hypothesis *directly* for the random vector $(b_1(h), \ldots, b_L(h))$ under μ_n (without deriving it from class-function arguments), then conjugacy invariance can be weakened or removed; what one loses is a robust toolkit for verifying the hypothesis from group structure. In particular, for non-normal H_n one should not expect uniform-on-H_n to behave like any conjugacy-invariant model unless additional averaging (e.g. conjugation by a random element of S_n) is introduced.

Logarithmic behavior. Even when μ_n is conjugacy-invariant (e.g. uniform on a conjugacy class, or on a union of classes), the Poissonized hypothesis may fail because the small-cycle counts need not be tight, nor approximately independent. The hypothesis implies, for each fixed ℓ, that $b_\ell(h)$ has bounded mean and in fact converges in distribution to $\mathrm{Poisson}(\alpha_\ell/\ell)$ after de-Poissonization. If $b_\ell(h)$ typically grows with n (or takes values of order n with non-negligible probability), then no choice of constants (α_ℓ) can make the approximation true.

This tightness requirement is exactly what breaks for several natural families of subgroups and structured sets, including cyclic subgroups, which we discuss next.

7.3 A non-logarithmic example: the cyclic subgroup C_n

Let $C_n = \langle (1\,2\,\cdots\,n) \rangle \leq S_n$ be the cyclic subgroup generated by an n-cycle, and consider $\mu_n = \mathrm{Unif}(C_n)$. This example illustrates two distinct obstructions.

Failure of conjugacy invariance. As noted above, C_n is not normal in S_n, hence $\mathrm{Unif}(C_n)$ is not conjugacy-invariant. Thus it lies outside our standing assumptions as stated.

Failure of logarithmic small-cycle statistics (even ignoring invariance). Write $g = (1\,2\cdots n)$ and sample $h = g^m$ with m uniform in $\{0, 1, \ldots, n-1\}$. Then h has cycle decomposition consisting of $\gcd(n, m)$ cycles, each of length $n/\gcd(n, m)$. In particular, for a fixed ℓ we have

$$b_\ell(h) \in \Big\{0, \frac{n}{\ell}\Big\} \qquad (\text{for } \ell \mid n),$$

and $b_\ell(h) = 0$ identically when $\ell \nmid n$, except for $\ell = 1$ where $b_1(h) = n$ occurs at the identity element.

More precisely, if $\ell \mid n$ then the event $\{b_\ell(h) = n/\ell\}$ occurs exactly when $\gcd(n, m) = n/\ell$, which happens for $\varphi(\ell)$ values of m modulo n. Therefore

$$\mathbb{P}\Big(b_\ell(h) = \frac{n}{\ell}\Big) = \frac{\varphi(\ell)}{n}, \qquad \mathbb{E}[b_\ell(h)] = \frac{n}{\ell} \cdot \frac{\varphi(\ell)}{n} = \frac{\varphi(\ell)}{\ell} \quad (\ell \mid n),$$

and $\mathbb{E}[b_\ell(h)] = 0$ if $\ell \nmid n$.

At the level of first moments, one might be tempted to read off a plausible "intensity" $\alpha_\ell = \varphi(\ell)$ along subsequences $n \equiv 0 \pmod{\ell}$. However, the distributional picture is incompatible with a Poisson limit: for fixed ℓ the random variable $b_\ell(h)$ is *not tight* (it equals n/ℓ with probability $\asymp 1/n$), whereas any Poisson limit would be supported on $\mathbb{N}$ with bounded tails. Even worse, the joint vector $(b_1, \ldots, b_L)$ is supported on a set of size at most n and is determined almost entirely by the single arithmetic quantity $\gcd(n, m)$, so asymptotic independence across different ℓ cannot occur.

This non-logarithmic behavior propagates to the induced cycle counts of σ in a way that defeats any fixed-B Poisson description unless the limit is trivial. For a fixed i, contributions to $a_i(\sigma)$ come from block-cycles of lengths dividing i, hence ultimately from those exponents m for which $h = g^m$ contains cycles of bounded length. Since such exponents have probability $O(1/n)$ for each fixed bounded length, it is plausible (and can be proved in concrete cases) that $a_i(\sigma) \to 0$ in probability for each fixed i, yielding a degenerate limit. But this "degeneracy" depends on delicate arithmetic of n (e.g. whether n has small divisors) and is not governed by a stable family of constants (α_ℓ) as required by our hypothesis. In particular, even if one restricts to subsequences $n \in q\mathbb{N}$, the remaining correlations among different ℓ are not Poissonian.

We view this as representative of a broader obstruction: models in which the base permutation is generated by a small-entropy parameter (here, $m \in \mathbb{Z}/n\mathbb{Z}$) frequently produce cycle counts that are either macroscopic or strongly correlated, and hence fall outside the logarithmic universality class.

7.4 Changing the ground set size: measures on S_{m_n} and the case of B_n

Our notation has fixed the base permutation h to lie in S_n and the wreath-product action to produce an element of S_{kn}. In many settings the natural base group acts on a different number of points, for example:

- the hyperoctahedral group $B_n \cong C_2 \wr S_n$, which has a canonical faithful action on $2n$ points (signed letters);
- imprimitive groups acting on dn points for a fixed d (already visible in iterated wreath products).

There are two essentially different ways to incorporate such examples.

(a) A notational generalization with m_n points. Suppose we have a sequence $m_n \to \infty$ and conjugacy-invariant measures μ_n on S_{m_n}. We may then form $\sigma_n \in S_{km_n}$ by sampling $h \sim \mu_n$ and $\gamma_1, \ldots, \gamma_{m_n} \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}(\Gamma)$ and letting $\sigma_n = (\gamma_1, \ldots, \gamma_{m_n}; h)$ act on $[km_n]$ in the usual imprimitive way. All definitions remain identical with n replaced by m_n in the wreath-product construction. The analytic Poissonization step, however, is indexed by n, not by m_n, so one must define

$$F_L^{(m)}(t; x_1, \ldots, x_L) := \sum_{n \geq 0} (1-t)t^n \, \mathbb{E}_{h \sim \mu_n} \Big[\prod_{\ell=1}^{L} x_\ell^{b_\ell(h)} \Big],$$

and assume a corresponding expansion for $\log F_L^{(m)}$. When m_n is linear, say $m_n = dn$, one expects the same Tauberian de-Poissonization mechanism to go through with minimal change; when m_n is irregular (e.g. sparse subsequences), one should not expect a uniform de-Poissonization statement without additional hypotheses.

(b) Intrinsic cycle indices for other base groups (example: B_n). If the base measure is uniform on B_n (or conjugacy-invariant within B_n), then viewing it as a measure on S_{2n} is problematic: it is not conjugacy-invariant in S_{2n}, and its cycle structure in S_{2n} has additional internal types. Concretely, conjugacy classes in B_n are described by *signed* cycle data, often encoded as a pair of partitions (λ^+, λ^-) where λ^+ lists lengths of "positive" cycles and λ^- lists lengths of "negative" cycles. Under the embedding into S_{2n}, positive and negative cycles contribute differently to the ordinary cycle counts: one obtains a mixture of ℓ-cycles and 2ℓ-cycles depending on the sign type.

To fit such a model into our framework, one should replace the unmarked cycle counts $b_\ell(h)$ by a vector of marked counts, say $(b_\ell^+(h), b_\ell^-(h))_{\ell \geq 1}$, and correspondingly replace the scalar variables x_ℓ in F_L by pairs (x_ℓ^+, x_ℓ^-). The appropriate hypothesis becomes a two-parameter Poissonized expansion of the form

$$\log F_L^B(t; \{x_\ell^+\}, \{x_\ell^-\}) = \sum_{\ell=1}^{L} \frac{t^\ell}{\ell} \Big(\alpha_\ell^+ (x_\ell^+ - 1) + \alpha_\ell^- (x_\ell^- - 1) \Big) + \text{controlled error},$$

for some intensities $(\alpha_\ell^\pm)$. Once such an expansion is available, the remainder of our argument is formal: one obtains independent Poisson limits for the marked block-cycle counts, and the final cycle counts in the induced action are obtained by a deterministic pushforward which now depends on the embedding rule (e.g. whether a negative ℓ-cycle contributes to a_ℓ or to $a_{2\ell}$). In other words, B_n forces one to enlarge the mark space, but the "Poisson $\Rightarrow$ pushforward-Poisson" mechanism persists.

We emphasize that this is not merely a cosmetic change. For B_n, the natural cycle index is not the S_n cycle index but its signed analogue, and the verification of the needed Poissonized hypothesis should be carried out in that language. The conclusion, once translated back into S_{2n}-cycle counts, is still a statement about the joint limit of bounded cycle counts, but with additional arithmetic constraints stemming from the sign structure.

7.5 What our method does *not* address

Finally, we record two limitations that are implicit in the preceding discussion.

Non-logarithmic regimes. When the base model exhibits either macroscopic fluctuations in $b_\ell(h)$ for bounded ℓ or strong dependence among different ℓ, our method provides no substitute limit theorem. The marked-Poisson description is, by construction, a universality statement for logarithmic combinatorial structures; cyclic subgroups and other low-entropy models are not expected to fall into this class.

Non-conjugacy-invariant base measures without direct hypotheses. If μ_n is not conjugacy-invariant and one cannot verify a Poissonized hypothesis for F_L (or an appropriate marked variant), then even the correct choice of statistics to control becomes unclear: the distribution of σ may depend on features of h beyond its cycle counts, and additional combinatorial invariants may be needed to describe the induced action. Any extension in this direction would require new ideas, either to identify a replacement for cycle index methods or to impose alternative structural conditions that recover approximate independence of the relevant local statistics.

7.6 Open problems and next steps: mesoscopic windows, sharper constants, and structural criteria

Our main theorem is deliberately "local": for each fixed B we obtain convergence of the joint law of $(a_1(\sigma), \ldots, a_B(\sigma))$ to a marked-Poisson pushforward, together with a total variation bound of the form

$$\left\| \mathcal{L}(a_1, \ldots, a_B) - \mathcal{L}(A_1, \ldots, A_B) \right\|_{\mathrm{TV}} \le CB\, \varepsilon_n.$$

We record here three directions in which it is natural to push further.

1. Mesoscopic windows $B = B(n)$. The displayed bound already yields a weak "mesoscopic" statement: if $B = B(n)$ is any sequence with $B(n)\varepsilon_n \to 0$, then the joint law of $(a_1, \ldots, a_{B(n)})$ converges in total variation to the corresponding truncation $(A_1, \ldots, A_{B(n)})$ of the limiting Poisson model. This is useful only insofar as one can quantify ε_n.

In our proof, ε_n ultimately comes from de-Poissonization at a parameter $t = t_n$ with $1 - t_n \asymp 1/n$ and from the uniform bound on the remainder term $\mathcal{E}_L$ on a complex polydisc. Heuristically, when one plugs $t = 1 - \Theta(1/n)$ into the hypothesis, the analytic error term $O(L^2(1-t))$ becomes $O(L^2/n)$. This suggests that the natural barrier for pushing L (and hence B) as a function of n is of order $n^{1/2}$. In other words, even in the most optimistic scenario where the hypothesis holds *uniformly* in L up to $L = L(n)$, the error term one would obtain by the same argument is meaningful only when $L(n) = o(n^{1/2})$.

This raises a concrete open problem: identify verifiable hypotheses under which one can take $B(n) \to \infty$, and determine the maximal growth rate for which the marked-Poisson description remains valid. A representative conjecture, tailored to our analytic framework, is the following.

> **Conjecture (mesoscopic marked-Poisson approximation).** Assume the Poissonized cycle-index hypothesis can be strengthened so that, for $L = L(n)$ and t in a complex domain with $1 - t \asymp 1/n$, one has
>
> $$\log F_L(t; x_1, \ldots, x_L) = \sum_{\ell \le L} \frac{\alpha_\ell t^\ell}{\ell}(x_\ell - 1) + O\left(\frac{L^2}{n}\right)$$
>
> uniformly for $|x_\ell| \le R$ and all $L \le L(n)$, with constants independent of n. Then for any $B(n) \le L(n)$ with $B(n) = o(n^{1/2})$, the joint law of $(a_1(\sigma), \ldots, a_{B(n)}(\sigma))$ is close in total variation to the corresponding truncation of the marked-Poisson pushforward, with an error tending to 0.

Even for the classical case $\Gamma = \{e\}$ and μ_n uniform on S_n, much sharper mesoscopic results are known by probabilistic couplings (e.g. the Feller coupling) and by Stein–Chen methods, and these typically work up to $B(n) = n^\beta$ for some $\beta < 1$ with explicit error terms. It would be interesting to develop an analogue for wreath products in which the "marks" coming from Γ are incorporated directly into the coupling. At a technical level, this would require a joint construction of block cycles of h and the internal permutations along those block cycles, so that the induced i-cycles of σ emerge with the correct dependencies.

A second aspect of the mesoscopic problem is conceptual. When B grows, the limiting object is no longer a fixed-dimensional vector but an array $(A_i)_{1\le i\le B(n)}$ whose coordinates are dependent through shared $Z_{\ell,\lambda}$ contributions. While the $Z_{\ell,\lambda}$ are independent, the pushforward produces relations such as

$$A_i = \sum_{\ell\mid i}\sum_{\lambda\vdash k} a_{i/\ell}(\lambda)\, Z_{\ell,\lambda}.$$

Thus, in large windows, one should view (A_i) as a linear image of a field of independent Poisson variables indexed by (ℓ,λ). Understanding which statistics of $(a_i)_{i\le B(n)}$ exhibit Gaussian behavior (via many independent contributions) and which remain genuinely Poissonian is itself a nontrivial question. For instance, the total number of points in cycles of length at most B,

$$T_B(\sigma) := \sum_{i\le B} i\, a_i(\sigma),$$

is a natural "mesoscopic mass" statistic. Under the limiting model,

$$T_B \approx \sum_{\ell\le B}\sum_{\lambda\vdash k} \ell\Big(\sum_{j\le B/\ell} j\, a_j(\lambda)\Big) Z_{\ell,\lambda},$$

which is a compound Poisson variable whose Lévy measure changes with B. Determining regimes of $B = B(n)$ in which T_B admits a central limit theorem (after centering and scaling) would connect our local picture to more global fluctuation theory.

2. Sharper constants and rates. Our total variation bound is adequate for qualitative convergence but is far from optimal, and in applications one may want an explicit rate as a function of (k,Γ) and of the base model. There are at least three sources of looseness.

(i) De-Poissonization. The analytic Tauberian step is robust but not tuned to any specific μ_n. In concrete models (e.g. Ewens-like measures or measures arising from analytic combinatorics), one often has sharper information on the singular behavior of the relevant generating functions, and this can yield explicit polynomial rates in n. It would be useful to isolate, within our argument, the minimal analytic input needed to obtain a specified rate, e.g. an error $O(n^{-\delta})$ for some $\delta > 0$.

(ii) Multivariate approximation in total variation. Even if one knows that $(b_1(h),\ldots,b_L(h))$ is close to a vector of independent Poisson variables, there remains a nontrivial step in propagating this approximation through the wreath-product construction and then bounding the distance between the induced cycle count vectors. Our present treatment bounds the discrepancy coordinatewise and then union-bounds in B, leading to a factor linear in B. In principle one could do better by exploiting the explicit linear structure

of the pushforward from $(Z_{\ell,\lambda})$ to (A_i). For example, since $(A_1, \dots, A_B)$ is an affine map of the vector $(Z_{\ell,\lambda})_{\ell \le B,\ \lambda \vdash k}$, any distance control on the latter could be transported using data-processing inequalities tailored to Poisson measures, potentially reducing the dependence on B.

(iii) Stein–Chen methods with marks. A natural next step is to adapt Stein's method for Poisson process approximation to the marked setting. One can regard the "atoms" as cycles of h with their induced internal mark (cycle type in Γ along the block cycle), yielding a random marked point process on $\mathbb{N} \times \{\lambda \vdash k\}$. If one can show that this point process is close (in an appropriate metric) to a Poisson point process with intensity $\alpha_\ell P_\Gamma(\lambda)/\ell$, then the finite-dimensional cycle counts follow by applying a deterministic functional. Such an approach has the potential to produce explicit constants and to handle $B = B(n)$ more flexibly, because one can work directly with process-level metrics (e.g. Wasserstein-type bounds) that behave better under truncation than total variation.

In the same vein, it would be useful to quantify the dependence of constants on the group Γ. Even for fixed k, the number of partitions $\lambda \vdash k$ grows, and a naive bound may introduce unnecessary factors depending on $p(k)$ (the partition number). Since k is fixed in our framework, this is not an asymptotic issue, but it matters for explicit estimates. A more careful accounting could express the constants in terms of a small collection of group parameters, such as

$$\max_{\lambda \vdash k} P_\Gamma(\lambda), \qquad \sum_{\lambda \vdash k} \sum_{j \ge 1} a_j(\lambda)^2 P_\Gamma(\lambda),$$

or other moments of the cycle-count profile of a uniform Γ-element. These quantities naturally appear when one bounds variances of A_i and can serve as proxies for the complexity of the mark distribution.

3. Structural criteria: which families (H_n) satisfy the hypothesis? The Poissonized cycle-index hypothesis is stated analytically, and in examples it is typically verified by a direct computation of a cycle index or by importing known generating-function asymptotics. An important missing piece is a higher-level classification: given a family of measures μ_n (or subgroups H_n with $\mu_n = \mathrm{Unif}(H_n)$), can one decide, without detailed cycle index manipulations, whether the hypothesis holds?

We separate three related problems.

(a) Characterizing the hypothesis in probabilistic terms. Our analytic assumption implies that for each fixed L the vector $(b_1, \dots, b_L)$ behaves like independent Poisson with means α_ℓ/ℓ. Conversely, one expects that a suitable family of factorial moment asymptotics implies the analytic statement (at least locally in the polydisc). It would be valuable to formulate an equivalence between the analytic Poissonized hypothesis and a finite list of probabilistic conditions, for example:

- boundedness and convergence of mixed factorial moments of $(b_1, \ldots, b_L)$;
- asymptotic factorization of these moments (approximate independence);
- a uniform integrability condition sufficient to justify analytic continuation in x_ℓ.

Such a reformulation would make it easier to verify the hypothesis by soft probabilistic arguments, and could provide a natural extension to non-conjugacy-invariant settings in which cycle index tools are unavailable.

(b) Closure properties and universality. Many "logarithmic" combinatorial structures are stable under operations such as superposition and conditioning, and their small-component counts admit Poisson limits with predictable intensities. In our context one can ask for analogous closure statements for measures on S_n: if μ_n and ν_n satisfy the hypothesis with parameters (α_ℓ) and (β_ℓ), does an appropriate convolution, mixture, or conditioning of these measures satisfy the hypothesis with parameters derived from (α_ℓ) and (β_ℓ)? A particularly natural operation is multiplication of independent permutations: if $h_1 \sim \mu_n$ and $h_2 \sim \nu_n$ are independent, what conditions ensure that $h_1 h_2$ has logarithmic small-cycle statistics? For uniform permutations this is trivial (closure under multiplication), but for structured measures it is not. Establishing such closure properties would allow one to build new examples systematically and would clarify which aspects of the hypothesis are genuinely restrictive.

(c) Subgroups H_n and group-theoretic criteria. As emphasized earlier, if one insists that μ_n be conjugacy-invariant as a measure on S_n, then uniform measure on a subgroup forces H_n to be normal, leaving essentially S_n and A_n (for $n \geq 5$) as the only nontrivial cases. In that strict sense, there is nothing to classify.

However, two broader classification problems remain meaningful.

First, one may drop conjugacy invariance and instead ask for conditions on H_n ensuring that the cycle counts $(b_1, \ldots, b_L)$ under $\mathrm{Unif}(H_n)$ have an approximately Poisson law. There is a substantial literature on cycle structure of random elements in permutation groups, often expressed via fixed-point ratios and minimal degree. A plausible conjectural dichotomy is that, for large primitive groups other than A_n and S_n, small cycles are typically *suppressed* so strongly that the limiting intensities are $\alpha_\ell = 0$, leading to a degenerate short-cycle limit. Proving (or refuting) such a statement in a form compatible with our wreath-product conclusions would require new input from permutation group theory.

Second, one may restore conjugacy invariance by averaging over conjugation: given a (possibly non-normal) subgroup H_n, consider the random element ghg^{-1} with $h \sim \mathrm{Unif}(H_n)$ and $g \sim \mathrm{Unif}(S_n)$ independent. This produces a conjugacy-invariant measure on S_n supported on the union of conjugacy classes intersecting H_n. It is then natural to ask: how do structural properties of H_n control the resulting intensities (α_ℓ)? In this formulation, one may hope for criteria in terms of the distribution of cycle types inside H_n rather than in terms of a full cycle index computation.

4. Interactions with the wreath-product layer. Finally, even when the base model is understood, there remains the question of how the internal group Γ modifies mesoscopic behavior. Because the mapping from $(Z_{\ell,\lambda})$ to (A_i) mixes many marks into each A_i (through the constraint $j\ell = i$), one expects that mesoscopic scaling limits may depend sensitively on the arithmetic of i and on the support of P_Γ on partitions of k. For example, if Γ is cyclic of order k, then λ is supported on partitions consisting of equal parts, whereas for $\Gamma = S_k$ all partitions occur. Understanding which features of Γ are identifiable from the induced short-cycle process $(a_i(\sigma))_{i\geq 1}$, and how this identifiability degrades as one moves from fixed B to $B = B(n)$, is a natural inverse problem. At a technical level, this may require controlling not only marginal distributions but also fine correlations among the $a_i(\sigma)$ across a growing range of i.

We expect that progress on any of the problems above—mesoscopic windows, sharper rates, and structural criteria—will require combining the analytic Poissonization viewpoint of the present work with more probabilistic tools (couplings, Stein's method, process approximations) and, in the subgroup direction, with input from the theory of permutation groups. The marked-Poisson pushforward mechanism is rigid and transparent; the main challenge is to identify hypotheses under which one can invoke it uniformly beyond fixed truncations and without bespoke cycle-index computations.