A Lévy–Khintchine Description of Wreath-Product Cycle Counts

Liz Lemma Future Detective

January 17, 2026

Abstract

We study the joint distribution of small cycle counts in random permutations drawn uniformly from the imprimitive wreath-product subgroup $G_n = \Gamma^n \rtimes S_n \le S_{kn}$, with fixed block size k and fixed internal group $\Gamma \le S_k$. Building on the Poissonization formula of Diaconis–Tung for the cycle index of wreath products, we recast the limiting cycle-count vector $A = (A_i)_{i\ge 1}$ as an explicitly infinitely divisible law on $\mathbb{N}^{\mathbb{N}}$ with a discrete Lévy measure supported on "cluster increments" indexed by block-cycle lengths ℓ and internal cycle types $\lambda \vdash k$. This yields a clean Lévy–Khintchine formula for all finite-dimensional Laplace transforms, explicit joint cumulants, and an exact criterion for when different cycle counts are independent (a divisibility-and-type condition). We further prove an identifiability theorem: within fixed k, the Lévy measure (equivalently the full limit law) determines and is determined by the cycle-type distribution of a uniform element of Γ. Several examples (e.g. $\Gamma = C_k$ vs $\Gamma = S_k$) illustrate how dependence arises and how to compute correlations. The paper packages wreath-product Poisson limits into a modern, reusable object suited for downstream applications (Stein bounds, large deviations, and algorithmic analysis).

Table of Contents

4. 4. Structural corollaries: joint cumulants and covariance formulas; dependence graph characterization; criteria for mutual independence blocks (e.g. odd indices, prime-power patterns) in terms of the support of ν_Γ.
5. 5. Identifiability and completeness: show ν_Γ (or the law of A) determines $\mathbb{P}_\Gamma(\lambda)$; discuss what is and is not identified about Γ as an abstract subgroup.
6. 6. Worked examples and computations: $\Gamma = C_k$ (independent coordinates), $\Gamma = S_k$ (explicit dependence), and one small-k table; optional computer verification for $k \leq 6$ (flagged as computational).
7. 7. Extensions and remarks: replacing S_n by other H_n with weights α_ℓ (brief), and how the Lévy measure interfaces with Stein's method and large deviations; open problems.

1 Introduction

We study the asymptotic cycle structure of the natural imprimitive action of the wreath product $\Gamma^n \rtimes S_n$ on kn points, where the block size k and the subgroup $\Gamma \leq S_k$ are fixed and $n \to \infty$. For uniform permutations in the full symmetric group, the classical Poisson limit for counts of short cycles is most naturally expressed as an infinite divisibility statement: the vector of cycle counts has a limiting law which is a product of Poisson distributions, hence compound Poisson with purely atomic Lévy measure supported on the standard basis vectors. In the wreath-product setting, one observes a comparable phenomenon at the level of blocks, but the induced cycle structure on $[kn]$ is no longer coordinatewise independent. Our purpose is to record a Lévy–Khintchine viewpoint that both packages the known limit formula into a single conceptual object (a discrete Lévy measure on finitely supported sequences) and yields immediate structural consequences, such as a sharp independence criterion and a completeness statement for the information carried by the limit.

The basic mechanism is as follows. A permutation in $\Gamma^n \rtimes S_n$ consists of a block permutation (an element of S_n) together with an internal element of Γ attached to each block. A cycle of the block permutation of length ℓ acts on ℓ blocks, and the internal labels along that block-cycle multiply to a random element of Γ. The cycle type of this product determines how points inside those ℓ blocks are permuted: a j-cycle in the internal product gives rise to cycles of length $j\ell$ on $[kn]$. Thus a single block-cycle contributes a deterministic increment to the cycle-count vector, depending only on ℓ and on the internal cycle type. Since the block permutation of S_n has, in the limit, asymptotically independent Poisson counts of cycles of each length ℓ, we are led to represent the limit of the cycle-count vector as a superposition of independent Poisson numbers of such increments. This is precisely the compound Poisson paradigm: an infinitely divisible law specified by its Lévy measure, with atoms indexed by the possible block-cycle lengths and internal cycle types.

A closely related limit description already appears in work of Diaconis–Tung, who compute a Poissonized cycle index and thereby obtain an explicit limiting joint probability generating function for the cycle counts in the wreath-product action. Their formula is exact and computationally effective, but it is presented in a way that does not emphasize the canonical decomposition into independent Poisson components. The present note reframes that limiting generating function as the log-Laplace functional of a discrete infinitely divisible law. Concretely, after taking logarithms, the limit exponent decomposes as a sum over the atoms corresponding to block-cycle length ℓ and internal cycle type λ. This decomposition is not merely cosmetic: it is the statement that the limit law admits a representation as a sum of independent Poisson random variables times deterministic increment vectors, and that all joint cumulant information is obtained by summing the corresponding contributions over the Lévy support. In this way, the wreath-product cycle index is read as a Lévy–Khintchine formula on the natural state space of cycle-count vectors.

One advantage of this perspective is that it isolates the dependence structure of the limiting cycle counts. In the symmetric-group case, each coordinate of the limit is driven by its own independent Poisson component, so all coordinates are independent. For wreath products, a single increment typically affects several coordinates at once (for instance, an internal type containing both 1-cycles and 2-cycles produces contributions simultaneously to cycle lengths ℓ and 2ℓ), and therefore dependence is inevitable in general. The Lévy decomposition makes the criterion for independence transparent: two coordinates are independent if and only if they are measurable with respect to disjoint families of Poisson components, equivalently if there is no atom of the Lévy measure whose increment has both coordinates positive. This yields a simple "dependence graph" on indices $i \geq 1$, computable directly from the cycle types that occur with positive probability in Γ. The criterion recovers, in particular, the known independence phenomenon for certain choices of Γ (notably cyclic Γ), and it also explains the typical patterns of dependence for large subgroups such as $\Gamma = S_k$, where many internal cycle types are present and shared divisors $\ell \mid \gcd(i, i')$ create unavoidable links between coordinates A_i and $A_{i'}$.

A second advantage is that the Lévy measure provides a minimal and complete parameterization of the limit law. By construction, the limit depends on Γ only through the distribution of the cycle type of a uniform element of Γ. This reduction is implicit in Diaconis–Tung's formula, since the internal group elements enter only via their cycle index. We go one step further and prove an identifiability statement at fixed block size k: the limit law (equivalently, the Lévy measure) determines the cycle-type distribution on Γ uniquely. The argument uses a "gcd slicing" of the Lévy support to separate contributions coming from block-cycles of different lengths ℓ, and then observes that the increments arising from $\ell = 1$ encode the internal cycle counts without ambiguity. In particular, while the limit law cannot distinguish non-isomorphic subgroups of S_k that have the same cycle-type distribution, it does distinguish different cycle-type distributions. This clarifies exactly what information about Γ is visible in the asymptotic cycle structure of the wreath-product action.

The Lévy viewpoint also streamlines the derivation of quantitative formulas. Once the limit is recognized as compound Poisson with known atoms and weights, joint factorial moments and cumulants follow from standard identities for Poisson sums. In particular, covariances and higher mixed cumulants reduce to finite sums over the relevant atoms whose increments hit the specified coordinates. This produces explicit, computable expressions and makes finiteness transparent at the level of each finite-dimensional marginal: for a fixed set of cycle lengths, only finitely many block-cycle lengths ℓ can contribute. Thus the measure-theoretic issues associated with working on $\mathbb{N}^{\mathbb{N}}$ are handled by a projective-limit philosophy: all statements are interpreted through finite-dimensional distributions, where the Lévy measure projects to a finite measure and the usual Lévy–Khintchine formalism applies without modification.

Methodologically, our proofs do not require new asymptotic enumeration beyond what is already encoded in the wreath-product cycle index. The core step is to take the limiting joint generating function furnished by the Poissonized cycle index and to rewrite its logarithm as a sum of terms of the form $\mu(e^{\langle t,v\rangle} - 1)$, which is the signature of a Poisson component with mean μ and jump (increment) v. From this rewriting we read off the atomic Lévy measure, obtain the independent-Poisson representation, and then deduce the independence criterion and cumulant formulas by general principles. The identifiability theorem then becomes a structural statement about the support of the Lévy measure: the increments are indexed by (ℓ, λ) in a way that is injective once ℓ is fixed appropriately, so the atomic weights recover the internal cycle-type probabilities.

The organization of the paper follows this logic. After setting notation and recalling the necessary background on cycle counts, partitions, and the wreath-product action, we introduce the space of finitely supported increment vectors and formalize the notion of discrete infinite divisibility appropriate for cycle-count vectors. We then establish the Lévy–Khintchine representation of the limit law and the independent-Poisson decomposition. Subsequent sections extract consequences: the dependence graph criterion and explicit cumulant/covariance formulas, followed by the identifiability result. Throughout, our aim is to treat the Lévy measure as the primary object: it is the compact encoding of the limit, it separates the roles of block structure and internal structure, and it provides a calculus for dependence and inference that is not apparent from the raw generating function alone.

2 Notation and background

2.1 Cycle counts and cycle type

For a permutation $\sigma \in S_N$ we write its cycle-count data as

$$\sigma \sim \prod_{i=1}^{N} i^{a_i(\sigma)},$$

meaning that σ has exactly $a_i(\sigma)$ cycles of length i. Thus $a_i(\sigma) \in \mathbb{N}$ and $\sum_{i\geq 1} i\, a_i(\sigma) = N$. When $N = kn$ varies with n we extend $(a_i(\sigma))_{i\geq 1}$ by zeros for $i > N$, and regard the cycle-count vector as an element of $\mathbb{N}^{\mathbb{N}}$.

We encode conjugacy classes in S_k by partitions $\lambda \vdash k$. Concretely, we use the cycle-count notation

$$\lambda = 1^{a_1(\lambda)} 2^{a_2(\lambda)} \cdots k^{a_k(\lambda)}, \qquad \text{where } \sum_{j=1}^{k} j\, a_j(\lambda) = k,$$

so that $a_j(\lambda)$ records the number of j-cycles in a permutation of type λ. If $\gamma \in S_k$, we write $\mathrm{ctype}(\gamma) = \lambda$ to denote its cycle type, and we set

$$\mathbb{P}_\Gamma(\lambda) := \mathbb{P}\big(\mathrm{ctype}(\gamma) = \lambda\big), \qquad \gamma \sim \mathrm{Unif}(\Gamma).$$

This distribution on partitions is the only group-specific input that will remain visible in the asymptotic laws.

2.2 Wreath products and the imprimitive block action

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

$$G_n := \Gamma^n \rtimes S_n,$$

where S_n acts on Γ^n by permuting coordinates:

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

An element of G_n will be written as (γ, π) with $\gamma = (\gamma_1, \ldots, \gamma_n) \in \Gamma^n$ and $\pi \in S_n$, and multiplication is given by

$$(\gamma, \pi)(\gamma', \pi') = (\gamma\,(\pi \cdot \gamma'),\, \pi\pi').$$

We embed G_n into S_{kn} via the standard imprimitive action on kn points. We identify

$$[kn] \cong [n] \times [k],$$

where $[n]$ indexes blocks and $[k]$ indexes positions within each block. The action of $(\gamma, \pi) \in G_n$ on $(b, x) \in [n] \times [k]$ is

$$(\gamma, \pi) \cdot (b, x) := \big(\pi(b),\, \gamma_{\pi(b)}(x)\big). \tag{1}$$

(Any equivalent convention, differing by an inversion in the index of γ, leads to the same cycle structure on $[kn]$; we fix (1) for definiteness.)

A basic structural feature of (1) is that the cycle structure on $[kn]$ is determined by the interaction between (i) the cycle decomposition of the block permutation $\pi \in S_n$ and (ii) the cycle types of certain products of the internal labels $\gamma_b \in \Gamma$ along the cycles of π. Indeed, let $C = (b_1\, b_2\, \ldots\, b_\ell)$ be a cycle of π of length ℓ. Restricting (1) to the union of blocks $\{b_1, \ldots, b_\ell\} \times [k]$, we see that one application of (γ, π) moves from block b_r to block b_{r+1} and applies the internal permutation attached to the target block. After ℓ steps we return to the initial block and have applied the product

$$g_C := \gamma_{b_1}\gamma_{b_\ell}\cdots\gamma_{b_2} \in \Gamma$$

(up to cyclic reordering, depending on the chosen convention). Consequently, the induced permutation on $\{b_1, \ldots, b_\ell\} \times [k]$ decomposes into cycles whose lengths are multiples of ℓ, with multiplicities determined by the cycle structure of g_C on $[k]$. In particular, if g_C has a_j cycles of length j on $[k]$, then the induced permutation on $[kn]$ has a_j cycles of length $j\ell$ supported on these ℓ blocks. This "ℓ-fold stretching" mechanism is the combinatorial origin of the compound-Poisson limit and will be formalized in the next section.

For $\sigma_n \sim \mathrm{Unif}(G_n)$, the block permutation π is uniform in S_n and independent of the internal labels $\gamma_1, \ldots, \gamma_n$, which are i.i.d. uniform in Γ. Moreover, for a fixed block-cycle C of length ℓ, the product g_C is uniform in Γ (a product of independent uniform elements in a finite group is uniform), and products corresponding to disjoint block-cycles are independent because they involve disjoint sets of labels. Thus the random contributions coming from distinct block-cycles separate cleanly.

2.3 State spaces and modes of convergence

We view cycle-count vectors as random elements of the countable product space $\mathbb{N}^{\mathbb{N}}$, equipped with its product σ-algebra. Since $\mathbb{N}^{\mathbb{N}}$ is not locally compact and we do not seek any topology beyond coordinatewise information, all distributional convergence statements are interpreted through finite-dimensional marginals. Namely, if $X^{(n)} = (X_i^{(n)})_{i\geq 1}$ and $X = (X_i)_{i\geq 1}$ are $\mathbb{N}^{\mathbb{N}}$-valued, we write

$$X^{(n)} \xrightarrow{\text{f.d.d.}} X$$

to mean that for every $m \geq 1$,

$$(X_1^{(n)}, \ldots, X_m^{(n)}) \xrightarrow{\text{law}} (X_1, \ldots, X_m) \quad \text{in } \mathbb{N}^m.$$

This is the natural framework for cycle counts: for each fixed m, only finitely many combinatorial structures can influence the vector of short-cycle counts up to length m, and explicit generating functions are available at that level.

Alongside $\mathbb{N}^{\mathbb{N}}$ we use the space of finitely supported sequences

$$\mathbb{N}^{(\mathbb{N})} := \Big\{ v = (v_i)_{i\geq 1} \in \mathbb{N}^{\mathbb{N}} : v_i = 0 \text{ for all but finitely many } i \Big\}.$$

This is the natural "jump space" for compound Poisson constructions: a single block-cycle affects only finitely many cycle lengths, hence contributes an increment in $\mathbb{N}^{(\mathbb{N})}$.

2.4 Discrete infinite divisibility and compound Poisson vectors

We recall the notion of infinite divisibility in the present discrete, infinite-dimensional setting. A random element $A = (A_i)_{i\geq 1}$ in $\mathbb{N}^{\mathbb{N}}$ is called infinitely divisible (in the sense relevant here) if for every $r \geq 1$ and every $m \geq 1$, the m-dimensional marginal $(A_1, \dots, A_m)$ is infinitely divisible as an $\mathbb{N}^m$-valued random vector. Equivalently, for each r and m there exist i.i.d. $\mathbb{N}^m$-valued random vectors $Y^{(r,1)}, \dots, Y^{(r,r)}$ such that

$$(A_1, \dots, A_m) \overset{d}{=} Y^{(r,1)} + \cdots + Y^{(r,r)}.$$

In our application, the limiting law will in fact be compound Poisson in each finite dimension, and hence infinitely divisible.

Let us therefore fix m and recall the standard compound Poisson form on $\mathbb{N}^m$. A random vector $X \in \mathbb{N}^m$ is compound Poisson if there exists a finite measure $\nu^{(m)}$ on $\mathbb{N}^m \setminus \{0\}$ such that

$$\log \mathbb{E}\big[e^{\langle t,X\rangle}\big] = \sum_{u\in\mathbb{N}^m\setminus\{0\}} \nu^{(m)}(\{u\})\big(e^{\langle t,u\rangle} - 1\big), \qquad t \in \mathbb{R}^m. \tag{2}$$

Equivalently, if $(Z_u)_{u\in\mathbb{N}^m\setminus\{0\}}$ are independent Poisson random variables with means $\nu^{(m)}(\{u\})$, then

$$X \overset{d}{=} \sum_{u\in\mathbb{N}^m\setminus\{0\}} Z_u\, u,$$

where the sum is almost surely finite because $\nu^{(m)}$ is finite.

For an $\mathbb{N}^{\mathbb{N}}$-valued vector A we say that A is (discrete) compound Poisson if for each m the marginal $(A_1, \dots, A_m)$ admits a representation of the form (2) with a finite measure $\nu^{(m)}$, and these measures are consistent under projection. One convenient way to encode consistency is via a σ-finite measure ν on $\mathbb{N}^{(\mathbb{N})} \setminus \{0\}$ such that for each m the pushforward $\nu^{(m)}$ obtained by projecting $v \mapsto (v_1, \dots, v_m)$ is finite. In that case (2) becomes, for $t \in \mathbb{R}^m$,

$$\log \mathbb{E}\Big[\exp\Big(\sum_{i=1}^m t_i A_i\Big)\Big] = \sum_{v\in\mathbb{N}^{(\mathbb{N})}\setminus\{0\}} \nu(\{v\})\Big(\exp\Big(\sum_{i=1}^m t_i v_i\Big) - 1\Big),$$

with the understanding that $v_i = 0$ for all $i > m$ inside the exponent. We will refer to such a ν as a (discrete) Lévy measure for A. The next section identifies the relevant atoms v arising from single block-cycles and computes their weights directly from the cycle-type distribution on Γ and the classical Poisson asymptotics for cycles in a uniform element of S_n.

3 The marked-Poisson construction and the Lévy measure

Fix $m \geq 1$ and write $t = (t_1, \dots, t_m) \in \mathbb{R}^m$, with the convention that $t_r = 0$ for $r > m$. For $\sigma_n = (\gamma, \pi) \in G_n$ we regard the short-cycle data $(a_1(\sigma_n), \dots, a_m(\sigma_n))$ as the statistic of interest and isolate the contribution coming from each cycle of the block permutation $\pi \in S_n$. The guiding observation is that a block-cycle of length ℓ can only create cycles in S_{kn} whose lengths lie in $\{\ell, 2\ell, \dots, k\ell\}$, hence it influences our m-dimensional marginal only when $\ell \leq m$.

Let $C = (b_1\, b_2\, \dots\, b_\ell)$ be a cycle of π of length ℓ. As recalled above, the restriction of σ_n to the union of the corresponding ℓ blocks is determined by the product of internal labels along C, which we denote by

$$g_C := \gamma_{b_1}\gamma_{b_\ell}\cdots\gamma_{b_2} \in \Gamma,$$

where the particular cyclic order is immaterial for cycle counts on $[k]$. The key distributional fact is that, for $\gamma_1, \dots, \gamma_n$ i.i.d. uniform on Γ, each g_C is uniform on Γ, and the family $(g_C)_C$ over distinct block-cycles C is independent. Indeed, for a fixed C this is the elementary identity that the product of independent uniform random variables on a finite group is uniform; independence across cycles holds because disjoint block-cycles use disjoint sets of coordinates of γ.

Write $\operatorname{ctype}(g_C) = \lambda \vdash k$. If g_C has $a_j(\lambda)$ cycles of length j on $[k]$, then the induced permutation on the ℓ blocks of C has exactly $a_j(\lambda)$ cycles of length $j\ell$ on $[kn]$. This motivates the increment vectors: for $\ell \in \mathbb{N}$ and $\lambda \vdash k$ we set

$$v(\ell,\lambda) \in \mathbb{N}^{(\mathbb{N})}, \qquad v(\ell,\lambda)_{j\ell} = a_j(\lambda)\ (1 \leq j \leq k), \qquad v(\ell,\lambda)_i = 0\ (i \notin \{\ell, 2\ell, \dots, k\ell\}).$$

Thus each block-cycle C contributes exactly the random increment $v(\ell,\lambda)$ with $\ell = |C|$ and $\lambda = \operatorname{ctype}(g_C)$, and the short-cycle vector is obtained by summing these contributions over all block-cycles. In particular, if we let $Z^{(n)}_{\ell,\lambda}$ denote the number of ℓ-cycles C of π such that $\operatorname{ctype}(g_C) = \lambda$, then for each fixed m we have the identity of $\mathbb{N}^m$-valued random vectors

$$(a_1(\sigma_n), \dots, a_m(\sigma_n)) = \sum_{\ell=1}^{m} \sum_{\lambda \vdash k} Z^{(n)}_{\ell,\lambda}\,(v(\ell,\lambda)_1, \dots, v(\ell,\lambda)_m), \tag{3}$$

since block-cycles with $\ell > m$ cannot contribute to coordinates $\leq m$.

We next identify the joint law of the marked counts $Z^{(n)}_{\ell,\lambda}$ as $n \to \infty$. Let $C_\ell(\pi)$ be the number of ℓ-cycles of π. Conditional on π, the random variables $(\operatorname{ctype}(g_C))_{C:\,|C|=\ell}$ are i.i.d. with law $\mathbb{P}_\Gamma$ on partitions of k, hence

$$(Z^{(n)}_{\ell,\lambda})_{\lambda\vdash k} \bigm| \pi \;\sim\; \text{Multinomial}\big(C_\ell(\pi);\ (\mathbb{P}_\Gamma(\lambda))_{\lambda\vdash k}\big), \qquad 1 \leq \ell \leq m.$$

Unconditionally, we combine this with the classical Poisson asymptotics for cycle counts in a uniform permutation: as $n \to \infty$,

$$(C_1(\pi), \ldots, C_m(\pi)) \xrightarrow{\text{law}} (\mathrm{Poi}(1), \mathrm{Poi}(1/2), \ldots, \mathrm{Poi}(1/m)),$$

with independence across ℓ. By the thinning property of Poisson variables, if $C_\ell \sim \mathrm{Poi}(1/\ell)$ and each of its points is independently marked with type λ with probability $\mathbb{P}_\Gamma(\lambda)$, then the resulting marked counts are independent Poisson with means $\mathbb{P}_\Gamma(\lambda)/\ell$. Consequently, for each fixed m,

$$\big(Z_{\ell,\lambda}^{(n)} : 1 \le \ell \le m,\ \lambda \vdash k\big) \xrightarrow{\text{law}} \big(Z_{\ell,\lambda} : 1 \le \ell \le m,\ \lambda \vdash k\big), \tag{4}$$

where the family $(Z_{\ell,\lambda})_{\ell \ge 1,\ \lambda \vdash k}$ is independent and

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

We now pass from (3) and (4) to the limiting Laplace transform. For fixed m and $t \in \mathbb{R}^m$ we compute, conditioning on π and using the independence of the marks across block-cycles,

$$\begin{aligned}\mathbb{E}\Big[\exp\Big(\sum_{i=1}^m t_i a_i(\sigma_n)\Big) \,\Big|\, \pi\Big] &= \prod_{C \text{ cycle of } \pi} \mathbb{E}\Big[\exp\Big(\sum_{i=1}^m t_i\, v(|C|, \mathrm{ctype}(g_C))_i\Big)\Big] \\ &= \prod_{\ell \ge 1}\Big(\sum_{\lambda \vdash k} \mathbb{P}_\Gamma(\lambda)\, \exp\Big(\sum_{j=1}^k t_{j\ell}\, a_j(\lambda)\Big)\Big)^{C_\ell(\pi)}.\end{aligned}$$

Only $\ell \le m$ contribute inside the exponent (since $t_{j\ell} = 0$ when $j\ell > m$), so the product may be truncated at $\ell = m$. Taking expectation over π and using the standard convergence of the joint pgf of $(C_1(\pi), \ldots, C_m(\pi))$, we obtain

$$\lim_{n\to\infty} \mathbb{E}\Big[\exp\Big(\sum_{i=1}^m t_i a_i(\sigma_n)\Big)\Big] = \exp\left(\sum_{\ell=1}^m \frac{1}{\ell}\Big(f_\ell(t) - 1\Big)\right),$$

where

$$f_\ell(t) := \sum_{\lambda \vdash k} \mathbb{P}_\Gamma(\lambda)\, \exp\Big(\sum_{j=1}^k t_{j\ell}\, a_j(\lambda)\Big).$$

Expanding the exponent gives the Lévy–Khintchine form

$$\log \mathbb{E}\Big[\exp\Big(\sum_{i=1}^m t_i A_i\Big)\Big] = \sum_{\ell \ge 1} \sum_{\lambda \vdash k} \frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\left(\exp\Big(\sum_{j=1}^k t_{j\ell}\, a_j(\lambda)\Big) - 1\right), \tag{5}$$

where A denotes the limiting cycle-count vector and we again interpret $t_r = 0$ for $r > m$. Since for fixed m the sum over ℓ is effectively restricted to $\ell \le m$, the right-hand side is finite and defines a proper Laplace transform on $\mathbb{R}^m$.

Formula (5) identifies the relevant discrete Lévy measure on the jump space $\mathbb{N}^{(\mathbb{N})} \setminus \{0\}$. Namely, we define ν_Γ by prescribing its atomic masses on the increments $v(\ell, \lambda)$:

$$\nu_\Gamma\big(\{v(\ell,\lambda)\}\big) := \frac{\mathbb{P}_\Gamma(\lambda)}{\ell}, \qquad \ell \geq 1,\ \lambda \vdash k,$$

and extending by countable additivity. For each m, the pushforward of ν_Γ under projection to the first m coordinates is a finite measure because only $\ell \leq m$ yield nonzero projected increments, and there are finitely many $\lambda \vdash k$. With this notation, (5) is exactly the compound-Poisson Laplace functional associated to ν_Γ.

Finally, the explicit Poisson-sum representation follows by the standard construction of a discrete compound Poisson vector. Let $(Z_{\ell,\lambda})_{\ell \geq 1,\ \lambda \vdash k}$ be independent with $Z_{\ell,\lambda} \sim \text{Poisson}(\mathbb{P}_\Gamma(\lambda)/\ell)$, and set

$$A := \sum_{\ell \geq 1} \sum_{\lambda \vdash k} Z_{\ell,\lambda}\, v(\ell, \lambda) \qquad \text{in } \mathbb{N}^{\mathbb{N}}.$$

This sum is well-defined coordinatewise: for a fixed i, only divisors $\ell \mid i$ can contribute to A_i, and for each such ℓ there are only finitely many $\lambda \vdash k$. Moreover, the Laplace transform of this A is given by the right-hand side of (5), since the log-mgf of a sum of independent Poisson contributions is the sum of their log-mgfs. Therefore the vector constructed above has the same finite-dimensional distributions as the limit of $(a_i(\sigma_n))_{i \geq 1}$, and we may take it as an explicit realization of the limiting law.

3.1 Structural corollaries: cumulants, covariances, and the dependence graph

We now record consequences of the representation

$$A = \sum_{\ell \geq 1} \sum_{\lambda \vdash k} Z_{\ell,\lambda}\, v(\ell, \lambda), \qquad Z_{\ell,\lambda} \sim \text{Poisson}\left(\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\right) \ \text{ independent},$$

which make explicit the dependence structure and provide closed formulas for joint cumulants. All statements below are understood at the level of finite-dimensional marginals: for any fixed collection of indices, only finitely many pairs (ℓ, λ) contribute.

Joint cumulants and factorial cumulants. Let $I \subset \mathbb{N}$ be finite and let $(r_i)_{i \in I}$ be nonnegative integers, not all zero. For an $\mathbb{N}$-valued random variable X we write $(X)_r := X(X-1)\cdots(X-r+1)$ for the falling factorial. Since each coordinate A_i is a sum of independent Poisson contributions with deterministic sizes, it is natural to work with factorial moments and factorial cumulants. Denote by

$$\kappa^{\text{fac}}\big((A_i)_{i \in I};\ (r_i)_{i \in I}\big)$$

the joint factorial cumulant of order $(r_i)_{i\in I}$, i.e. the mixed coefficient obtained by expanding the logarithm of the factorial moment generating function

$$\log \mathbb{E}\Big[\prod_{i\in I}(1+s_i)^{A_i}\Big] \quad \text{around } s_i = 0.$$

Equivalently (and more concretely for our purposes), factorial cumulants are the multilinear forms characterized by additivity over independent sums and by the rule that if $Y \sim \mathrm{Poisson}(\theta)$ then $\kappa^{\mathrm{fac}}(Y; r) = \theta\, \mathbf{1}_{\{r=1\}}$.

Corollary 3.1 (Explicit joint factorial cumulants)**.** *For any finite $I \subset \mathbb{N}$ and integers $(r_i)_{i\in I}$ not all zero,*

$$\kappa^{\mathrm{fac}}\big((A_i)_{i\in I};\ (r_i)_{i\in I}\big) = \sum_{\ell\ge 1}\sum_{\lambda\vdash k}\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\prod_{i\in I}\big(a_{i/\ell}(\lambda)\big)_{r_i},$$

with the convention that $a_q(\lambda) = 0$ if $q \notin \{1,\dots,k\}$ or if $q \notin \mathbb{N}$ (in particular when $\ell \nmid i$).

The sum is finite for fixed I because a term can only be nonzero when $\ell \mid i$ for all $i \in I$, hence $\ell \le \min I$, and there are finitely many $\lambda \vdash k$. The proof is the standard compound-Poisson identity: for each atom $v = v(\ell, \lambda)$, the contribution of $Z_{\ell,\lambda} v$ to the factorial cumulant equals the Poisson mean $\mathbb{P}_\Gamma(\lambda)/\ell$ times the product of the relevant factorial powers of the jump sizes $v_i = a_{i/\ell}(\lambda)$, and cumulants add over independent summands.

If one prefers ordinary (non-factorial) cumulants, one may differentiate the log-mgf in (5). Writing $\kappa(A_{i_1}, \dots, A_{i_r})$ for the joint cumulant, one obtains similarly

$$\kappa(A_{i_1},\dots,A_{i_r}) = \sum_{\ell\ge 1}\sum_{\lambda\vdash k}\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\prod_{q=1}^{r} a_{i_q/\ell}(\lambda),$$

again with the same divisibility convention; this is the specialization of Corollary 3.1 to $r_i \in \{0, 1\}$.

Covariance and a positivity principle. The case $|I| = 2$ yields particularly transparent formulas.

Corollary 3.2 (Covariances)**.** *For any $i, i' \ge 1$,*

$$\mathrm{Cov}(A_i, A_{i'}) = \sum_{\ell\mid \gcd(i,i')}\sum_{\lambda\vdash k}\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\, a_{i/\ell}(\lambda)\, a_{i'/\ell}(\lambda).$$

In particular,

$$\mathrm{Var}(A_i) = \sum_{\ell\mid i}\sum_{\lambda\vdash k}\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\big(a_{i/\ell}(\lambda)\big)^2.$$

All summands in Corollary 3.2 are nonnegative. Consequently, for each pair (i, i') we have the equivalence

$$\mathrm{Cov}(A_i, A_{i'}) = 0 \quad \Longleftrightarrow \quad \text{no atom } v(\ell, \lambda) \text{ satisfies } v_i(\ell, \lambda) > 0 \text{ and } v_{i'}(\ell, \lambda) > 0. \tag{6}$$

This is stronger than what holds for a general multivariate law: here the covariance detects the presence of shared Poisson components because the jumps are supported on $\mathbb{N}^{(\mathbb{N})}$ and the Lévy measure is purely atomic.

The dependence graph and sigma-field factorization. For each index $i \geq 1$ define the set of Poisson components that can affect the ith coordinate by

$$\mathcal{U}_i := \Big\{(\ell, \lambda) : \ \ell \mid i, \ \lambda \vdash k, \ a_{i/\ell}(\lambda) > 0, \ \mathbb{P}_\Gamma(\lambda) > 0\Big\}.$$

Then A_i is measurable with respect to the independent family $(Z_u)_{u \in \mathcal{U}_i}$, and the representation immediately yields:

Proposition 3.3 (Dependence graph criterion)**.** *For distinct $i, i' \geq 1$, the following are equivalent:*

1. *A_i and $A_{i'}$ are independent;*
2. *$\mathcal{U}_i \cap \mathcal{U}_{i'} = \varnothing$;*
3. *there is no pair (ℓ, λ) with $\ell \mid i$ and $\ell \mid i'$ such that $\mathbb{P}_\Gamma(\lambda) > 0$ and $a_{i/\ell}(\lambda), a_{i'/\ell}(\lambda) > 0$;*
4. *$\mathrm{Cov}(A_i, A_{i'}) = 0$.*

It is convenient to encode this as a graph on vertex set $\mathbb{N}$: we connect i and i' by an edge when $\mathcal{U}_i \cap \mathcal{U}_{i'} \neq \varnothing$, equivalently when (6) fails. In this language, Proposition 3.3 states that two coordinates are independent if and only if they are nonadjacent. More generally, if $S, T \subset \mathbb{N}$ are disjoint sets of indices, then the subvectors $(A_i)_{i \in S}$ and $(A_i)_{i \in T}$ are independent if and only if there is no edge between S and T, i.e. if and only if

$$\Big(\bigcup_{i \in S} \mathcal{U}_i\Big) \cap \Big(\bigcup_{i \in T} \mathcal{U}_i\Big) = \varnothing.$$

This is a sigma-field factorization statement: each block of indices depends only on the Poisson components attached to it, and disjointness of these component sets is exactly independence.

Arithmetic blocks from the support of ν_Γ. The dependence criterion becomes especially transparent when we impose structural restrictions on the internal cycle types that occur with positive probability under $\mathbb{P}_\Gamma$. Write

$$\mathrm{Supp}(\mathbb{P}_\Gamma) := \{\lambda \vdash k : \ \mathbb{P}_\Gamma(\lambda) > 0\}, \qquad \mathrm{Parts}(\lambda) := \{j : \ a_j(\lambda) > 0\}.$$

Then an atom $v(\ell, \lambda)$ has support exactly on $\{j\ell : \ j \in \mathrm{Parts}(\lambda)\}$, and all dependence questions reduce to whether the support of ν_Γ contains atoms that hit two prescribed indices simultaneously.

A basic instance is a parity decomposition. Assume that

$$\mathrm{Parts}(\lambda) \subset 2\mathbb{Z} + 1 \qquad \text{for all } \lambda \in \mathrm{Supp}(\mathbb{P}_\Gamma), \tag{7}$$

i.e. every element of Γ has only odd cycle lengths on $[k]$. Then each atom $v(\ell, \lambda)$ is supported either entirely on odd indices (when ℓ is odd) or entirely on even indices (when ℓ is even). Hence no atom of ν_Γ can contribute simultaneously to an odd coordinate and an even coordinate, and we obtain the independence of the two parity blocks:

$$(A_i)_{i \text{ odd}} \ \perp\!\!\!\perp \ (A_i)_{i \text{ even}}.$$

More generally, fix a prime p and suppose that all cycle lengths occurring in Γ are coprime to p, i.e.

$$p \nmid j \quad \text{for all } j \in \mathrm{Parts}(\lambda), \ \lambda \in \mathrm{Supp}(\mathbb{P}_\Gamma). \tag{8}$$

Then for any atom $v(\ell, \lambda)$ and any index i in its support we have $i = j\ell$ with $p \nmid j$, hence $v_p(i) = v_p(\ell)$. It follows that an atom cannot simultaneously charge two indices with different p-adic valuations. Consequently the blocks

$$\mathcal{B}_r^{(p)} := \{i \geq 1 : \ v_p(i) = r\}, \qquad r \geq 0,$$

are independent in the sense that the subvectors $(A_i)_{i \in \mathcal{B}_r^{(p)}}$ are mutually independent over different r. This provides a family of "prime-power pattern" decompositions: the jump support of ν_Γ enforces independence across p-adic layers whenever (8) holds.

At the opposite extreme, if $\mathrm{Supp}(\mathbb{P}_\Gamma)$ contains a partition λ with at least two distinct part sizes $j \neq j'$, then for every ℓ the atom $v(\ell, \lambda)$ forces dependence between the coordinates $j\ell$ and $j'\ell$ (and in fact produces positive covariance by Corollary 3.2). Thus the presence or absence of mixed part sizes in the internal cycle types governs whether dependence propagates along arithmetic rays $\ell\mathbb{N}$.

We emphasize that all such conclusions are read directly from ν_Γ: a proposed decomposition of the index set into blocks yields independent subvectors precisely when no atom in the support of ν_Γ places positive mass on two blocks. This viewpoint will be used in the next section, where we show that the same atomic data in fact determine the internal cycle-type distribution $\mathbb{P}_\Gamma$ itself.

3.2 Identifiability and completeness

We now justify the "completeness" assertion that the limiting law (equivalently the discrete Lévy measure) determines the internal cycle-type distribution

$$\mathbb{P}_\Gamma(\lambda) = \mathbb{P}\big(\mathrm{ctype}(\gamma) = \lambda \text{ for } \gamma \sim \mathrm{Unif}(\Gamma)\big), \qquad \lambda \vdash k.$$

Since all finite-dimensional marginals of A are compound Poisson with finite Lévy measure (Lemma 3), the Lévy–Khintchine representation is unique in each dimension. Concretely, for each $m \geq 1$ the log-Laplace transform

$$\log \mathbb{E}\Big[\exp\Big(\sum_{i=1}^{m} t_i A_i\Big)\Big] = \int_{\mathbb{N}^{(\mathbb{N})}\setminus\{0\}} \Big(\exp\Big(\sum_{i=1}^{m} t_i v_i\Big) - 1\Big)\, \nu_\Gamma^{(m)}(dv)$$

determines the finite measure $\nu_\Gamma^{(m)}$ on the projected state space by standard uniqueness of Laplace transforms for finite measures on a countable set. Passing over all m identifies ν_Γ itself (as a measure on $\mathbb{N}^{(\mathbb{N})} \setminus \{0\}$) because the atoms $v(\ell, \lambda)$ are finitely supported and hence visible in some finite projection. Thus, knowing $\mathrm{Law}(A)$ is equivalent to knowing ν_Γ.

The key point is that ν_Γ decomposes canonically into disjoint "ℓ-slices" which isolate the $\ell = 1$ atoms, and these atoms encode $\mathbb{P}_\Gamma$ without collisions. To make this precise, for $v \in \mathbb{N}^{(\mathbb{N})} \setminus \{0\}$ define

$$d(v) := \gcd\{\, i \geq 1 : \ v_i > 0 \,\}.$$

By Lemma 2 we have, for every atom of our form,

$$d\big(v(\ell, \lambda)\big) = \ell.$$

Hence the sets

$$\mathcal{V}_\ell := \{\, v \in \mathbb{N}^{(\mathbb{N})} \setminus \{0\} : \ d(v) = \ell \,\}, \qquad \ell \geq 1,$$

form a partition of $\mathbb{N}^{(\mathbb{N})} \setminus \{0\}$ into measurable pieces on which the Lévy measure is supported on disjoint families of atoms. In particular, the restriction of ν_Γ to $\mathcal{V}_1$ sees precisely the $\ell = 1$ contributions:

$$\nu_\Gamma \restriction_{\mathcal{V}_1} = \sum_{\lambda \vdash k} \mathbb{P}_\Gamma(\lambda)\, \delta_{v(1,\lambda)}. \tag{9}$$

The weights are exactly $\mathbb{P}_\Gamma(\lambda)$ because $\nu_\Gamma(\{v(\ell, \lambda)\}) = \mathbb{P}_\Gamma(\lambda)/\ell$ by definition.

It remains to argue that the map $\lambda \mapsto v(1, \lambda)$ is injective, so that the atomic masses in (9) can be read unambiguously as the probabilities of the corresponding cycle types. But for $\ell = 1$ we have

$$v(1, \lambda)_j = a_j(\lambda), \qquad 1 \leq j \leq k,$$

so $v(1, \lambda)$ is exactly the cycle-count vector $(a_1(\lambda), \ldots, a_k(\lambda))$ (padded by zeros thereafter). Lemma 1 asserts that λ is uniquely determined by these counts, i.e. λ is recovered by taking $a_j(\lambda)$ parts of size j for each j. Thus the atoms $v(1, \lambda)$ are all distinct as elements of $\mathbb{N}^{(\mathbb{N})}$, and we conclude:

Theorem 3.4 (Identifiability at fixed block size). *Fix k. If $\Gamma, \Gamma' \le S_k$ satisfy $\nu_\Gamma = \nu_{\Gamma'}$ (equivalently $\mathrm{Law}(A^\Gamma) = \mathrm{Law}(A^{\Gamma'})$), then*

$$\mathbb{P}_\Gamma(\lambda) = \mathbb{P}_{\Gamma'}(\lambda) \qquad \textit{for all } \lambda \vdash k.$$

Moreover, given ν_Γ we recover $\mathbb{P}_\Gamma$ by the explicit formula

$$\mathbb{P}_\Gamma(\lambda) = \nu_\Gamma\big(\{v(1,\lambda)\}\big), \qquad \lambda \vdash k.$$

Proof. By the gcd-slicing described above, ν_Γ determines its restriction to $\mathcal{V}_1$, which by (9) is the purely atomic measure $\sum_\lambda \mathbb{P}_\Gamma(\lambda)\delta_{v(1,\lambda)}$. Since $\lambda \mapsto v(1,\lambda)$ is injective (Lemma 1), the mass of the singleton atom $\{v(1,\lambda)\}$ equals $\mathbb{P}_\Gamma(\lambda)$. The same argument applies to Γ', and equality of Lévy measures yields equality of the recovered masses. ☐

We emphasize what this does and does not say about Γ as a subgroup. The recovered data are precisely the proportions of elements of Γ in each conjugacy class of S_k, i.e. the class function

$$\lambda \longmapsto \mathbb{P}_\Gamma(\lambda) = \frac{|\Gamma \cap C_\lambda|}{|\Gamma|},$$

where $C_\lambda \subset S_k$ denotes the conjugacy class of cycle type λ. Equivalently, we recover the cycle index polynomial

$$Z_\Gamma(x_1,\dots,x_k) := \frac{1}{|\Gamma|}\sum_{\gamma\in\Gamma}\prod_{j=1}^{k} x_j^{a_j(\gamma)} = \sum_{\lambda\vdash k}\mathbb{P}_\Gamma(\lambda)\prod_{j=1}^{k} x_j^{a_j(\lambda)}.$$

Thus, the limit law determines exactly the averaged cycle structure of Γ (and conversely, by Theorem A, this averaged cycle structure determines the limit law).

However, $\mathbb{P}_\Gamma$ is much weaker than the subgroup itself. First, it is invariant under conjugation: if $\Gamma' = \tau\Gamma\tau^{-1}$ for some $\tau \in S_k$, then Γ and Γ' have identical intersections with each conjugacy class, hence the same $\mathbb{P}_\Gamma$ and the same limit law. More significantly, it may happen that two nonconjugate subgroups $\Gamma, \Gamma' \le S_k$ have the same class distribution $\mathbb{P}_\Gamma = \mathbb{P}_{\Gamma'}$; such pairs are often called *almost conjugate* or *Gassmann equivalent.* For such a pair, no statistic built solely from the cycle type of a uniform subgroup element can distinguish Γ from Γ', and in particular our limiting cycle-count vector A cannot do so. Theorem 3.4 therefore identifies Γ only up to the equivalence relation

$$\Gamma \sim \Gamma' \quad \Longleftrightarrow \quad \mathbb{P}_\Gamma = \mathbb{P}_{\Gamma'}.$$

One can nevertheless view the preceding theorem as a completeness statement for the limit: there is no additional "hidden" parameter of Γ influencing $\mathrm{Law}(A)$ beyond $\mathbb{P}_\Gamma$. In practical terms, any invariant of Γ that can be expressed as an average of a class function on S_k (for instance, the expected number of j-cycles in a uniform element of Γ, or the probability of having a fixed point) is determined by the limit law. By contrast, invariants depending on the internal multiplication structure of Γ (e.g. whether Γ is abelian, solvable, or simple) are generally not determined by $\mathbb{P}_\Gamma$ and hence are invisible to A.

Finally, we remark that the identifiability argument is genuinely arithmetic: the ability to isolate $\ell = 1$ rests on the fact that each atom $v(\ell, \lambda)$ has its support contained in $\ell\{1, 2, \ldots, k\}$, so its support-gcd is exactly ℓ. This separation mechanism fails for more general infinitely divisible laws where different jump types may overlap in a way that cannot be disentangled from the full Lévy measure. Here the imprimitive block structure forces a rigid support pattern, and the $\ell = 1$ slice functions as a "fingerprint" of the internal cycle-type distribution.

In the next section we exploit the explicit formulas above in concrete cases, computing $\mathbb{P}_\Gamma$ (hence ν_Γ and the dependence graph) for several standard choices of Γ and illustrating how independence and dependence manifest in the coordinates $(A_i)_{i\geq 1}$.

3.3 Worked examples and computations

We record three concrete families illustrating how the internal distribution $\mathbb{P}_\Gamma$ propagates to the Lévy atoms $v(\ell, \lambda)$ and hence to the dependence pattern among the coordinates $(A_i)_{i\geq 1}$. Throughout we use the Poisson–sum representation

$$A \;=\; \sum_{\ell\geq 1} \sum_{\lambda\vdash k} Z_{\ell,\lambda}\, v(\ell, \lambda), \qquad Z_{\ell,\lambda} \sim \text{Poisson}\left(\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\right) \;\; \text{independent.}$$

3.3.1 The cyclic case $\Gamma = C_k$: independent coordinates

Let $\Gamma = C_k = \langle c\rangle \leq S_k$, where c is a k-cycle. The cycle type of c^r is determined by $d = \gcd(k, r)$: it consists of exactly d disjoint cycles each of length k/d, i.e. the partition

$$\lambda_d \;=\; (k/d)^d \qquad (d \mid k).$$

The number of residues $r \in \{0, 1, \ldots, k-1\}$ with $\gcd(k, r) = d$ equals $\varphi(k/d)$, so

$$\mathbb{P}_{C_k}(\lambda_d) \;=\; \frac{\varphi(k/d)}{k}, \qquad \mathbb{P}_{C_k}(\lambda) = 0 \;\text{ if } \lambda \neq \lambda_d \;\forall d \mid k. \tag{10}$$

For $\lambda = \lambda_d$ we have $a_{k/d}(\lambda_d) = d$ and $a_j(\lambda_d) = 0$ for $j \neq k/d$. Consequently each increment vector $v(\ell, \lambda_d)$ is supported on a *single* coordinate:

$$v(\ell, \lambda_d)_{(k/d)\ell} = d, \qquad v(\ell, \lambda_d)_i = 0 \;\text{ for } i \neq (k/d)\ell.$$

In particular, no Lévy atom simultaneously affects two distinct coordinates. By the dependence criterion (Proposition B), it follows that all coordinates $(A_i)_{i\geq 1}$ are mutually independent.

It is useful to make the one-dimensional laws explicit. Fix $i \geq 1$. The only contributions to A_i come from pairs (ℓ, d) with $d \mid k$ and $(k/d)\ell = i$, i.e. $\ell = id/k$ (so necessarily $k/d \mid i$). Writing $j = k/d$ (so $j \mid k$), we obtain

$$A_i \;=\; \sum_{\substack{j\mid k\\ j\mid i}} \frac{k}{j}\, Z_{\,i/j,\;\lambda_{k/j}}, \qquad Z_{\,i/j,\;\lambda_{k/j}} \sim \mathrm{Poisson}\!\left(\frac{\varphi(j)}{k}\cdot\frac{j}{i}\right) = \mathrm{Poisson}\!\left(\frac{\varphi(j)}{i}\right), \tag{11}$$

independently across different j. Thus A_i is a (one-dimensional) compound Poisson variable with jump sizes $\{k/j : j \mid k,\, j \mid i\}$, and the family is independent across i. In particular,

$$\mathbb{E}[A_i] = \sum_{\substack{j\mid k\\ j\mid i}} \frac{k}{j}\cdot\frac{\varphi(j)}{i} = \frac{1}{i}\sum_{\substack{j\mid k\\ j\mid i}} \varphi(j)\,\frac{k}{j}.$$

When k is prime, (11) reduces to a sum of at most two independent Poisson terms, corresponding to $j = 1$ and $j = k$.

3.3.2 The full symmetric group $\Gamma = S_k$: explicit dependence via shared atoms

Now take $\Gamma = S_k$. Then the cycle type of a uniform element is distributed according to conjugacy class sizes:

$$\mathbb{P}_{S_k}(\lambda) = \frac{|C_\lambda|}{k!} = \frac{1}{z_\lambda}, \qquad z_\lambda := \prod_{j=1}^{k} j^{a_j(\lambda)}\, a_j(\lambda)!. \tag{12}$$

In this case many partitions λ have multiple nonzero counts $(a_1(\lambda), a_2(\lambda), \dots)$, hence many Lévy atoms $v(\ell, \lambda)$ have support on several coordinates $\ell, 2\ell, \dots, k\ell$, producing systematic dependence.

A basic instance is the dependence between A_j and A_{2j} (assuming $k \geq 2$). Consider any partition $\lambda \vdash k$ with $a_1(\lambda) > 0$ and $a_2(\lambda) > 0$; for example $\lambda = (1, 2, 1^{k-3})$ exists for all $k \geq 3$. For such λ and $\ell = j$ we have

$$v(j,\lambda)_j = a_1(\lambda) > 0, \qquad v(j,\lambda)_{2j} = a_2(\lambda) > 0,$$

so the common Poisson component $Z_{j,\lambda}$ appears in both A_j and A_{2j}. Hence A_j and A_{2j} are not independent whenever

$$\sum_{\lambda:\; a_1(\lambda)>0,\; a_2(\lambda)>0} \mathbb{P}_{S_k}(\lambda) \;>\; 0,$$

which holds for every $k \geq 3$ (and for $k = 2$ the dependence is instead between A_j and A_{2j} through $\lambda = (2)$ with $a_2 = 1$ and $\lambda = (1, 1)$ with $a_1 = 2$, affecting different coordinates but still generating nontrivial structure across multiples).

Covariances can be computed directly from Corollary D. For $i, i' \geq 1$,

$$\mathrm{Cov}(A_i, A_{i'}) = \sum_{\ell | \gcd(i,i')} \frac{1}{\ell} \sum_{\lambda \vdash k} \frac{1}{z_\lambda} a_{i/\ell}(\lambda)\, a_{i'/\ell}(\lambda), \tag{13}$$

with the convention that $a_r(\lambda) = 0$ for $r > k$. In particular, for $k \geq 3$ and $i' = 2i$ we see that the divisor $\ell = i$ contributes

$$\frac{1}{i} \sum_{\lambda \vdash k} \frac{1}{z_\lambda} a_1(\lambda)\, a_2(\lambda),$$

which is strictly positive, proving $\mathrm{Cov}(A_i, A_{2i}) > 0$ without further calculation. More generally, dependence is prevalent along the divisor poset: if i and i' share a nontrivial gcd ℓ and there exists λ with both $a_{i/\ell}(\lambda)$ and $a_{i'/\ell}(\lambda)$ nonzero, then the shared components at slice ℓ enforce correlation.

We also note that (12) permits closed forms for many averaged internal statistics. For instance, if $\gamma \sim \mathrm{Unif}(S_k)$ then $\mathbb{E}[a_j(\gamma)] = 1/j$ for $1 \leq j \leq k$. Plugging this into the Lévy representation yields

$$\mathbb{E}[A_i] = \sum_{\ell | i} \sum_{\lambda \vdash k} \frac{1}{\ell} \cdot \frac{1}{z_\lambda} a_{i/\ell}(\lambda) = \sum_{\substack{\ell | i \\ i/\ell \leq k}} \frac{1}{\ell} \cdot \frac{1}{i/\ell} = \frac{1}{i} \sum_{\substack{\ell | i \\ i/\ell \leq k}} 1,$$

so $\mathbb{E}[A_i] = d_{\leq k}(i)/i$, where $d_{\leq k}(i)$ counts divisors ℓ of i with $i/\ell \leq k$.

3.3.3 A small-k table: $k = 3$

For $k = 3$ the partitions are

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

Writing $p_1 = \mathbb{P}_\Gamma(1^3)$, $p_2 = \mathbb{P}_\Gamma(1\,2)$, $p_3 = \mathbb{P}_\Gamma(3)$ (so $p_1 + p_2 + p_3 = 1$), the corresponding increment vectors satisfy, for each $\ell \geq 1$,

$$v(\ell, 1^3)_\ell = 3, \qquad v(\ell, 1\,2)_\ell = 1, \;\; v(\ell, 1\,2)_{2\ell} = 1, \qquad v(\ell, 3)_{3\ell} = 1,$$

with all other coordinates equal to 0. Therefore the only Lévy atoms that simultaneously affect two coordinates are those of type $(\ell, 1\,2)$, and they always link the pair $(\ell, 2\ell)$. By Proposition B we obtain the equivalence

$$A_\ell \text{ and } A_{2\ell} \text{ are dependent} \quad \Longleftrightarrow \quad p_2 > 0,$$

while A_ℓ is independent of $A_{3\ell}$ regardless of p_3, since the type 3 atom only hits the coordinate 3ℓ and the type 1^3 atom only hits ℓ.

Two specializations are worth recording. If $\Gamma = C_3$, then $p_1 = p_3 = 1/3$ and $p_2 = 0$ (since no power of a 3-cycle has cycle type 1 2), and hence all coordinates are independent, consistent with the cyclic discussion above. If $\Gamma = S_3$, then by class sizes we have $p_1 = 1/6$, $p_2 = 1/2$, $p_3 = 1/3$, so A_ℓ and $A_{2\ell}$ are dependent for every ℓ. Moreover, the covariance is particularly simple: from Corollary D,

$$\mathrm{Cov}(A_\ell, A_{2\ell}) = \frac{1}{\ell} \sum_{\lambda \vdash 3} \mathbb{P}_{S_3}(\lambda)\, a_1(\lambda)\, a_2(\lambda) = \frac{1}{\ell} \cdot \frac{1}{2} \cdot 1 \cdot 1 = \frac{1}{2\ell}.$$

Optional computational verification (computational). For fixed $k \leq 6$ one can verify these formulas by brute force in a computer algebra system. Given $\Gamma \leq S_k$ specified by generators, we enumerate $\gamma \in \Gamma$, compute its cycle type, and tabulate $\mathbb{P}_\Gamma(\lambda)$ over $\lambda \vdash k$. This immediately yields the Lévy atoms and their weights, hence (i) predicted marginal means and covariances via Corollary D, and (ii) the dependence graph via Proposition B. Independently, for large n we may sample $\sigma_n \sim \mathrm{Unif}(\Gamma^n \rtimes S_n)$ by drawing $\pi \sim \mathrm{Unif}(S_n)$ and independent $\gamma_1, \dots, \gamma_n \sim \mathrm{Unif}(\Gamma)$ and forming the standard imprimitive action, then computing cycle counts $(a_i(\sigma_n))_{1 \leq i \leq m}$ for moderate m. Empirical moments and pairwise correlations converge rapidly in n for small m, and the presence or absence of dependence typically matches the Lévy-support criterion already at modest n (e.g. $n \approx 10^3$), while quantitative agreement for covariances improves as n grows.

3.4 Extensions and remarks

7.1. Replacing S_n by other block-permutation models. Our arguments isolate the contribution of the *block permutation* and the *internal* Γ–types through the same combinatorial mechanism: each ℓ–cycle of the block permutation yields an increment of the form $v(\ell, \lambda)$, where λ is the cycle type of the product of the internal labels around that block cycle. For σ_n uniform on $\Gamma^n \rtimes S_n$ the asymptotic number of ℓ–cycles in the block permutation is Poisson$(1/\ell)$, which is exactly the origin of the factor $1/\ell$ in ν_Γ.

This suggests an immediate generalization in which the block permutation π_n is sampled from a conjugacy-invariant measure on S_n whose small-cycle counts converge to independent Poisson variables with means α_ℓ/ℓ, for some prescribed weights $(\alpha_\ell)_{\ell \geq 1}$ with $\alpha_\ell \geq 0$. A canonical example is the Ewens measure with parameter $\theta > 0$, for which $\alpha_\ell \equiv \theta$. More generally, one may consider logarithmic combinatorial structures (in the sense of Arratia–Barbour–Tavaré) where the cycle index has the asymptotic factorization property leading to Poisson limits for fixed ℓ.

Assume concretely that for each fixed L,

$$\big(C_1^{(n)}, \dots, C_L^{(n)}\big) \;\Rightarrow\; \big(C_1, \dots, C_L\big), \qquad C_\ell \sim \mathrm{Poisson}\Big(\frac{\alpha_\ell}{\ell}\Big) \;\; \text{independent},$$

where $C_\ell^{(n)}$ counts ℓ–cycles of π_n. If we then sample independent internal labels $\gamma_1, \dots, \gamma_n$ from a fixed distribution on Γ (uniform being the special case) and form the standard imprimitive action, the same decomposition by block cycles yields, for each fixed m, a limiting vector $(A_1, \dots, A_m)$ which remains compound Poisson with *the same* increment family $\{v(\ell, \lambda)\}$ but with reweighted intensity

$$\nu_{\Gamma,\alpha}\big(\{v(\ell,\lambda)\}\big) \;=\; \frac{\alpha_\ell}{\ell}\,\mathbb{P}_\Gamma(\lambda). \tag{14}$$

Equivalently, in the Poisson–sum representation we simply replace

$$Z_{\ell,\lambda} \sim \mathrm{Poisson}\bigg(\frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\bigg) \quad \text{by} \quad Z_{\ell,\lambda} \sim \mathrm{Poisson}\bigg(\frac{\alpha_\ell\,\mathbb{P}_\Gamma(\lambda)}{\ell}\bigg),$$

still independent over (ℓ, λ). In this form, the dependence criterion (Proposition B) is unchanged at the level of *support* (only weights change): two coordinates are linked if and only if there exists an increment affecting both. The weights (α_ℓ) only determine the *strength* of dependence through the magnitudes of the shared Poisson components.

Two remarks are in order. First, the identifiability statement (Theorem C) adapts provided $\alpha_1 > 0$: the $\ell = 1$ slice still exposes $\mathbb{P}_\Gamma(\lambda)$ through the injective map $\lambda \mapsto v(1, \lambda)$. If $\alpha_1 = 0$, then the $\ell = 1$ slice disappears and identifiability must be reformulated in terms of the smallest ℓ with $\alpha_\ell > 0$. Second, if the internal labels are not uniform on Γ but are i.i.d. with some law μ on Γ, then $\mathbb{P}_\Gamma(\lambda)$ in (14) is replaced by the induced cycle-type distribution $\mathbb{P}_\mu(\lambda)$ of a μ–distributed element of Γ; no other modification is needed.

7.2. A Stein operator suggested by the Lévy measure. The explicit Lévy measure provides a natural route to quantitative approximation of $(a_1(\sigma_n), \dots, a_m(\sigma_n))$ by its limit $(A_1, \dots, A_m)$ via Stein's method for compound Poisson approximation. For the m–dimensional marginal, write the projected increment vectors as

$$u_{\ell,\lambda}^{(m)} \;:=\; \big(v(\ell,\lambda)_1, \dots, v(\ell,\lambda)_m\big) \in \mathbb{N}^m,$$

and define the finite measure

$$\nu_\Gamma^{(m)} \;:=\; \sum_{\ell \geq 1} \;\sum_{\lambda \vdash k:\ \min\{j\ell:\ a_j(\lambda)>0\} \leq m} \frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\,\delta_{u_{\ell,\lambda}^{(m)}}.$$

Then $(A_1, \dots, A_m)$ is the stationary distribution of the pure-jump process with generator

$$(\mathcal{A}f)(x) \;=\; \sum_{u \in \mathbb{N}^m \setminus \{0\}} \nu_\Gamma^{(m)}(\{u\})\big(f(x+u) - f(x)\big), \qquad x \in \mathbb{N}^m. \tag{15}$$

Accordingly, for a test function h on $\mathbb{N}^m$ one may seek f solving the Stein equation

$$\mathcal{A}f(x) \;=\; h(x) - \mathbb{E}\big[h(A_1,\ldots,A_m)\big],$$

and then bound

$$\big|\mathbb{E}h(a_1(\sigma_n),\ldots,a_m(\sigma_n)) - \mathbb{E}h(A_1,\ldots,A_m)\big| \;=\; \big|\mathbb{E}(\mathcal{A}f)(a_1(\sigma_n),\ldots,a_m(\sigma_n))\big|$$

by controlling the discrepancy between the true jump structure of the finite-n object and the idealized jumps encoded in $\nu_\Gamma^{(m)}$. The dependence graph in Proposition B is useful here: it identifies exactly which coordinates can be simultaneously altered by a single jump, hence which local couplings are required. In particular, bounds in Wasserstein or total variation metrics for compound Poisson approximation typically depend on (i) the total intensity $\nu_\Gamma^{(m)}(\mathbb{N}^m \setminus \{0\})$, and (ii) the maximal jump size $\max\{\|u\|_1 : \nu_\Gamma^{(m)}(\{u\}) > 0\}$, both of which are explicit from the support description $u_{\ell,\lambda}^{(m)}$.

A second, complementary approach is to couple block cycles directly: in the wreath product, small cycles in $[kn]$ are created by short block cycles in $[n]$ together with internal types. Since short block cycles are rare and nearly independent, one can often obtain explicit rates of convergence (in n) by truncating to block cycles of length at most m and estimating the error of excluding longer block cycles, which contribute only to coordinates $> m$.

7.3. Large deviations and exponential tilting. Because we possess the log-mgf in closed form, finite-dimensional large deviation statements for the limit law follow formally from standard convex duality. For $m \geq 1$ define

$$\kappa_m(t_1,\ldots,t_m) \;:=\; \log \mathbb{E}\Big[\exp\Big(\sum_{i=1}^m t_i A_i\Big)\Big] \;=\; \sum_{\ell\geq 1}\sum_{\lambda\vdash k} \frac{\mathbb{P}_\Gamma(\lambda)}{\ell}\Big(\exp\Big(\sum_{j=1}^k t_{j\ell} a_j(\lambda)\Big) - 1\Big),$$

with $t_r = 0$ for $r > m$. The convex conjugate

$$I_m(x) \;:=\; \sup_{t\in\mathbb{R}^m} \big\{\langle t, x\rangle - \kappa_m(t)\big\}, \qquad x \in \mathbb{R}^m,$$

is the natural rate function governing exponential tilts of $(A_1,\ldots,A_m)$ and tail estimates via Chernoff bounds. While $(A_1,\ldots,A_m)$ itself is not a scaled quantity, these bounds are practically effective for rare-event probabilities such as $\mathbb{P}(A_i \geq r)$ for large r, and they make transparent how the extremal behavior depends on the set of feasible jumps $u_{\ell,\lambda}^{(m)}$.

For the original finite-n models, one may ask for large deviations of the *small-cycle process* $(a_1(\sigma_n),\ldots,a_m(\sigma_n))$ at a scale depending on n. Since these counts converge to an $O(1)$ limit, a classical n–speed LDP is not expected for fixed m; however, an LDP can emerge when $m = m(n)$ grows (for instance $m = o(n)$), or when one studies the *empirical measure* of cycle lengths up to a cutoff. In such regimes the Lévy measure viewpoint remains useful: it suggests that deviations are driven by atypical numbers of short block cycles (a deviation in the block model) combined with atypical internal types (a deviation in the Γ–labels), which can often be separated by exponential tilting at the level of cycle indices.

7.4. Open problems. We conclude with a non-exhaustive list of directions suggested by the preceding structure.

1. *Quantitative convergence.* Establish explicit error bounds, uniform in n, for the approximation of $(a_1(\sigma_n), \dots, a_m(\sigma_n))$ by $(A_1, \dots, A_m)$ in strong metrics (e.g. total variation). The Lévy atoms indicate that the relevant couplings should be organized by the block-cycle structure, but sharp constants appear to require refined control of the dependence created by collisions among short block cycles.

2. *Partial identifiability from truncated data.* Theorem C uses the full Lévy measure (equivalently, all coordinates) to recover $\mathbb{P}_\Gamma$. Given only the law of $(A_1, \dots, A_m)$ for fixed m, determine which linear functionals of $\mathbb{P}_\Gamma$ are identifiable, and characterize the kernel of the map $\mathbb{P}_\Gamma \mapsto \mathrm{Law}(A_1, \dots, A_m)$.

3. *Classification of dependence graphs.* Proposition B reduces dependence to the existence of atoms simultaneously hitting two coordinates. For fixed k, which graphs on $\mathbb{N}$ can occur as the dependence graph of (A_i) as Γ varies over subgroups of S_k (or more generally as $\mathbb{P}_\Gamma$ varies over class measures supported on partitions of k)?

4. *Beyond fixed block size.* Our analysis keeps k fixed as $n \to \infty$. If $k = k(n) \to \infty$, the internal cycle-type space changes with n and the current Lévy description no longer applies verbatim. A natural problem is to find scaling regimes in which a limiting Lévy measure persists (possibly on a different state space), and to understand whether universality phenomena appear.

5. *Other actions and other wreath products.* The standard imprimitive action is particularly rigid. It would be useful to develop analogous Lévy descriptions for other permutation representations of $\Gamma \wr S_n$, and to determine which features (such as gcd-slicing of atoms) survive.

In all these questions, the guiding principle is that the limit is governed by an explicitly enumerable set of jumps $v(\ell, \lambda)$ together with their intensities; any refinement must ultimately quantify how rapidly the finite-n cycle structure approaches this ideal compound Poisson superposition.