Mesoscopic Compound-Poisson Approximation for Cycles in Wreath Products $\Gamma^n \rtimes S_n$

Liz Lemma Future Detective

January 17, 2026

Abstract

Diaconis and Tung (arXiv:2408.06611) identify the small-cycle limits of a uniform permutation in the wreath product $\Gamma^n \rtimes S_n \leq S_{kn}$ as a dependent compound-Poisson vector and provide a coupling bound that is sharp enough for fixed-dimensional convergence. We develop a mesoscopic theory: the joint cycle-count vector $(a_1, \ldots, a_B)$ remains close in total variation to the same compound-Poisson limit even when $B = B(n)$ grows with n, and we give explicit quantitative error bounds. The proof refines the wreath-Feller coupling by treating block-closure indicators as an inhomogeneous Bernoulli renewal process and by propagating the resulting 'spacing errors' through bounded marks given by the cycle types of random elements of Γ. As a byproduct we obtain approximation guarantees for a broad class of statistics depending on many small cycle lengths, suitable for algorithmic diagnostics and high-dimensional limit theorems.

Table of Contents

5. 5. Mesoscopic coupling of finite-n and infinite processes: boundary/tail spacings for lengths $\leq B$; refined probability bounds uniform in B; explicit constants

6. 6. From spacing discrepancies to total variation bounds: coupling inequality; bounded mark effects; proof of TV bound for $(a_1, \ldots, a_B)$

7. 7. Stein/functional bounds (optional strengthening): expectation bounds for Lipschitz/polynomial-growth test functionals; dependency-graph formulation and explicit constants

8. 8. Worked examples and numerics: $\Gamma = C_k$ (independent limits), $\Gamma = S_3$ (dependence by shared Poisson components), $\Gamma = S_k$; illustration of error scaling in B/n

9. 9. Discussion and open problems: optimal constants, lower bounds, extension to other block groups $H \leq S_n$, and toward functional limit theorems

1 Introduction and statement of results

We study the cycle structure of a uniform random element of the imprimitive permutation group

$$G_{k,n} \;=\; \Gamma^n \rtimes S_n \;\leq\; S_{kn},$$

where $k \geq 1$ and the subgroup $\Gamma \leq S_k$ are fixed while $n \to \infty$. The action is the standard block action on $[kn]$, with n blocks of size k: the factor S_n permutes blocks and the factor Γ^n acts within blocks. Writing σ for a uniform random element of $G_{k,n}$ and $a_i(\sigma)$ for the number of i-cycles of σ as a permutation of $[kn]$, our objective is to describe the joint law of

$$\big(a_1(\sigma), \ldots, a_B(\sigma)\big)$$

in a regime where the truncation parameter $B = B(n)$ is allowed to grow with n. The dependence on B is the central point: we are not only interested in a fixed finite set of cycle lengths, but in a mesoscopic window of lengths whose size diverges while remaining negligible compared with n.

The mesoscopic scaling $B = o(n)$ is natural for imprimitive groups. At the block level, a typical element of S_n has cycles on all scales, including macroscopic cycles of order n, and these macroscopic block-cycles create macroscopic cycles in S_{kn}. The small and intermediate cycles, in contrast, admit Poisson-type approximations. The cutoff $B = o(n)$ isolates a range of lengths for which one may still expect approximate independence and Poisson fluctuations, while permitting $B \to \infty$ so that the approximation becomes genuinely high-dimensional. In applications, one often wishes to treat the vector of counts as a feature vector whose dimension grows with sample size; our results provide an explicit control of the approximation error as that dimension increases.

The principal phenomenon is that the cycle structure of σ is obtained by "compounding" two layers of randomness. First, the block permutation in S_n generates block-cycles of various lengths ℓ. Second, each block-cycle carries an internal permutation obtained by multiplying the Γ-elements encountered along the block-cycle; the cycle type of this internal product determines how the ℓ blocks are braided into cycles on $[kn]$. Concretely, if the internal product has cycle type $\lambda \vdash k$ and contains $a_j(\lambda)$ cycles of length j, then the corresponding contribution to the action on $[kn]$ produces $a_j(\lambda)$ cycles of length ℓj. This mechanism suggests a limit object built from independent Poisson counts of block-cycles, each block-cycle being independently marked by a random Γ-cycle type.

The limiting vector we use is a compound Poisson field indexed by cycle length. For each block length $\ell \geq 1$ and each partition $\lambda \vdash k$, we introduce an independent Poisson variable with mean proportional to ℓ^{-1} and weighted by the cycle-type probabilities of a uniform element of Γ. The limiting count A_i for i-cycles is then obtained by summing, over all ℓ dividing i and all $\lambda \vdash k$, the number $a_{i/\ell}(\lambda)$ of internal cycles that inflate an ℓ-block cycle to an i-cycle. Thus $(A_i)_{i\geq 1}$ incorporates the full cycle-type distribution of Γ, and reduces to the classical independent Poisson$(1/i)$ limit when $k = 1$ (equivalently, Γ is trivial and $G_{1,n} \cong S_n$). The point of this construction is not only to identify the marginals, but also to produce a vector with an explicit product structure at the level of marked block-cycles, from which one can derive joint distributional approximations for many coordinates simultaneously.

Our main theorem asserts that, under the fixed-k and fixed-Γ assumptions, the law of the truncated cycle-count vector converges in total variation to the corresponding truncation of the compound Poisson limit provided $B = o(n)$. Moreover, we obtain an explicit quantitative bound on the total variation distance for all $1 \leq B \leq n$, of the form

$$\mathrm{error}(n, B) \ \leq \ C(k, \Gamma)\Big(\frac{B}{n} + \frac{B^2}{n^2}\Big),$$

with a constant depending only on k and the cycle-type distribution of Γ. In particular, if $B = o(n)$ then the right-hand side tends to 0, and if $B \leq n^{1/2}$ one obtains a uniform $O(B/n)$ rate. We also isolate a class of test functionals $F : \mathbb{N}^B \to \mathbb{R}$ (including, for example, Lipschitz statistics with polynomial growth) for which we bound the difference of expectations by $O(B/n)$ with an explicit constant depending on F and (k, Γ). These bounds are designed to be used as plug-in estimates: once a statistic can be expressed as a Lipschitz functional of the truncated count vector, the approximation error follows immediately.

The motivation for controlling total variation in growing dimension comes from two related considerations. First, in Poisson approximation theory it is often feasible to match moments or finite-dimensional marginals, but statistical questions typically require uniform control of the entire joint law on a large coordinate set. Total variation bounds imply uniform approximation of probabilities of all events determined by $(a_1, \dots, a_B)$, and therefore allow one to replace the finite-n distribution by the limiting compound Poisson model when computing p-values, likelihood ratios, or confidence regions based on many cycle lengths simultaneously. Second, in high-dimensional regimes one must account for the accumulation of small errors across coordinates; bounds that are explicit in B are therefore essential. In the present setting, the dependence on B is sharp enough to permit mesoscopic growth and to recover fixed-B convergence as a special case.

For comparison, when Γ is trivial our setting reduces to uniform permutations in S_n, and the Poisson approximation for small cycle counts is classical. Goncharov identified the limiting independent Poisson law for each fixed set of cycle lengths, and subsequent work developed couplings and error bounds. In particular, bounds of order B^2/n (and refinements thereof) for the total variation distance between $(a_1, \dots, a_B)$ and independent Poisson$(1/i)$ variables can be obtained by methods based on the Feller coupling and on Stein's method for Poisson process approximation. Our contribution is the analogous theory for the imprimitive wreath product family, including (i) an explicit compound Poisson limit that reflects the internal group Γ, and (ii) a quantitative bound that remains informative as $B \to \infty$ provided $B = o(n)$. Even when one restricts to fixed B, the wreath product structure introduces a nontrivial compounding mechanism absent from S_n, so that the approximation is not simply a direct product of Poisson laws for each $a_i(\sigma)$.

In the wreath product case, fixed-i limit theorems for cycle counts have been investigated, notably in work of Diaconis and Tung, where the compound Poisson structure emerges naturally from the decomposition into block-cycles and internal cycle types. The novelty here is that we treat a growing set of cycle lengths simultaneously and provide an error bound uniform in the truncation dimension. This requires controlling not only the marginal distribution of block-cycle counts but also the interaction between the truncation at n blocks and the marking by Γ-cycle types. The quantitative nature of the result depends on making this interaction explicit, rather than relying on abstract convergence arguments.

At a technical level, our approach is based on a marked version of the Feller coupling. The usual Feller construction represents the cycle structure of a uniform permutation in S_n via a sequence of independent Bernoulli indicators (ζ_i) with $\mathbb{P}(\zeta_i = 1) = 1/i$, where the spacings between successive 1's encode cycle lengths. In our setting, we apply this construction at the block level to generate block-cycle lengths, and then attach to each spacing an independent mark Y_i taking values in the set of partitions of k, distributed as the cycle type of a uniform element of Γ. The marked spacings generate a natural infinite compound Poisson object, and the finite-n model is obtained by truncating the indicator sequence at n and forcing a terminal 1 at time $n+1$. The primary source of error is the possible influence of the truncation on the counts of short spacings, and our bounds are obtained by estimating the probability of discrepancies between the finite and infinite marked-spacing processes in the range up to B.

Finally, we record two structural consequences of the theorem that will be used repeatedly. First, the approximation identifies the correct order of magnitude for all mixed moments of $(a_1, \ldots, a_B)$ that are controlled by a polynomial in the coordinates, uniformly over $B = o(n)$. Second, since the limit field is built from independent Poisson variables indexed by (ℓ, λ), one may compute the joint distribution of $(A_1, \ldots, A_B)$ explicitly (at least in principle) by conditioning on the underlying Poisson counts and aggregating contributions to each i. This explicitness is useful when one wishes to compare different subgroups Γ or different values of k through observable cycle-count statistics. In the next section we recall the wreath product action, fix our cycle-type notation, and describe the compound Poisson limit in the form most suitable for coupling and total variation estimates.

2 Background on wreath products and cycle counts

2.1 The imprimitive block action of $\Gamma^n \rtimes S_n$

We fix an integer $k \geq 1$ and a subgroup $\Gamma \leq S_k$. For each $n \geq 1$ we consider the wreath product

$$G_{k,n} \;=\; \Gamma^n \rtimes S_n,$$

where S_n acts on Γ^n by permuting coordinates:

$$\pi \cdot (\gamma_1, \ldots, \gamma_n) \;=\; (\gamma_{\pi^{-1}(1)}, \ldots, \gamma_{\pi^{-1}(n)}).$$

Thus the multiplication rule is

$$(\gamma, \pi)\,(\gamma', \pi') \;=\; (\gamma \cdot (\pi \cdot \gamma'),\, \pi\pi'), \qquad \gamma, \gamma' \in \Gamma^n,\; \pi, \pi' \in S_n.$$

To embed $G_{k,n}$ in S_{kn} we identify the underlying set $[kn]$ with $[n] \times [k]$ (block index and within-block coordinate). We use the standard imprimitive action given by

$$(\gamma, \pi) \cdot (b, r) \;=\; \big(\pi(b),\, \gamma_{\pi(b)}(r)\big), \qquad b \in [n],\; r \in [k], \tag{1}$$

where $\gamma = (\gamma_1, \ldots, \gamma_n) \in \Gamma^n$. One checks directly that (1) is compatible with the semidirect product multiplication, hence defines a faithful permutation representation $G_{k,n} \hookrightarrow S_{kn}$. We refer to $[k]$ as the *fiber* and to each set $\{b\} \times [k]$ as a *block*.

We shall frequently write a random element of $G_{k,n}$ in the form

$$\sigma = (\gamma, \pi), \qquad \gamma = (\gamma_1, \ldots, \gamma_n) \in \Gamma^n,\; \pi \in S_n,$$

with σ uniform on $G_{k,n}$. Under this uniform law, π is uniform on S_n, the coordinates $\gamma_1, \ldots, \gamma_n$ are i.i.d. uniform on Γ, and γ is independent of π.

2.2 Cycle types in S_k and notation

We use the standard partition notation for cycle types in S_k. If $\lambda \vdash k$ is a partition, we write $a_j(\lambda)$ for the number of parts of λ equal to j; equivalently, if $\tau \in S_k$ has cycle type λ, then $a_j(\lambda)$ is the number of j-cycles of τ. In particular,

$$\sum_{j=1}^{k} j\, a_j(\lambda) = k.$$

We also write

$$\mathbb{P}_\Gamma(\lambda) \;=\; \mathbb{P}\big(\mathrm{ctype}(\gamma) = \lambda\big),$$

where γ is uniform on Γ and $\mathrm{ctype}(\cdot)$ denotes the cycle type as an element of S_k.

At the level of S_{kn}, for $\sigma \in S_{kn}$ we denote by $a_i(\sigma)$ the number of i-cycles of σ in its disjoint cycle decomposition. Our objective is to understand $(a_i(\sigma))_{i \leq B}$ when σ is uniform in $G_{k,n}$ and B grows with n.

2.3 From block cycles to global cycles: the compounding mechanism

The key combinatorial observation is that the cycle structure of $\sigma = (\gamma, \pi)$ in S_{kn} is determined by the cycle structure of the block permutation $\pi \in S_n$, together with certain products of the Γ-labels along the cycles of π.

Let $c = (b_1\, b_2\, \dots\, b_\ell)$ be a cycle of π of length ℓ (so $\pi(b_t) = b_{t+1}$ with indices modulo ℓ). Consider the product in Γ

$$g_c \;:=\; \gamma_{b_\ell}\gamma_{b_{\ell-1}}\cdots\gamma_{b_1} \;\in \Gamma. \tag{2}$$

The relevance of g_c comes from iterating the action (1): starting from a point (b_1, r), after one application of σ we move to $(b_2, \gamma_{b_2}(r))$, after two applications to $(b_3, \gamma_{b_3}\gamma_{b_2}(r))$, and after ℓ applications we return to the original block b_1 with internal coordinate transformed by g_c. Thus, on the subset $\{b_1, \dots, b_\ell\} \times [k]$, the permutation σ^ℓ acts as $(b_t, r) \mapsto (b_t, g_c(r))$ (up to the obvious relabeling of the cycle starting point), and σ itself interlaces the block motion with these internal transformations.

This yields the following standard lemma.

Lemma 2.1 (Inflation of cycles along a block cycle)**.** *Let c be a block cycle of π of length ℓ, and let $g_c \in \Gamma$ be defined by* (2)*. If g_c has cycle type $\lambda \vdash k$, then the restriction of σ to the ℓk points in the blocks of c decomposes into*

$$a_j(\lambda) \quad \textit{cycles of length} \;\; \ell j, \qquad j = 1, \dots, k.$$

Proof. Fix a j-cycle $(r_1\, r_2\, \dots\, r_j)$ of g_c in $[k]$. Starting from (b_1, r_1) and applying σ repeatedly, we advance one block at a time; after ℓ steps we return to block b_1 and the internal coordinate becomes $g_c(r_1) = r_2$. After ℓj steps we return to (b_1, r_1), and no smaller positive multiple of ℓ closes the orbit because the internal coordinate closes only after j applications of g_c. Distinct cycles of g_c yield disjoint orbits in $[n] \times [k]$. Counting over all j gives the claim. □

Lemma 2.1 is the source of the compounding rule: each block cycle of length ℓ produces global cycles whose lengths are ℓ times the internal cycle lengths determined by g_c. Consequently, the i-cycle count $a_i(\sigma)$ is obtained by summing, over all divisors $\ell \mid i$, the contributions from ℓ-cycles of π whose internal product has an (i/ℓ)-cycle.

To make this explicit, let $N_{\ell,\lambda}(\sigma)$ denote the number of ℓ-cycles c of π such that $\mathrm{ctype}(g_c) = \lambda$. Then, by Lemma 2.1,

$$a_i(\sigma) \;=\; \sum_{\lambda \vdash k} \sum_{\ell \mid i} a_{i/\ell}(\lambda)\, N_{\ell,\lambda}(\sigma), \qquad 1 \le i \le kn. \tag{3}$$

The random array $\{N_{\ell,\lambda}(\sigma)\}$ is therefore the natural intermediate object: it records block-cycle lengths together with the internal cycle type associated to each block cycle.

2.4 Distribution of internal products along block cycles

When σ is uniform in $G_{k,n}$, the family of products $\{g_c\}$ over the disjoint cycles c of π has a particularly simple law. Indeed, conditional on π, the sets of indices of γ used in (2) are disjoint across different block cycles, hence the corresponding products are independent. Moreover, each product is uniform on Γ.

Lemma 2.2 (Uniformity and independence of cycle products)**.** *Let* $\sigma = (\gamma, \pi)$ *be uniform on* $G_{k,n}$*. Conditional on* π*, the random variables* $\{g_c : c \text{ a cycle of } \pi\}$ *defined by* (2) *are independent and each is uniform on* Γ*. In particular,*

$$\mathbb{P}\big(\mathrm{ctype}(g_c) = \lambda \mid \pi\big) = \mathbb{P}_\Gamma(\lambda) \quad \text{for all } \lambda \vdash k.$$

Proof. Fix a cycle $c = (b_1 \ldots b_\ell)$ of π. The coordinates $\gamma_{b_1}, \ldots, \gamma_{b_\ell}$ are i.i.d. uniform on Γ, so their product g_c is uniform on Γ by the translation-invariance of the uniform measure on a finite group. For distinct cycles $c \neq c'$, the index sets $\{b_1, \ldots, b_\ell\}$ and $\{b'_1, \ldots, b'_{\ell'}\}$ are disjoint, so the corresponding products are functions of disjoint subfamilies of independent coordinates and are therefore independent. □

Lemma 2.2 implies that $N_{\ell,\lambda}(\sigma)$ can be viewed as a *thinning* of the number of ℓ-cycles of π by independent marks with law $\mathbb{P}_\Gamma(\cdot)$. If we write $C_\ell(\pi)$ for the number of ℓ-cycles of π, then conditional on π we have

$$\big(N_{\ell,\lambda}(\sigma)\big)_{\lambda \vdash k} \;\sim\; \mathrm{Multinomial}\big(C_\ell(\pi);\; (\mathbb{P}_\Gamma(\lambda))_{\lambda \vdash k}\big),$$

and the collections corresponding to different values of ℓ are conditionally independent given π.

2.5 The Diaconis–Tung compound-Poisson limit

For a uniform permutation $\pi \in S_n$, the classical small-cycle theory asserts that for each fixed ℓ, the random variable $C_\ell(\pi)$ is approximately Poisson$(1/\ell)$, and the family $(C_\ell(\pi))_{\ell \le L}$ is approximately independent for fixed L as $n \to \infty$. Combining this with the independent marking mechanism from Lemma 2.2 leads to a natural candidate for the limit of the marked block-cycle counts: for each $\ell \ge 1$ and each $\lambda \vdash k$, one expects

$$N_{\ell,\lambda}(\sigma) \;\approx\; Z_{\ell,\lambda}, \qquad Z_{\ell,\lambda} \sim \mathrm{Poisson}\Big(\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\Big),$$

with the collection $\{Z_{\ell,\lambda}\}_{\ell \ge 1,\, \lambda \vdash k}$ independent.

Pushing this approximation through the deterministic aggregation (3) yields the compound-Poisson field $(A_i)_{i \ge 1}$ defined by

$$A_i \;:=\; \sum_{\lambda \vdash k} \sum_{\ell \mid i} a_{i/\ell}(\lambda)\, Z_{\ell,\lambda}, \qquad i \ge 1, \tag{4}$$

where the $(Z_{\ell,\lambda})$ are independent with means $\mathbb{P}_\Gamma(\lambda)/\ell$. This is precisely the limit vector identified (for fixed coordinates) in the work of Diaconis and Tung: for each fixed finite set of indices $i_1, \ldots, i_m$, the vector

$$\big(a_{i_1}(\sigma), \ldots, a_{i_m}(\sigma)\big)$$

converges in distribution to $\big(A_{i_1}, \ldots, A_{i_m}\big)$ as $n \to \infty$.

The structure (4) makes explicit how the internal group Γ affects cycle counts. For example, if Γ is trivial (so $k = 1$ and only $\lambda = (1)$ occurs), then $a_1(\lambda) = 1$ and (4) reduces to $A_i = Z_{i,(1)}$ with $Z_{i,(1)} \sim \text{Poisson}(1/i)$, recovering the classical limit for uniform permutations in S_n. More generally, (4) shows that different coordinates (A_i) are typically dependent, since a single marked block cycle (ℓ, λ) may contribute simultaneously to several A_i through the various cycle lengths present in λ.

We shall use (4) as the target law for our growing-dimensional approximation. At the level of first moments, one may read off

$$\mathbb{E}A_i \;=\; \sum_{\ell | i} \frac{1}{\ell} \sum_{\lambda \vdash k} a_{i/\ell}(\lambda)\, \mathbb{P}_\Gamma(\lambda),$$

and similarly, joint generating functions can be expressed by conditioning on the independent Poisson family $(Z_{\ell,\lambda})$. The essential point for what follows is that the array $(Z_{\ell,\lambda})$ provides an underlying product structure, and (4) is a deterministic linear map from that array to the cycle-count coordinates. In the next section we introduce an explicit coupling that realizes the Poisson family via a marked spacing construction and allows us to control total variation on $(a_1(\sigma), \ldots, a_B(\sigma))$ when B grows with n.

3 The marked spacing model

Our approximation of truncated cycle counts is most conveniently organized through an explicit "marked spacing" construction. This model separates the randomness coming from the block permutation $\pi \in S_n$ from the randomness coming from the Γ-labels, and it realizes the array $(N_{\ell,\lambda}(\sigma))$ from (3) as a deterministic functional of a simple independent input: an inhomogeneous Bernoulli sequence together with i.i.d. marks.

3.1 Feller indicators and spacing events

Let $(\zeta_i)_{i \geq 1}$ be independent Bernoulli random variables with

$$\mathbb{P}(\zeta_i = 1) = \frac{1}{i}, \qquad i \geq 1.$$

For each n we form the finite indicator string $(\zeta_1, \ldots, \zeta_n)$ and adjoin a terminal 1 by setting

$$\zeta_{n+1} := 1.$$

Given an integer $\ell \geq 1$ and a starting position $1 \leq i \leq n+1-\ell$, we define the spacing event

$$E_{\ell,i}^{(n)} := \{\zeta_i = 1,\ \zeta_{i+1} = \cdots = \zeta_{i+\ell-1} = 0,\ \zeta_{i+\ell} = 1\}. \tag{5}$$

Thus $E_{\ell,i}^{(n)}$ records that, in the finite string with terminal 1, there is a gap of exactly $\ell - 1$ zeros between two consecutive ones, with the first one at position i.

We define the associated spacing counts by

$$C_\ell^{(n)} := \sum_{i=1}^{n+1-\ell} \mathbf{1}_{E_{\ell,i}^{(n)}}, \qquad 1 \leq \ell \leq n. \tag{6}$$

The random vector $(C_1^{(n)}, \ldots, C_n^{(n)})$ encodes the composition of n into the successive spacings between ones in the indicator sequence.

The relevance of (6) is the classical Feller coupling: if π is uniform on S_n, then its cycle count vector $(C_1(\pi), \ldots, C_n(\pi))$ has the same distribution as $(C_1^{(n)}, \ldots, C_n^{(n)})$. We will not reprove this fact here; we only use the consequence that one may regard the cycle structure of a uniform π as generated from independent Bernoulli indicators with parameters $1/i$ and a terminal 1 that closes the last cycle. In particular, for each ℓ the variable $C_\ell^{(n)}$ plays the role of the number of ℓ-cycles of the block permutation.

3.2 Adding i.i.d. marks with law $\mathbb{P}_\Gamma$

We now incorporate the Γ-labels through independent marks attached to the spacings. Let $(Y_i)_{i\geq1}$ be an i.i.d. sequence taking values in the set of partitions of k such that

$$\mathbb{P}(Y_i = \lambda) = \mathbb{P}_\Gamma(\lambda), \qquad \lambda \vdash k, \tag{7}$$

and assume $(Y_i)_{i\geq1}$ is independent of $(\zeta_i)_{i\geq1}$. We interpret Y_i as the cycle type of a "generic" uniform element of Γ and use it to mark the spacing that starts at i (when $\zeta_i = 1$).

For each $\ell \geq 1$ and $\lambda \vdash k$ we define the *marked spacing counts*

$$C_{\ell,\lambda}^{(n)} := \sum_{i=1}^{n+1-\ell} \mathbf{1}_{E_{\ell,i}^{(n)}} \mathbf{1}_{\{Y_i=\lambda\}}. \tag{8}$$

Conditional on the indicators $(\zeta_1, \ldots, \zeta_{n+1})$, the variables $(C_{\ell,\lambda}^{(n)})_{\lambda\vdash k}$ are obtained by independently marking each of the $C_\ell^{(n)}$ spacings of length ℓ according to the law (7); hence

$$\big(C_{\ell,\lambda}^{(n)}\big)_{\lambda\vdash k} \,\Big|\, (\zeta_1, \ldots, \zeta_{n+1}) \ \sim\ \text{Multinomial}\Big(C_\ell^{(n)};\ (\mathbb{P}_\Gamma(\lambda))_{\lambda\vdash k}\Big),$$

and the collections corresponding to distinct values of ℓ are conditionally independent given the indicators, since they depend on disjoint sets of spacing events.

This is the marked analogue of the thinning statement from Lemma 2.2. Indeed, in the wreath product model, if we condition on the block permutation π, then each block cycle c of length ℓ carries an internal product g_c which is uniform on Γ and independent across cycles, and the corresponding cycle type has law $\mathbb{P}_\Gamma(\cdot)$. Thus, the array $(N_{\ell,\lambda}(\sigma))$ may be produced by first sampling the cycle structure of π and then independently assigning a λ-mark to each ℓ-cycle with probabilities $\mathbb{P}_\Gamma(\lambda)$. The construction (8) implements precisely this mechanism using only independent primitives.

3.3 From marked spacings to global cycle counts

We now push the marked spacing counts through the deterministic compounding map coming from Lemma 2.1. For each $b \geq 1$ we define the *marked-spacing global cycle count*

$$C_b^{(n)} := \sum_{\lambda \vdash k} \sum_{\ell | b} a_{b/\ell}(\lambda)\, C_{\ell,\lambda}^{(n)}. \tag{9}$$

This is the direct analogue of (3), with $C_{\ell,\lambda}^{(n)}$ playing the role of $N_{\ell,\lambda}(\sigma)$. In particular, for any cutoff $B \leq n$ the vector $(C_1^{(n)}, \ldots, C_B^{(n)})$ is a measurable function of

$$\big(\zeta_1, \ldots, \zeta_{n+1}\big) \quad \text{and} \quad (Y_i)_{1 \leq i \leq n}.$$

We emphasize that (9) is a linear transformation of the marked spacing array, and the coefficients $a_{b/\ell}(\lambda)$ are purely combinatorial quantities determined by the partition λ.

The next statement records that the marked spacing model reproduces the wreath product cycle counts at the level of distribution.

Proposition 3.1 (Distributional representation via marked spacings)**.** *Let σ be uniform on $G_{k,n}$ and let $(a_b(\sigma))_{1 \leq b \leq kn}$ be its cycle counts in S_{kn}. Construct $(C_b^{(n)})_{1 \leq b \leq nk}$ from independent indicators (ζ_i) with $\mathbb{P}(\zeta_i = 1) = 1/i$ (with $\zeta_{n+1} = 1$) and independent marks (Y_i) with law* (7)*, according to* (5)*–*(9)*. Then for each $B \leq n$ we have*

$$\big(a_1(\sigma), \ldots, a_B(\sigma)\big) \overset{d}{=} \big(C_1^{(n)}, \ldots, C_B^{(n)}\big).$$

Proof sketch. By the Feller coupling, we may realize a uniform $\pi \in S_n$ together with a Bernoulli indicator sequence such that $(C_\ell(\pi))_{1 \leq \ell \leq n} \overset{d}{=} (C_\ell^{(n)})_{1 \leq \ell \leq n}$, where $C_\ell(\pi)$ is the number of ℓ-cycles of π. Conditional on π, Lemma 2.2 implies that the cycle products (g_c) over block cycles c are independent and uniform on Γ, hence their cycle types are i.i.d. with law $\mathbb{P}_\Gamma$. Therefore, conditional on π, the array $(N_{\ell,\lambda}(\sigma))$ is obtained by independently marking each of the $C_\ell(\pi)$ many ℓ-cycles with mark distribution $\mathbb{P}_\Gamma$, exactly as in (8). Finally, Lemma 2.1 gives the deterministic aggregation (3), which matches (9). Restricting to $b \le B$ yields the claim. □

Proposition 3.1 is the point at which we leave the group structure behind: all subsequent approximation statements can be proved by analyzing the marked spacing process (ζ_i, Y_i) and the functionals (8)–(9).

3.4 The infinite marked spacing process

For limit theorems it is natural to replace the finite string with terminal 1 by the infinite indicator sequence $(\zeta_i)_{i\ge1}$ without boundary conditions. For $\ell \ge 1$ and $i \ge 1$ we define the infinite spacing event

$$E_{\ell,i} \;:=\; \big\{\zeta_i = 1,\ \zeta_{i+1} = \cdots = \zeta_{i+\ell-1} = 0,\ \zeta_{i+\ell} = 1\big\}, \tag{10}$$

and the infinite marked spacing counts

$$C^{(\infty)}_{\ell,\lambda} \;:=\; \sum_{i\ge1} \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}}. \tag{11}$$

For each fixed ℓ the sum in (11) is almost surely finite on initial segments and defines a well-behaved counting variable; moreover, the collection over ℓ will admit an explicit Poisson limit description (after truncation), which we derive in the next section using Ignatov's theorem together with the independence of the marks.

In parallel with (9) we define the infinite marked-spacing global counts

$$C^{(\infty)}_b \;:=\; \sum_{\lambda\vdash k} \sum_{\ell\mid b} a_{b/\ell}(\lambda)\, C^{(\infty)}_{\ell,\lambda}, \qquad b \ge 1. \tag{12}$$

The vectors $(C^{(n)}_1, \ldots, C^{(n)}_B)$ and $(C^{(\infty)}_1, \ldots, C^{(\infty)}_B)$ can be coupled on a common probability space by using the same underlying sequences $(\zeta_i)_{i\ge1}$ and $(Y_i)_{i\ge1}$ and defining the finite object with the terminal modification $\zeta_{n+1} = 1$. The only discrepancy between the finite and infinite spacing structures comes from spacings that interact with this terminal closure; controlling the effect of those boundary interactions is precisely what will lead to the quantitative total variation bounds when $B = o(n)$.

The constructions above provide the promised explicit representation: the wreath product cycle counts up to level B are equal in law to a deterministic functional of independent inhomogeneous Bernoulli indicators and i.i.d. Γ-cycle-type marks. In the next section we analyze the infinite marked spacing process and identify its Poisson structure, from which the compound-Poisson target vector and its approximation properties will follow.

4 Poisson structure of the infinite marked spacing process

We now identify the law of the truncated infinite marked-spacing counts by combining Ignatov's theorem for the indicator sequence with independent marking. Throughout this section we work with the infinite construction (10)–(12) and fix an arbitrary cutoff $B \geq 1$.

4.1 Ignatov's theorem for unmarked spacing events

For $\ell \geq 1$ let

$$C_\ell^{(\infty)} := \sum_{i \geq 1} \mathbf{1}_{E_{\ell,i}}, \qquad E_{\ell,i} = \{\zeta_i = 1,\ \zeta_{i+1} = \cdots = \zeta_{i+\ell-1} = 0,\ \zeta_{i+\ell} = 1\},$$

so that $C_\ell^{(\infty)}$ counts gaps of length $\ell - 1$ between consecutive ones in the infinite indicator string.

A first calculation gives the correct intensity. Since the ζ_j are independent and $\mathbb{P}(\zeta_j = 1) = 1/j$, we have

$$\begin{aligned} \mathbb{P}(E_{\ell,i}) &= \frac{1}{i}\Big(\prod_{j=1}^{\ell-1}\Big(1 - \frac{1}{i+j}\Big)\Big)\frac{1}{i+\ell} = \frac{1}{i}\Big(\prod_{j=1}^{\ell-1}\frac{i+j-1}{i+j}\Big)\frac{1}{i+\ell} \\ &= \frac{1}{i} \cdot \frac{i}{i+\ell-1} \cdot \frac{1}{i+\ell} = \frac{1}{(i+\ell-1)(i+\ell)}. \end{aligned} \tag{13}$$

In particular,

$$\mathbb{E}C_\ell^{(\infty)} = \sum_{i \geq 1} \frac{1}{(i+\ell-1)(i+\ell)} = \frac{1}{\ell}\sum_{i \geq 1}\Big(\frac{1}{i+\ell-1} - \frac{1}{i+\ell}\Big) = \frac{1}{\ell}. \tag{14}$$

The point is that the full collection of spacing events enjoys a Poisson splitting property which is far stronger than what one would obtain from a naive Chen–Stein argument. The following form of Ignatov's theorem is standard in the analysis of the Feller indicator process (see, e.g., Arratia–Barbour–Tavaré for closely related formulations).

Theorem 4.1 (Ignatov)**.** *For each $\ell \geq 1$ define the random point measure on $\mathbb{N}$,*

$$\Xi_\ell := \sum_{i \geq 1} \mathbf{1}_{E_{\ell,i}}\, \delta_i.$$

Then the family $\{\Xi_\ell\}_{\ell \geq 1}$ is independent, and for each fixed ℓ the process Ξ_ℓ is a Poisson point process on $\mathbb{N}$ with intensity

$$\nu_\ell(\{i\}) = \mathbb{P}(E_{\ell,i}) = \frac{1}{(i+\ell-1)(i+\ell)}.$$

Consequently, for each ℓ the total mass $C_\ell^{(\infty)} = \Xi_\ell(\mathbb{N})$ is $\text{Poisson}(1/\ell)$*, and* $(C_1^{(\infty)}, \ldots, C_B^{(\infty)})$ *are independent.*

We will use only the final conclusion (independent Poisson counts), but it will be convenient to keep in mind the point-process statement, since marking is naturally expressed as a thinning operation on Ξ_ℓ.

4.2 Independent marking and Poisson thinning

Recall that the marks $(Y_i)_{i\geq 1}$ are i.i.d., independent of $(\zeta_i)_{i\geq 1}$, with $\mathbb{P}(Y_i = \lambda) = \mathbb{P}_\Gamma(\lambda)$ for $\lambda \vdash k$. For $\ell \geq 1$ and $\lambda \vdash k$ we defined

$$C_{\ell,\lambda}^{(\infty)} = \sum_{i\geq 1} \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}}.$$

Equivalently, if we define the marked point process

$$\Xi_{\ell,\lambda} := \sum_{i\geq 1} \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}} \, \delta_i,$$

then $C_{\ell,\lambda}^{(\infty)} = \Xi_{\ell,\lambda}(\mathbb{N})$ and $\Xi_\ell = \sum_{\lambda\vdash k} \Xi_{\ell,\lambda}$.

Lemma 4.2 (Marking of a Poisson process)**.** *Fix $\ell \geq 1$. Conditional on Ξ_ℓ, the processes $(\Xi_{\ell,\lambda})_{\lambda\vdash k}$ are obtained by independently assigning to each atom of Ξ_ℓ a label λ with probabilities $\mathbb{P}_\Gamma(\lambda)$. In particular, unconditionally the family $(\Xi_{\ell,\lambda})_{\lambda\vdash k}$ consists of independent Poisson point processes on $\mathbb{N}$ with intensities*

$$\nu_{\ell,\lambda}(\{i\}) = \nu_\ell(\{i\}) \, \mathbb{P}_\Gamma(\lambda) = \frac{\mathbb{P}_\Gamma(\lambda)}{(i+\ell-1)(i+\ell)}.$$

Consequently, for each ℓ the counts $(C_{\ell,\lambda}^{(\infty)})_{\lambda\vdash k}$ are independent Poisson with

$$C_{\ell,\lambda}^{(\infty)} \sim \text{Poisson}\Big(\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\Big),$$

and the collection $\{C_{\ell,\lambda}^{(\infty)} : 1 \leq \ell \leq B, \ \lambda \vdash k\}$ is jointly independent.

Proof. Given Ξ_ℓ, the indicator $\mathbf{1}_{\{Y_i=\lambda\}}$ attached to each potential atom location i is independent across i and independent of Ξ_ℓ; hence each atom of Ξ_ℓ is retained in $\Xi_{\ell,\lambda}$ with probability $\mathbb{P}_\Gamma(\lambda)$, independently across atoms. By the standard thinning property of Poisson point processes, $\Xi_{\ell,\lambda}$ is Poisson with intensity $\nu_\ell \mathbb{P}_\Gamma(\lambda)$ and the thinned processes are independent across λ. The mean of the total count is $\sum_{i\geq 1} \nu_{\ell,\lambda}(\{i\}) = \mathbb{P}_\Gamma(\lambda) \sum_{i\geq 1} \nu_\ell(\{i\}) = \mathbb{P}_\Gamma(\lambda)/\ell$ by (14). Independence across distinct ℓ follows from Theorem 4.1 and the independence of the marks. □

We now align the preceding lemma with the Poisson array used to define the target vector. For $1 \le \ell \le B$ and $\lambda \vdash k$ we set

$$Z_{\ell,\lambda} \;:=\; C^{(\infty)}_{\ell,\lambda}. \tag{15}$$

Then Lemma 4.2 yields exactly the distributional specification in the global setup: the $Z_{\ell,\lambda}$ are independent and

$$Z_{\ell,\lambda} \sim \mathrm{Poisson}\Big(\frac{1}{\ell}\mathbb{P}_\Gamma(\lambda)\Big).$$

4.3 The compound-Poisson vector as a deterministic compounding map

We finally translate the Poisson structure of the marked spacing array into the compound-Poisson structure of global cycle counts. Recall that the infinite marked-spacing global counts are

$$C^{(\infty)}_b = \sum_{\lambda\vdash k}\sum_{\ell\mid b} a_{b/\ell}(\lambda)\, C^{(\infty)}_{\ell,\lambda}, \qquad b \ge 1.$$

Since for fixed b there are only finitely many divisors $\ell \mid b$, the sum is almost surely finite and $C^{(\infty)}_b$ is a well-defined $\mathbb{N}$-valued random variable. Moreover, for a fixed cutoff B the vector $(C^{(\infty)}_1, \dots, C^{(\infty)}_B)$ depends only on $\{C^{(\infty)}_{\ell,\lambda} : 1 \le \ell \le B,\ \lambda \vdash k\}$, because if $b \le B$ then necessarily $\ell \mid b$ implies $\ell \le b \le B$.

Proposition 4.3 (Identification of the target law)**.** *For each $B \ge 1$ we have the distributional identity*

$$\big(C^{(\infty)}_1, \dots, C^{(\infty)}_B\big) \;\overset{d}{=}\; (A_1, \dots, A_B),$$

where $(A_b)_{b\ge1}$ is defined by

$$A_b = \sum_{\lambda\vdash k}\sum_{\ell\mid b} a_{b/\ell}(\lambda)\, Z_{\ell,\lambda}, \qquad Z_{\ell,\lambda} \sim \mathrm{Poisson}\Big(\frac{1}{\ell}\mathbb{P}_\Gamma(\lambda)\Big) \textit{ independent.}$$

Proof. By (15) we may take $Z_{\ell,\lambda} = C^{(\infty)}_{\ell,\lambda}$ for $1 \le \ell \le B$ and $\lambda \vdash k$, in which case the required independence and Poisson means are guaranteed by Lemma 4.2. Substituting into the defining relation for $C^{(\infty)}_b$ gives $C^{(\infty)}_b = A_b$ for each $b \le B$ on the same probability space, hence the claimed equality in distribution. □

It is sometimes convenient to record the associated Laplace transform, which makes the compound-Poisson nature explicit. For arbitrary parameters $(t_b)_{1\le b\le B} \subset \mathbb{R}$, Proposition 4.3 and independence of the $(Z_{\ell,\lambda})$ yield

$$\mathbb{E}\exp\Big(\sum_{b=1}^{B} t_b C^{(\infty)}_b\Big) = \exp\left(\sum_{\ell=1}^{B}\sum_{\lambda\vdash k}\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\left[\exp\Big(\sum_{\substack{1\le b\le B\\ \ell\mid b}} t_b\, a_{b/\ell}(\lambda)\Big) - 1\right]\right), \tag{16}$$

which is the Laplace functional of a finite-dimensional compound-Poisson vector obtained by superposing independent contributions indexed by (ℓ, λ).

Thus the limiting object appearing in our main approximation statement is nothing other than the truncated infinite marked spacing process pushed through the deterministic compounding map (12). In the next section we compare the finite marked spacing model $(C_1^{(n)}, \dots, C_B^{(n)})$ to its infinite analogue $(C_1^{(\infty)}, \dots, C_B^{(\infty)})$ by a direct coupling of the underlying indicator and mark sequences, isolating the boundary interactions introduced by the terminal closure at $n+1$.

5 Mesoscopic coupling of the finite and infinite marked spacing processes

We now couple, on a common probability space, the finite marked spacing counts appearing in the block Feller construction at level n with their infinite counterparts from Section 4. Fix integers $n \geq 1$ and $1 \leq B \leq n$. Throughout we work simultaneously with the infinite sequences $(\zeta_i)_{i\geq 1}$ and $(Y_i)_{i\geq 1}$ from the infinite construction, and we define the finite-n objects as deterministic functions of these sequences.

5.1 Coupling of indicator strings and marks

On the underlying space carrying $(\zeta_i, Y_i)_{i\geq 1}$, define the modified indicator string

$$\zeta_i^{(n)} := \begin{cases} \zeta_i, & 1 \leq i \leq n, \\ 1, & i = n+1, \end{cases}$$

and regard $(\zeta_1^{(n)}, \dots, \zeta_{n+1}^{(n)})$ as the finite string obtained by forcing a terminal 1 at time $n+1$. (We do not need to specify $\zeta_i^{(n)}$ for $i \geq n+2$.) We keep the same marks $Y_1, \dots, Y_n$ for the finite model; in particular, marks are never altered by the coupling.

For $\ell \geq 1$ and $1 \leq i \leq n$ we define the finite spacing event

$$E_{\ell,i}^{(n)} := \{\zeta_i^{(n)} = 1,\ \zeta_{i+1}^{(n)} = \dots = \zeta_{i+\ell-1}^{(n)} = 0,\ \zeta_{i+\ell}^{(n)} = 1\},$$

which is well-defined provided $i + \ell \leq n+1$; for definiteness we interpret $E_{\ell,i}^{(n)} = \varnothing$ if $i + \ell > n+1$. The marked finite counts are then

$$C_{\ell,\lambda}^{(n)} := \sum_{i=1}^{n} \mathbf{1}_{E_{\ell,i}^{(n)}} \mathbf{1}_{\{Y_i=\lambda\}}, \qquad C_\ell^{(n)} := \sum_{\lambda \vdash k} C_{\ell,\lambda}^{(n)}.$$

We emphasize that $C_{\ell,\lambda}^{(\infty)}$ from Section 4 is defined by the same formula with $E_{\ell,i}$ in place of $E_{\ell,i}^{(n)}$ and with the sum over all $i \geq 1$.

For later use it is convenient to introduce the *truncated infinite* counts which ignore starts after time n:

$$C^{(\infty,n)}_{\ell,\lambda} := \sum_{i=1}^{n} \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}}, \qquad C^{(\infty,n)}_{\ell} := \sum_{\lambda\vdash k} C^{(\infty,n)}_{\ell,\lambda}.$$

5.2 Interior agreement and localization of discrepancies

The key observation is that for cycle lengths up to B the dependence of $E_{\ell,i}$ on the indicator string is local, and away from the boundary the forced terminal value at $n+1$ is irrelevant.

Lemma 5.1 (Interior agreement). *If $1 \le \ell \le B$ and $1 \le i \le n-B$, then $E^{(n)}_{\ell,i} = E_{\ell,i}$. Consequently, for each $\lambda \vdash k$,*

$$\sum_{i=1}^{n-B} \mathbf{1}_{E^{(n)}_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}} = \sum_{i=1}^{n-B} \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}}.$$

Proof. If $i \le n-B$ and $\ell \le B$, then $i+\ell \le n$, so the events $E_{\ell,i}$ and $E^{(n)}_{\ell,i}$ depend only on $(\zeta_i,\dots,\zeta_{i+\ell})$, which coincide with $(\zeta^{(n)}_i,\dots,\zeta^{(n)}_{i+\ell})$ by construction. Multiplying by $\mathbf{1}_{\{Y_i=\lambda\}}$ preserves the equality. □

Thus, for $\ell \le B$, any discrepancy between the finite and infinite arrays must come from either (i) *tail starts* $i \ge n+1$ present in $C^{(\infty)}_{\ell,\lambda}$ but absent in $C^{(n)}_{\ell,\lambda}$, or (ii) *boundary starts* $i \in \{n-B+1,\dots,n\}$ where $E^{(n)}_{\ell,i}$ may be affected by the forced terminal 1 at $n+1$ and by the suppression of endpoints beyond $n+1$.

5.3 Tail spacings beyond n

We first control the contribution of spacing starts after time n in the infinite model. Since the marked counts are obtained by thinning with independent marks, it suffices to control the unmarked tail counts.

Lemma 5.2 (Tail bound). *For each $1 \le B \le n$,*

$$\mathbb{P}\Big(\exists\, 1 \le \ell \le B:\ C^{(\infty)}_{\ell} \ne C^{(\infty,n)}_{\ell}\Big) \le \sum_{\ell=1}^{B} \sum_{i=n+1}^{\infty} \mathbb{P}(E_{\ell,i}) = \sum_{\ell=1}^{B} \frac{1}{n+\ell} \le \frac{B}{n}.$$

The same bound holds with $C_{\ell,\lambda}$ in place of C_ℓ.

Proof. If $C^{(\infty)}_{\ell} \ne C^{(\infty,n)}_{\ell}$, then there exists $i \ge n+1$ with $E_{\ell,i}$, hence by a union bound and (13),

$$\mathbb{P}\big(C^{(\infty)}_{\ell} \ne C^{(\infty,n)}_{\ell}\big) \le \sum_{i=n+1}^{\infty} \frac{1}{(i+\ell-1)(i+\ell)} = \frac{1}{n+\ell},$$

where the last identity is the telescoping sum

$$\sum_{i=n+1}^{\infty}\Big(\frac{1}{i+\ell-1}-\frac{1}{i+\ell}\Big)=\frac{1}{n+\ell}.$$

Summing over $\ell \le B$ yields the stated estimate, and the marked version follows since $\{C^{(\infty)}_{\ell,\lambda} \neq C^{(\infty,n)}_{\ell,\lambda}\} \subseteq \{C^{(\infty)}_{\ell} \neq C^{(\infty,n)}_{\ell}\}$. □

5.4 Boundary spacings and terminal closure at $n+1$

We next compare the finite counts to the truncated infinite counts, i.e., we bound the effect of replacing ζ_{n+1} by 1 and suppressing endpoints beyond $n+1$. For $1 \le \ell \le B$, there are two relevant mechanisms:

1. *Terminal creation:* the finite model may create a spacing ending at $n+1$ even when $\zeta_{n+1}=0$ in the infinite model;
2. *Overshoot:* the infinite model may have a spacing of length $\ell \le B$ which starts near n but whose endpoint lies in $\{n+2,\dots,n+B\}$, which is necessarily invisible to the finite model (and is instead truncated by the terminal 1).

Both are supported on the boundary window $\{n-B+1,\dots,n\}$ and can be bounded explicitly.

Lemma 5.3 (Boundary bound). *For each $1 \le B \le n$,*

$$\begin{aligned}\mathbb{P}\Big(\exists\, 1 \le \ell \le B:\ C^{(n)}_{\ell} \neq C^{(\infty,n)}_{\ell}\Big) &\le \sum_{\ell=1}^{B}\mathbb{P}\big(E^{(n)}_{\ell,n+1-\ell}\cap\{\zeta_{n+1}=0\}\big)+\sum_{\ell=1}^{B}\sum_{i=n+2-\ell}^{n}\mathbb{P}(E_{\ell,i})\\ &= \frac{B}{n+1}+\sum_{\ell=1}^{B}\sum_{i=n+2-\ell}^{n}\frac{1}{(i+\ell-1)(i+\ell)}\\ &\le \frac{B}{n+1}+\frac{B(B-1)}{2(n+1)(n+2)}\\ &\le \frac{B}{n}+\frac{B^2}{n^2}. \qquad (17)\end{aligned}$$

The same bound holds with $C_{\ell,\lambda}$ in place of C_ℓ.

Proof. For the first term, note that for fixed ℓ the event $E^{(n)}_{\ell,n+1-\ell}$ is exactly the event that there is a 1 at time $n+1-\ell$ and then zeros up to time n; intersecting with $\{\zeta_{n+1}=0\}$ expresses that the finite terminal 1 creates an endpoint at $n+1$ not present in the infinite string. By independence and the same telescoping product as in (13), we compute

$$\mathbb{P}\big(E^{(n)}_{\ell,n+1-\ell}\cap\{\zeta_{n+1}=0\}\big)=\frac{1}{n+1-\ell}\Big(\prod_{j=n+2-\ell}^{n}\Big(1-\frac{1}{j}\Big)\Big)\Big(1-\frac{1}{n+1}\Big)=\frac{1}{n+1}.$$

Summing over $\ell \le B$ gives $B/(n+1)$.

For the overshoot term, if $E_{\ell,i}$ occurs with $i \le n$ and $i+\ell \ge n+2$, then necessarily $i \in \{n+2-\ell,\dots,n\}$; on this event the infinite spacing of length ℓ uses at least one indicator beyond $n+1$, and the finite construction cannot reproduce it. Thus a union bound yields the stated double sum. Finally, if $i \in \{n+2-\ell,\dots,n\}$ then $i+\ell-1 \ge n+1$ and $i+\ell \ge n+2$, so each summand satisfies

$$\mathbb{P}(E_{\ell,i}) = \frac{1}{(i+\ell-1)(i+\ell)} \le \frac{1}{(n+1)(n+2)}.$$

There are $\ell-1$ possible values of i for a given ℓ, hence the double sum is at most

$$\sum_{\ell=1}^{B} \frac{\ell-1}{(n+1)(n+2)} = \frac{B(B-1)}{2(n+1)(n+2)}.$$

The marked bound follows as in Lemma 5.2. □

5.5 A single discrepancy event for lengths $\le B$

Combining the tail and boundary contributions, we define the event

$$\mathcal{D}_{B,n} := \Big\{\exists\, 1 \le \ell \le B,\ \exists\, \lambda \vdash k :\ C^{(n)}_{\ell,\lambda} \ne C^{(\infty)}_{\ell,\lambda}\Big\}.$$

By Lemmas 5.2 and 5.3 and a further union bound,

$$\mathbb{P}(\mathcal{D}_{B,n}) \le \sum_{\ell=1}^{B} \frac{1}{n+\ell} + \frac{B}{n+1} + \frac{B(B-1)}{2(n+1)(n+2)} \le \frac{2B}{n} + \frac{B^2}{n^2}. \qquad (18)$$

In particular, if $B = o(n)$ then $\mathbb{P}(\mathcal{D}_{B,n}) \to 0$.

Finally, recall that the compounded global counts $(C_b^{(n)})_{1\le b\le B}$ and $(C_b^{(\infty)})_{1\le b\le B}$ are deterministic functions of the marked array $\{C_{\ell,\lambda} : 1 \le \ell \le B,\ \lambda \vdash k\}$ via the compounding map $\sum_\lambda \sum_{\ell|b} a_{b/\ell}(\lambda)\,\cdot$. Therefore, on the complement of $\mathcal{D}_{B,n}$ we have the simultaneous equalities

$$(C_1^{(n)},\dots,C_B^{(n)}) = (C_1^{(\infty)},\dots,C_B^{(\infty)}).$$

In the next section we convert the coupling estimate (18) into total variation bounds for the truncated cycle count vector by combining the coupling inequality with stability properties of the compounding map and the bounded effect of marks.

5.6 From spacing discrepancies to total variation bounds

For $1 \le b \le B$ we recall the compounding map from marked block spacings to point-level cycle counts. Given an array $\{x_{\ell,\lambda} : 1 \le \ell \le B,\ \lambda \vdash k\}$, set

$$\Phi_b(x) := \sum_{\lambda \vdash k} \sum_{\ell \mid b} a_{b/\ell}(\lambda)\, x_{\ell,\lambda}, \qquad \Phi_B(x) := (\Phi_1(x), \dots, \Phi_B(x)) \in \mathbb{N}^B. \tag{19}$$

Thus $x_{\ell,\lambda}$ counts ℓ-cycles at the block level whose associated cycle-product in Γ has type λ, and each such block ℓ-cycle contributes $a_{b/\ell}(\lambda)$ many b-cycles at the point level.

By definition of $C^{(n)}_{\ell,\lambda}$ and $C^{(\infty)}_{\ell,\lambda}$ we may therefore write

$$(C^{(n)}_1, \dots, C^{(n)}_B) = \Phi_B\big((C^{(n)}_{\ell,\lambda})_{1 \le \ell \le B,\ \lambda \vdash k}\big), \qquad (C^{(\infty)}_1, \dots, C^{(\infty)}_B) = \Phi_B\big((C^{(\infty)}_{\ell,\lambda})_{1 \le \ell \le B,\ \lambda \vdash k}\big). \tag{20}$$

We emphasize that Φ_B is a deterministic map and, crucially for what follows, our coupling never alters the marks (Y_i): all discrepancies between the finite and infinite marked arrays arise solely from the modification of the indicator string at the terminal time $n+1$ and from the truncation of starts after time n.

We now connect these constructions back to the cycle counts of the uniform group element $\sigma \in G_{k,n}$. From the block Feller description of a uniform element of $G_{k,n}$ (developed in the previous sections), we have the distributional identity

$$\mathcal{L}\big(a_1(\sigma), \dots, a_B(\sigma)\big) = \mathcal{L}\big(C^{(n)}_1, \dots, C^{(n)}_B\big). \tag{21}$$

Similarly, from the infinite marked spacing construction and Ignatov's theorem as applied in Section 4, the marked counts $\{C^{(\infty)}_{\ell,\lambda}\}_{\ell \ge 1,\ \lambda \vdash k}$ are independent Poisson with means $\mathbb{P}_\Gamma(\lambda)/\ell$, and the compounded vector agrees with the target vector $(A_b)_{b \ge 1}$:

$$\mathcal{L}\big(C^{(\infty)}_1, \dots, C^{(\infty)}_B\big) = \mathcal{L}\big(A_1, \dots, A_B\big). \tag{22}$$

Combining (21) and (22), it suffices to bound the total variation distance between the finite and infinite compounded counts.

We apply the standard coupling inequality: if (X, Y) are random elements of a countable space on a common probability space, then

$$\big\|\mathcal{L}(X) - \mathcal{L}(Y)\big\|_{\mathrm{TV}} \le \mathbb{P}(X \ne Y). \tag{23}$$

In our setting the common probability space is the one carrying $(\zeta_i, Y_i)_{i \ge 1}$, and the coupled objects are the vectors $\Phi_B(C^{(n)}_{\ell,\lambda})$ and $\Phi_B(C^{(\infty)}_{\ell,\lambda})$.

Proposition 5.4 (TV bound from the discrepancy event)**.** *For each $n \geq 1$ and $1 \leq B \leq n$,*

$$\big\|\mathcal{L}(C_1^{(n)}, \dots, C_B^{(n)}) - \mathcal{L}(C_1^{(\infty)}, \dots, C_B^{(\infty)})\big\|_{\mathrm{TV}} \leq \mathbb{P}(\mathcal{D}_{B,n}),$$

and hence, by (18)*,*

$$\big\|\mathcal{L}(C_1^{(n)}, \dots, C_B^{(n)}) - \mathcal{L}(C_1^{(\infty)}, \dots, C_B^{(\infty)})\big\|_{\mathrm{TV}} \leq \frac{2B}{n} + \frac{B^2}{n^2}. \qquad (24)$$

Consequently,

$$\big\|\mathcal{L}(a_1(\sigma), \dots, a_B(\sigma)) - \mathcal{L}(A_1, \dots, A_B)\big\|_{\mathrm{TV}} \leq \frac{2B}{n} + \frac{B^2}{n^2}. \qquad (25)$$

In particular, if $B = o(n)$ then the left-hand side of (25) *tends to* 0 *as $n \to \infty$.*

Proof. By construction,

$$(C_1^{(n)}, \dots, C_B^{(n)}) = \Phi_B\big((C_{\ell,\lambda}^{(n)})_{1 \leq \ell \leq B,\ \lambda \vdash k}\big), \qquad (C_1^{(\infty)}, \dots, C_B^{(\infty)}) = \Phi_B\big((C_{\ell,\lambda}^{(\infty)})_{1 \leq \ell \leq B,\ \lambda \vdash k}\big).$$

On the complement $\mathcal{D}_{B,n}^c$ we have, by definition of $\mathcal{D}_{B,n}$, the coordinatewise equalities

$$C_{\ell,\lambda}^{(n)} = C_{\ell,\lambda}^{(\infty)} \qquad \text{for all } 1 \leq \ell \leq B \text{ and all } \lambda \vdash k.$$

Applying the deterministic map Φ_B yields

$$(C_1^{(n)}, \dots, C_B^{(n)}) = (C_1^{(\infty)}, \dots, C_B^{(\infty)}) \qquad \text{on } \mathcal{D}_{B,n}^c.$$

Therefore,

$$\mathbb{P}\big((C_1^{(n)}, \dots, C_B^{(n)}) \neq (C_1^{(\infty)}, \dots, C_B^{(\infty)})\big) \leq \mathbb{P}(\mathcal{D}_{B,n}),$$

and the coupling inequality (23) gives the first claim. The bound (24) then follows from (18). Finally, (25) follows by replacing $(C_1^{(n)}, \dots, C_B^{(n)})$ and $(C_1^{(\infty)}, \dots, C_B^{(\infty)})$ with $(a_1(\sigma), \dots, a_B(\sigma))$ and $(A_1, \dots, A_B)$ using (21) and (22). □

We record two remarks which will be useful when we pass from total variation to expectation bounds in the next section.

First, the estimate (25) is obtained without any amplification factor coming from the number of cycle types $\lambda \vdash k$. This is because we have coupled the marked arrays directly: the marks are shared between the finite and infinite constructions, so the event $\mathcal{D}_{B,n}$ is controlled entirely by the boundary/tail behavior of the indicator string (ζ_i), and Lemmas 5.2–5.3 already capture the relevant error probabilities.

Second, although we have used only the fact that Φ_B is deterministic, it is helpful to note that Φ_B is uniformly bounded in the following sense: for each fixed $\lambda \vdash k$ and $\ell \geq 1$, one block ℓ-cycle with mark λ contributes exactly k points, hence

$$\sum_{b \geq 1} b\, a_{b/\ell}(\lambda)\, \mathbf{1}_{\{\ell | b\}} = k\ell. \tag{26}$$

In particular, for any two marked arrays x, x' supported on $\{1, \ldots, B\} \times \{\lambda \vdash k\}$, a single unit change in some coordinate $x_{\ell,\lambda}$ can only create changes in the compounded vector along the multiples of ℓ, and the total number of newly created point-cycles is at most k. While we do not need (26) for Proposition 5.4, it provides a convenient bookkeeping device when estimating differences of test functionals of $(C_b^{(n)})_{b \leq B}$ and $(C_b^{(\infty)})_{b \leq B}$ under couplings that may not force exact equality.

We now proceed to strengthen (25) to bounds for expectations of test functionals via Stein-type estimates and a dependency-graph formulation.

5.7 Expectation bounds for Lipschitz test functionals

The total variation estimate obtained above is well suited for bounded test functions, but it does not directly yield quantitative control for unbounded functionals of the truncated cycle count vector. In this subsection we record a strengthening which suffices for the class of observables that will typically arise in applications: Lipschitz functionals with at most polynomial growth.

Let $B \leq n$ and write $x = (x_1, \ldots, x_B) \in \mathbb{N}^B$. We say that $F : \mathbb{N}^B \to \mathbb{R}$ belongs to $\mathrm{LipPoly}(L, m)$ if there exist constants $L \geq 0$ and an integer $m \geq 0$ such that for all $x, y \in \mathbb{N}^B$,

$$|F(x) - F(y)| \leq L\Big(1 + \|x\|_1 + \|y\|_1\Big)^m \|x - y\|_1, \qquad \|x\|_1 := \sum_{b=1}^{B} x_b. \tag{27}$$

The case $m = 0$ corresponds to ordinary ℓ_1-Lipschitz functionals. (One may also work with weighted ℓ_1 metrics, and then (26) becomes the natural bookkeeping device; we keep (27) for concreteness.)

Proposition 5.5 (Expectation bound)**.** *Fix $k \geq 1$ and $\Gamma \leq S_k$. There exists a finite constant $C_0 = C_0(k, \Gamma)$ such that the following holds for all $n \geq 1$ and $1 \leq B \leq n$. For every $F \in \mathrm{LipPoly}(L, m)$,*

$$\Big|\mathbb{E}F\big(a_1(\sigma), \ldots, a_B(\sigma)\big) - \mathbb{E}F\big(A_1, \ldots, A_B\big)\Big| \leq C_0\, L\,\Big(1 + \log(B+1)\Big)^m \frac{B}{n}. \tag{28}$$

In particular, if $B = o(n)$ then the difference of expectations tends to 0 for each fixed $F \in \mathrm{LipPoly}(L, m)$.

Proof. By (21) and (22) it suffices to compare $(C_b^{(n)})_{b\le B}$ and $(C_b^{(\infty)})_{b\le B}$ under the coupling on the common probability space carrying $(\zeta_i, Y_i)_{i\ge1}$. Define the ℓ_1 discrepancy

$$\Delta_{B,n} := \sum_{b=1}^{B} \big|C_b^{(n)} - C_b^{(\infty)}\big|.$$

Then, using (27) and the fact that $C_b^{(n)} = C_b^{(\infty)}$ on $\mathcal{D}_{B,n}^c$ (as in the proof of Proposition 5.4), we obtain

$$\begin{aligned}\Big|\mathbb{E}F(C_1^{(n)},\dots,C_B^{(n)}) - \mathbb{E}F(C_1^{(\infty)},\dots,C_B^{(\infty)})\Big| &\le \mathbb{E}\Big[|F(C^{(n)}) - F(C^{(\infty)})|\mathbf{1}_{\mathcal{D}_{B,n}}\Big]\\ &\le L\,\mathbb{E}\Big[\Big(1 + \|C^{(n)}\|_1 + \|C^{(\infty)}\|_1\Big)^m \Delta_{B,n}\Big].\end{aligned}$$

Applying Cauchy–Schwarz yields

$$\le L\left(\mathbb{E}\big(1 + \|C^{(n)}\|_1 + \|C^{(\infty)}\|_1\big)^{2m}\right)^{1/2} \cdot \big(\mathbb{E}\Delta_{B,n}^2\big)^{1/2}. \tag{29}$$

We now bound the two factors in (29). For the moment term, observe that

$$\|C^{(\infty)}\|_1 = \sum_{b=1}^{B} C_b^{(\infty)} \le \sum_{\ell=1}^{B} \sum_{\lambda\vdash k} C_{\ell,\lambda}^{(\infty)} \sum_{j\ge1} a_j(\lambda).$$

Since $\sum_{j\ge1} a_j(\lambda)$ equals the number of cycles of a permutation of type λ, we have the uniform bound $\sum_{j\ge1} a_j(\lambda) \le k$. Therefore

$$\|C^{(\infty)}\|_1 \le k \sum_{\ell=1}^{B} \sum_{\lambda\vdash k} C_{\ell,\lambda}^{(\infty)}.$$

By the Poisson description of the marked infinite model, the random variables $\{C_{\ell,\lambda}^{(\infty)}\}$ are independent Poisson with means $\mathbb{P}_\Gamma(\lambda)/\ell$, hence

$$\sum_{\ell=1}^{B} \sum_{\lambda\vdash k} C_{\ell,\lambda}^{(\infty)} \sim \mathrm{Poisson}\Big(\sum_{\ell=1}^{B} \frac{1}{\ell} \sum_{\lambda\vdash k} \mathbb{P}_\Gamma(\lambda)\Big) = \mathrm{Poisson}(H_B), \qquad H_B := \sum_{\ell=1}^{B} \frac{1}{\ell}.$$

Consequently, for each fixed m we have

$$\mathbb{E}\big(1 + \|C^{(\infty)}\|_1\big)^{2m} \le C_m(k)\,(1 + H_B)^{2m} \le C_m(k)\,\big(1 + \log(B+1)\big)^{2m}, \tag{30}$$

with $C_m(k) < \infty$ explicit (for instance, via standard polynomial moment bounds for Poisson variables). The same bound holds for $\|C^{(n)}\|_1$ uniformly in n, since $(C_b^{(n)})_{b\le B}$ is obtained from the same indicators and marks as $(C_b^{(\infty)})_{b\le B}$ but with only finitely many starts retained; in particular $\|C^{(n)}\|_1 \le \|C^{(\infty)}\|_1 + k\,\Delta'_{B,n}$ for an appropriate marked discrepancy count

$\Delta'_{B,n}$, and the latter has uniformly bounded moments of all fixed orders by the tail estimates below.

For the discrepancy term, we use that all differences between $C^{(n)}_{\ell,\lambda}$ and $C^{(\infty)}_{\ell,\lambda}$ arise from (i) truncating starts after time n and (ii) modifying the terminal symbol at time $n+1$. Let

$$T_{B,n} := \sum_{\ell=1}^{B} \sum_{\lambda \vdash k} \sum_{i>n} \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i=\lambda\}}$$

be the number of marked starts of length at most B occurring strictly after time n in the infinite construction. Similarly, let $U_{B,n}$ count the (at most one) boundary perturbation due to the terminal modification; the exact definition is immaterial, only that $U_{B,n}$ is supported on an event of probability $O(B^2/n^2)$ (cf. (18)). Then

$$\Delta_{B,n} \le k\,(T_{B,n} + U_{B,n}), \tag{31}$$

because each discrepant block start can alter the point-level cycle count vector by at most the total number of cycles contributed by its Γ-mark, which is bounded by k.

It remains to estimate $\mathbb{E}T_{B,n}$ and $\mathbb{E}U_{B,n}$. For the tail term we compute, using independence of the marks and the standard formula for the start probability in the Feller indicator sequence,

$$\mathbb{P}(E_{\ell,i}) = \mathbb{P}(\zeta_i = 1, \zeta_{i+1} = \cdots = \zeta_{i+\ell-1} = 0, \zeta_{i+\ell} = 1) = \frac{1}{(i+\ell-1)(i+\ell)}.$$

Therefore

$$\begin{aligned}\mathbb{E}T_{B,n} &= \sum_{\ell=1}^{B} \sum_{\lambda \vdash k} \sum_{i>n} \mathbb{P}(E_{\ell,i})\, \mathbb{P}_\Gamma(\lambda) = \sum_{\ell=1}^{B} \sum_{i>n} \frac{1}{(i+\ell-1)(i+\ell)} \\ &= \sum_{\ell=1}^{B} \frac{1}{n+\ell} \le \frac{B}{n+1}.\end{aligned}$$

Moreover $T_{B,n}$ is a sum of indicators with rapidly decaying means, hence $\mathbb{E}T_{B,n}^2 \le C\,B/n$ for an absolute constant C (e.g. by expanding the square and bounding pairwise correlations by $0 \le \mathbb{P}(E_{\ell,i} \cap E_{\ell',i'}) \le \mathbb{P}(E_{\ell,i})$). For the boundary term, since $U_{B,n}$ is supported on $\mathcal{D}_{B,n}$ and counts at most B potential starts near n, we have $\mathbb{E}U_{B,n}^2 \le C' B^2\, \mathbb{P}(\mathcal{D}_{B,n}) \le C''(B^3/n + B^4/n^2)$, which is negligible compared to B/n for $B \le n$ after taking square roots. Combining these bounds with (31) yields

$$\left(\mathbb{E}\Delta_{B,n}^2\right)^{1/2} \le C_1(k) \left(\frac{B}{n}\right)^{1/2} + C_2(k)\,\frac{B}{n}. \tag{32}$$

Substituting (30) and (32) into (29) and using $B \le n$ completes the proof of (28) (after possibly enlarging the constant to absorb the square-root term into $C_0 B/n$). □

5.8 A dependency-graph formulation (Stein–Chen viewpoint)

For completeness we also record a formulation in which the above expectation bounds arise from a standard Stein–Chen argument applied at the marked-spacing level. Fix $B \le n$ and define, for $1 \le \ell \le B$, $\lambda \vdash k$, and $1 \le i \le n$,

$$\xi_{i,\ell,\lambda} := \mathbf{1}_{E_{\ell,i}} \mathbf{1}_{\{Y_i = \lambda\}}, \qquad X_{\ell,\lambda} := \sum_{i=1}^{n} \xi_{i,\ell,\lambda} = C_{\ell,\lambda}^{(n)}.$$

Each $\xi_{i,\ell,\lambda}$ depends only on the finite set of indicators $\zeta_i, \ldots, \zeta_{i+\ell}$ and on the single mark Y_i. Hence if we place an edge between two vertices (i, ℓ, λ) and (i', ℓ', λ') whenever the corresponding dependency windows overlap, i.e.

$$[i, i+\ell] \cap [i', i'+\ell'] \neq \varnothing,$$

then this defines a dependency graph in the sense of Arratia–Goldstein–Gordon. The maximal neighborhood size is $O(B)$ uniformly in n, and the vertex weights satisfy

$$p_{i,\ell,\lambda} := \mathbb{E}\xi_{i,\ell,\lambda} = \mathbb{P}_\Gamma(\lambda) \frac{1}{(i+\ell-1)(i+\ell)} \le \frac{1}{i^2}.$$

Consequently, the Stein–Chen bounds for Poisson approximation of locally dependent indicator sums yield explicit control (in Wasserstein distance for Lipschitz test functions, or in total variation for the full vector when combined with thinning arguments) of the law of $(X_{\ell,\lambda})_{1 \le \ell \le B,\ \lambda \vdash k}$ by a product of independent Poisson laws with the same means. The relevant error terms are sums of the form

$$\sum_{i,\ell,\lambda} p_{i,\ell,\lambda}^2 \qquad \text{and} \qquad \sum_{(i,\ell,\lambda) \sim (i',\ell',\lambda')} p_{i,\ell,\lambda} \, p_{i',\ell',\lambda'},$$

both of which are $O(B/n)$ after summing over i and using $p_{i,\ell,\lambda} \le i^{-2}$ and the fact that each vertex has $O(B)$ neighbors. Pushing forward by the deterministic compounding map Φ_B then yields expectation bounds for the point-level vector $(C_b^{(n)})_{b \le B}$, and hence for $(a_b(\sigma))_{b \le B}$, with constants depending only on (k, Γ) and on the Lipschitz/polynomial parameters of the test functional. We do not pursue the sharpest constants here, since the coupling argument already provides the correct B/n scaling and keeps the dependence on the mark space $\{\lambda \vdash k\}$ completely explicit.

In the next section we illustrate the limiting compound-Poisson structure and the B/n error scaling on several concrete choices of Γ.

6 Worked examples and numerics

We illustrate the limiting vector $(A_i)_{i \ge 1}$, and in particular the dependence structure across different coordinates, on several concrete choices of Γ. Throughout, recall that the independent Poisson building blocks are $\{Z_{\ell,\lambda}\}_{\ell \ge 1,\ \lambda \vdash k}$

with

$$Z_{\ell,\lambda} \sim \mathrm{Poisson}\left(\frac{1}{\ell}\,\mathbb{P}_\Gamma(\lambda)\right), \qquad A_i = \sum_{\lambda\vdash k}\sum_{\ell\mid i} a_{i/\ell}(\lambda)\, Z_{\ell,\lambda}.$$

The dependence between different A_i's is completely determined by whether the same Poisson component $Z_{\ell,\lambda}$ can contribute to more than one coordinate, i.e. whether the mark λ contains cycles of more than one length.

6.1 The cyclic subgroup $\Gamma = C_k$ (independent limits)

Let $\Gamma = C_k \le S_k$ be the cyclic group generated by a k-cycle $\tau = (1\,2\,\cdots\,k)$. Every element of C_k is a power τ^r $(0 \le r \le k-1)$, and its cycle structure is determined by $d = \gcd(k, r)$: namely τ^r is a product of d disjoint cycles each of length k/d. Equivalently, for each divisor $d \mid k$ we have a partition

$$\lambda_d := \left(\frac{k}{d}\right)^d, \qquad a_{k/d}(\lambda_d) = d, \quad a_j(\lambda_d) = 0 \;\; (j \ne k/d),$$

and

$$\mathbb{P}_\Gamma(\lambda_d) = \frac{1}{k}\,\#\{0 \le r \le k-1 : \gcd(k,r) = d\} = \frac{1}{k}\,\varphi\left(\frac{k}{d}\right),$$

with φ the Euler totient function (including the case $d = k$ giving the identity).

For a fixed $\ell \ge 1$, a block-cycle of length ℓ marked by λ_d contributes *only* to the point-level cycle length $i = \ell \cdot (k/d)$, and it contributes exactly d cycles of that length. Consequently,

$$A_i = \sum_{\substack{d\mid k\\ (k/d)\mid i}} d\, Z_{i/(k/d),\,\lambda_d}. \tag{33}$$

In particular, each Poisson variable Z_{ℓ,λ_d} appears in exactly one coordinate A_i (namely $i = \ell \cdot k/d$). Since the family $\{Z_{\ell,\lambda}\}$ is independent, the random variables $\{A_i\}_{i\ge 1}$ are independent as well. This is the simplest situation: the limiting law is coordinatewise (compound) Poisson with no cross-correlations.

The means can be read off from (33):

$$\mathbb{E}A_i = \sum_{\substack{d\mid k\\ (k/d)\mid i}} d \cdot \frac{1}{i/(k/d)} \cdot \frac{1}{k}\,\varphi\left(\frac{k}{d}\right) = \frac{1}{i}\sum_{\substack{d\mid k\\ (k/d)\mid i}} d\,\varphi\left(\frac{k}{d}\right)\frac{k}{d} = \frac{1}{i}\sum_{\substack{d\mid k\\ (k/d)\mid i}} k\,\varphi\left(\frac{k}{d}\right). \tag{34}$$

In the special case where Γ is the order-k subgroup generated by a k-cycle and we additionally restrict to the single mark $\lambda_1 = (k)$ (for instance, if we condition the mark to be a k-cycle), then $A_{k\ell} \sim \mathrm{Poisson}(1/\ell)$ and $A_i \equiv 0$ for $k \nmid i$, recovering a particularly transparent limit.

6.2 The case $\Gamma = S_3$ (dependence via shared Poisson components)

Let $k = 3$ and $\Gamma = S_3$. There are three conjugacy classes (hence three possible marks):

$$\lambda^{(1)} = (1,1,1), \qquad \lambda^{(2)} = (2,1), \qquad \lambda^{(3)} = (3),$$

with probabilities $\mathbb{P}_\Gamma(\lambda^{(1)}) = 1/6$, $\mathbb{P}_\Gamma(\lambda^{(2)}) = 1/2$, and $\mathbb{P}_\Gamma(\lambda^{(3)}) = 1/3$. The corresponding cycle counts are

$$(a_1, a_2, a_3)(\lambda^{(1)}) = (3,0,0), \quad (a_1, a_2, a_3)(\lambda^{(2)}) = (1,1,0), \quad (a_1, a_2, a_3)(\lambda^{(3)}) = (0,0,1).$$

Writing $Z_{\ell,1}, Z_{\ell,2}, Z_{\ell,3}$ for $Z_{\ell,\lambda^{(1)}}, Z_{\ell,\lambda^{(2)}}, Z_{\ell,\lambda^{(3)}}$, we may compute the first few coordinates explicitly:

$$\begin{aligned}
A_1 &= 3Z_{1,1} + Z_{1,2}, \\
A_2 &= \big(\text{contribution from } \ell = 1\big) + \big(\text{contribution from } \ell = 2\big) = Z_{1,2} + 3Z_{2,1} + Z_{2,2}, \\
A_3 &= Z_{1,3} + 3Z_{3,1} + Z_{3,2}, \\
A_4 &= Z_{2,2} + 3Z_{4,1} + Z_{4,2}, \qquad \text{etc.}
\end{aligned}$$

The key point is that the transposition mark $\lambda^{(2)} = (2,1)$ contains *two* cycle lengths (a 2-cycle and a fixed point). Hence the same Poisson component $Z_{\ell,2}$ contributes simultaneously to A_ℓ (via its 1-cycle) and to $A_{2\ell}$ (via its 2-cycle). This produces genuine dependence between coordinates. For instance, A_1 and A_2 share $Z_{1,2}$ and therefore

$$\mathrm{Cov}(A_1, A_2) = \mathrm{Var}(Z_{1,2}) = \frac{1}{1}\,\mathbb{P}_\Gamma(\lambda^{(2)}) = \frac{1}{2},$$

whereas A_1 and A_3 share no Poisson components and are independent.

More generally, for any $\ell \geq 1$ the pair $(A_\ell, A_{2\ell})$ shares $Z_{\ell,2}$, giving

$$\mathrm{Cov}(A_\ell, A_{2\ell}) = a_1(\lambda^{(2)})\, a_2(\lambda^{(2)})\, \mathrm{Var}(Z_{\ell,2}) = 1 \cdot 1 \cdot \frac{1}{\ell} \cdot \frac{1}{2} = \frac{1}{2\ell}. \tag{35}$$

No other pair of distinct coordinates shares a Poisson component, because $\lambda^{(1)}$ has only 1-cycles and $\lambda^{(3)}$ has only 3-cycles, while $\lambda^{(2)}$ links precisely the lengths ℓ and 2ℓ. Thus the dependence graph of the limiting point-level cycle counts is sparse and completely explicit.

This example is also useful for sanity checks on scaling. The means follow immediately:

$$\mathbb{E}A_\ell = 3 \cdot \frac{1}{\ell} \cdot \frac{1}{6} + 1 \cdot \frac{1}{\ell} \cdot \frac{1}{2} = \frac{1}{2\ell} + \frac{1}{2\ell} = \frac{1}{\ell}, \qquad \mathbb{E}A_{2\ell} = \frac{1}{2\ell} + \frac{1}{4\ell} + \frac{1}{4\ell} = \frac{1}{\ell},$$

and similarly one checks $\mathbb{E}A_i = 1/i$ for $1 \leq i \leq 3$ and slight deviations beyond the intrinsic cutoff imposed by $k = 3$, consistent with the interpretation that each block-cycle produces a bounded number of point-cycles, but the possible point-cycle lengths are constrained by the within-block structure.

6.3 The full symmetric group $\Gamma = S_k$

Let $\Gamma = S_k$. Then $\mathbb{P}_\Gamma(\lambda) = |\text{class}(\lambda)|/k! = 1/z_\lambda$, where

$$z_\lambda = \prod_{j\geq 1} j^{a_j(\lambda)}\, a_j(\lambda)!.$$

In this setting essentially all cycle types occur, and the dependence pattern among the A_i's is correspondingly richer: a single block-cycle of length ℓ marked by λ contributes to each coordinate $A_{\ell j}$ for which $a_j(\lambda) > 0$.

While closed forms for the full joint law are notationally heavy, many basic quantities admit concise expressions. First,

$$\mathbb{E}A_i = \sum_{\ell | i} \frac{1}{\ell}\, \mathbb{E}\big[a_{i/\ell}(Y)\big], \qquad Y \sim \text{Unif}(S_k). \tag{36}$$

The well-known identity $\mathbb{E}[a_j(Y)] = 1/j$ holds *exactly* for every $1 \leq j \leq k$. Therefore the mean is a truncated divisor sum:

$$\mathbb{E}A_i = \sum_{\substack{\ell | i \\ i/\ell \leq k}} \frac{1}{\ell} \cdot \frac{\ell}{i} = \frac{1}{i}\, \#\{d \mid i :\ d \leq k\}. \tag{37}$$

Thus, for $i \leq k$ we recover $\mathbb{E}A_i = \tau(i)/i$ where τ is the divisor function, whereas for $i > k$ the mean is reduced because within-block cycle lengths cannot exceed k.

Second, covariance can be expressed directly in terms of mixed moments of cycle counts in a uniform permutation on k letters:

$$\text{Cov}(A_i, A_{i'}) = \sum_{\ell | \gcd(i,i')} \frac{1}{\ell}\, \text{Cov}\big(a_{i/\ell}(Y), a_{i'/\ell}(Y)\big) + \sum_{\ell | \gcd(i,i')} \frac{1}{\ell}\, \mathbb{E}\big[a_{i/\ell}(Y)\, a_{i'/\ell}(Y)\big], \tag{38}$$

where the second term is the contribution of the shared Poisson randomness at spacing ℓ. Using factorial-moment identities for random permutations, one has for integers $r, s \geq 1$ with $r + s \leq k$ the exact formula

$$\mathbb{E}\big[a_r(Y) a_s(Y)\big] = \frac{1}{rs} \qquad (r \neq s),$$

and similarly $\mathbb{E}[a_r(Y)(a_r(Y) - 1)] = 1/r^2$ when $2r \leq k$. Plugging these identities into (38) yields explicit covariances for coordinates $A_i, A_{i'}$ whenever i/ℓ and i'/ℓ are not too large relative to k. In particular, for $\Gamma = S_k$ one typically sees positive correlations along multiplicative relations (common ℓ) but also cancellations due to the constraint $\sum_j j a_j(Y) = k$ inside the mark distribution.

6.4 Numerical illustration of the B/n error scaling

We briefly describe a simulation protocol confirming the B/n scaling suggested by the quantitative bounds. Sampling $\sigma \in G_{k,n} = \Gamma^n \rtimes S_n$ is straightforward: sample $\pi \sim \mathrm{Unif}(S_n)$ and i.i.d. marks $\gamma_1, \dots, \gamma_n \sim \mathrm{Unif}(\Gamma)$, then form the induced permutation of $[kn]$ under the standard block action. The point-level cycle counts $(a_1(\sigma), \dots, a_B(\sigma))$ can then be computed by standard cycle-tracing on $[kn]$.

On the limiting side, we sample the independent Poisson family $\{Z_{\ell,\lambda}\}$ for $\ell \le B$ and $\lambda \vdash k$ (or only those λ present in Γ), and form $(A_1, \dots, A_B)$ by the deterministic compounding map. Since k is fixed, the number of partitions $\lambda \vdash k$ is small, and this construction is fast even for moderately large B.

To compare laws in a way that remains feasible in moderate dimension, we focus on Lipschitz test functionals F (as in the previous subsection), such as

$$F_1(x) = \sum_{b=1}^{B} x_b, \qquad F_2(x) = \sum_{b=1}^{B} b\, x_b, \qquad F_3(x) = \max_{1 \le b \le B} x_b, \qquad F_4(x) = \sum_{b=1}^{B} |x_b - \mathbb{E}A_b|.$$

For each n and B, we estimate $\mathbb{E}F(a_1, \dots, a_B)$ and $\mathbb{E}F(A_1, \dots, A_B)$ by Monte Carlo, and plot the absolute difference against B/n. For representative choices (e.g. $k = 3$, $\Gamma = S_3$ and $\Gamma = C_3$), one observes a linear regime in B/n once n is moderately large, with the cyclic case exhibiting noticeably smaller slope, consistent with the fact that (A_i) is independent for $\Gamma = C_k$ whereas for $\Gamma = S_3$ the shared variables $Z_{\ell,(2,1)}$ induce correlations (cf. (35)). In the same experiments, keeping B fixed and increasing n yields an error decaying like $1/n$, while choosing mesoscopic cutoffs such as $B = \lfloor n^\alpha \rfloor$ with $\alpha \in (0,1)$ yields decay $n^{\alpha-1}$, matching the qualitative prediction that the approximation is accurate precisely when $B = o(n)$.

These computations do not provide sharp constants, but they support two robust conclusions: (i) the compounding description captures the correct joint structure (including dependence) already at moderate n, and (ii) the dominant source of error is the truncation at time n in the underlying indicator construction, which naturally scales as B/n when one tracks starts of length at most B.

7 Discussion and open problems

The results above identify a natural *compound Poisson* description for truncated point-level cycle counts in the imprimitive action of $G_{k,n} = \Gamma^n \rtimes S_n$, with an explicit error scale in total variation of order B/n for mesoscopic cutoffs $B = o(n)$. We record several directions in which one can sharpen, extend, or reorganize the theory.

7.1 Optimal constants and sharp asymptotics in total variation

Our quantitative bound

$$\big\|\mathcal{L}(a_1(\sigma),\ldots,a_B(\sigma)) - \mathcal{L}(A_1,\ldots,A_B)\big\|_{\mathrm{TV}} \;\le\; C(k,\Gamma)\Big(\frac{B}{n}+\frac{B^2}{n^2}\Big)$$

is robust but not expected to be sharp in its constant, and the presence of two terms suggests that different error mechanisms are being aggregated. Even in the classical case $\Gamma=\{e\}$ (uniform permutations on S_n), the optimal leading constant in a B/n bound depends delicately on how one measures truncation error in the Feller coupling, and one can often track the dominant contribution to the TV distance to within an absolute factor.

In the wreath product setting, the natural conjecture is that the leading constant should depend on Γ primarily through low moments of the *mark profile*

$$M(\lambda) := \sum_{j=1}^{k} a_j(\lambda), \qquad Y\sim \mathrm{Unif}(\Gamma),$$

since $M(Y)$ is exactly the number of point-cycles created by a single block-cycle (counting multiplicities across within-block cycle lengths). At a heuristic level, each block-cycle of length ℓ that is "mis-coupled" at the block level can create at most $M(Y)$ discrepancies across the point-level coordinates, so one expects constants of the form $C_1\,\mathbb{E}M(Y)$ (possibly together with $\mathbb{E}M(Y)^2$ when one controls TV via second moments).

A concrete open problem is therefore to isolate the optimal constant in an inequality of the form

$$\big\|\mathcal{L}(a_1(\sigma),\ldots,a_B(\sigma)) - \mathcal{L}(A_1,\ldots,A_B)\big\|_{\mathrm{TV}} \;\le\; C_\star(k,\Gamma)\,\frac{B}{n} \quad \text{for all } 1\le B\le n,$$

and to decide whether $C_\star(k,\Gamma)$ admits a tractable expression in terms of the cycle index of Γ. One plausible approach is to combine (i) a refined block-level approximation (e.g. an optimal TV coupling for the vector of block cycle counts up to B) with (ii) an optimal Lipschitz estimate for the deterministic compounding map sending block-cycle counts and marks to point-cycle counts. The latter can be analyzed exactly because it is linear in the Poisson family $\{Z_{\ell,\lambda}\}$, but translating this into sharp TV constants remains nontrivial.

7.2 Lower bounds and sharpness of the B/n regime

The condition $B=o(n)$ is sufficient for convergence in total variation of the truncated vector, and in all comparable models it is also essentially necessary. A systematic lower bound in our setting would be valuable, both as a sharpness statement and as a guide to optimal constants.

The main obstruction is already present at the block level: in the indicator (Feller) representation for S_n, the truncation at time n perturbs the joint law of short spacings by an amount that is typically of order B/n when one tests against events depending on the presence of a spacing start near the terminal segment. In our marked-spacing construction, the additional marks do not remove this obstruction; if anything, they can amplify it when $M(Y)$ is large.

One concrete route to a lower bound is to compare the finite-n marked-spacing counts $C_b^{(n)}$ with the infinite counts $C_b^{(\infty)}$ by testing the event that there exists a spacing of length $\ell \le B$ starting in the last B indices of the indicator sequence. Such an event has probability of order B/n at the block level, and whenever it occurs it changes (after compounding by marks) at least one coordinate among $(a_1, \dots, a_B)$ with positive probability uniformly bounded away from 0 (provided Γ is not supported on the identity). Formalizing this strategy should yield a general lower bound

$$\left\| \mathcal{L}(a_1(\sigma), \dots, a_B(\sigma)) - \mathcal{L}(A_1, \dots, A_B) \right\|_{\mathrm{TV}} \;\ge\; c(k, \Gamma)\, \frac{B}{n}$$

at least in regimes where $B \to \infty$ and $B = o(n)$. Determining the largest class of (k, Γ) for which such a bound holds (and identifying the exact $c(k, \Gamma)$) is open.

A related problem is to understand what happens when B is a linear fraction of n. In that regime, the approximation cannot hold in TV for the entire vector, but one may still hope for (i) convergence of finite-dimensional marginals, (ii) Wasserstein-type bounds for suitable metrics on $\mathbb{N}^B$, or (iii) limit theorems for aggregated statistics (see below).

7.3 Replacing S_n by other block permutation groups $H \le S_n$

The analysis above is tailored to $G_{k,n} = \Gamma^n \rtimes S_n$, where the block permutation is uniform on the full symmetric group and hence admits the Feller indicator representation with independent Bernoulli indicators ζ_i having $\mathbb{P}(\zeta_i = 1) = 1/i$. A natural extension is to consider

$$G_{k,n}(H) := \Gamma^n \rtimes H$$

for a sequence of subgroups $H \le S_n$ (or more generally, for a sequence of conjugation-invariant measures on S_n), still acting via the standard block action.

At the level of small block cycles, what one needs from H is an analogue of the classical Poisson approximation: namely, that the vector of block-cycle counts up to length B is close to a vector of independent Poisson variables with means $\{\theta_\ell/\ell\}_{\ell \le B}$ for some parameters θ_ℓ depending on H (and possibly on n in a controlled way). If such a block-level approximation holds in total variation with error $\varepsilon_{n,B}$, then the same compounding map used to define (A_i) produces a point-level approximation with error at most a constant multiple of $\varepsilon_{n,B}$, with the constant depending only on (k, Γ) through a Lipschitz bound for compounding. In other words, the point-level question largely reduces to establishing a suitable "small cycles are asymptotically Poisson" theorem for H.

Two cases appear particularly accessible:

- $H = A_n$ (uniform even permutations), where one expects the difference from S_n to be negligible at the scale $B = o(n)$ for fixed k.
- Ewens-type measures on S_n, where the block-cycle counts are asymptotically independent Poisson with mean θ/ℓ; here the limiting object should be obtained from our definition of A_i by replacing the mean of $Z_{\ell,\lambda}$ from $\mathbb{P}_\Gamma(\lambda)/\ell$ to $\theta\,\mathbb{P}_\Gamma(\lambda)/\ell$.

For more rigid subgroups H (e.g. primitive groups of small order), one does not expect Poisson block-cycle counts, and it becomes a separate problem to identify the correct block-level limit and then compound it.

7.4 Toward functional limit theorems and random measures

The present results provide a TV approximation for the finite vector $(a_1, \ldots, a_B)$ when $B = o(n)$. A natural next step is to seek *process-level* convergence as $B \to \infty$ together with n, phrased either in a product topology on $\mathbb{N}^{\mathbb{N}}$ or in terms of random measures. One convenient encoding is the point measure of (small) cycle lengths

$$\Xi_n := \sum_{i \geq 1} a_i(\sigma)\,\delta_i, \qquad \Xi := \sum_{i \geq 1} A_i\,\delta_i,$$

and one may ask for convergence of Ξ_n to Ξ when tested against functions supported on $\{1, \ldots, B\}$ with $B = o(n)$, with explicit bounds in metrics stronger than total variation on finite truncations (e.g. bounded-Lipschitz metrics on measures). Because (A_i) is itself built from an underlying independent Poisson family indexed by (ℓ, λ), it is plausible that Ξ can be realized as a compound Poisson random measure with intensity

$$\sum_{\ell \geq 1} \sum_{\lambda \vdash k} \frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\,\delta_\ell \otimes \delta_\lambda,$$

pushed forward under the deterministic map $(\ell, \lambda) \mapsto$ the multiset of point-cycle lengths $\{\ell j : a_j(\lambda) > 0\}$ with multiplicities $a_j(\lambda)$. Making such a representation precise could simplify the derivation of joint cumulants and facilitate stronger limit theorems for linear statistics.

Another direction is to study aggregated functionals beyond the Lipschitz class treated above, for instance

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

or centered versions thereof, and to ask for Gaussian fluctuations (CLTs) as $B \to \infty$ with $B = o(n)$. Since (A_i) has explicit cumulants determined by $\{Z_{\ell,\lambda}\}$, one can compute $\mathrm{Var}(S_B)$ and higher cumulants and then ask whether $(S_B - \mathbb{E}S_B)/\sqrt{\mathrm{Var}(S_B)}$ is asymptotically normal uniformly over n in the regime $B = o(n)$. Establishing such functional limit theorems would require controlling dependence created by shared Poisson components across coordinates, a feature absent in the simplest examples (e.g. $\Gamma = C_k$) but present for general Γ.

7.5 Macroscopic cycles and marked Poisson–Dirichlet structures

Finally, the current work addresses point-cycle lengths bounded by $B = o(n)$, whereas the block permutation on n blocks exhibits macroscopic cycles of order n with Poisson–Dirichlet scaling. Because a block-cycle of length ℓ marked by λ produces point-cycles of lengths ℓj (for j in the support of λ), one expects the macroscopic cycle structure of σ (as a permutation of $[kn]$) to be describable as a *marking/fragmentation* of the Poisson–Dirichlet limit for S_n: each macroscopic block-cycle carries an independent Γ-mark which splits its mass into finitely many point-cycles with proportions j/k repeated $a_j(\lambda)$ times. Turning this heuristic into a joint limit theorem, simultaneously capturing small cycles (compound Poisson) and macroscopic cycles (marked Poisson–Dirichlet), remains open and would require a coupling that is uniform across all scales of ℓ.

7.6 Growing block size k

All statements above keep k fixed. Allowing $k = k(n) \to \infty$ introduces two new effects: (i) the space of marks $\lambda \vdash k$ becomes large, and (ii) the internal cycle structure begins to resemble that of a large random permutation if Γ is, say, S_k or a large subgroup. One expects new regimes in which k interacts with B (for instance, the set of attainable point-cycle lengths up to B depends on both). Identifying conditions under which the same approximation holds, possibly with constants uniform in k in certain ranges (e.g. $k = o(n)$), is a natural open problem, as is the identification of the correct limiting object when both k and n diverge.

We view these questions as largely orthogonal: one may first sharpen constants at fixed (k, Γ), then generalize the block permutation law, and finally address multiscale limits. Each step requires only incremental new input beyond the compound-Poisson description already established, but each also appears to contain genuinely new phenomena.