Boundedly exchangeable graphs

Leon-Roman Rinke * Stefan Weber

Leibniz Universität Hannover

August 18, 2026

Abstract

An infinite jointly exchangeable adjacency matrix is a mixture of graphon models, and its graph is almost surely empty or dense; classical matrix exchangeability therefore excludes sparse nonempty graphs. We study a point-process alternative, *graph processes*, whose vertices are identified with the atoms of a locally finite point process on a Polish latent space and whose edges are drawn conditionally independently from a symmetric connection kernel W. For a Poisson graph process with diffuse σ-finite intensity, we show that its edge measure is jointly exchangeable as a random measure if and only if W is essentially constant. Non-constant kernels therefore require a different exchangeability notion. We introduce *bounded exchangeability*, under which every nontrivial bounded restriction determines an infinite jointly exchangeable adjacency matrix that is unique in distribution, with an associated local graphon representation. We then study *threshold models* on homogeneous Riemannian manifolds, in which Poisson-distributed vertices are connected whenever their intrinsic distance is at most a fixed threshold. On every infinite-volume manifold in this class, the model is sparse, its empirical degree distribution converges to a Poisson law, and its limiting clustering coefficient depends on an additional geometric parameter that can vary while the limiting degree law is held fixed.

Keywords: Random graphs, infinite graphs, exchangeable arrays, latent space models, random connection models

2020 Mathematics Subject Classification: Primary 05C80, 60G09; secondary 60D05, 60G55, 05C63.

*Corresponding author. Leibniz Universität Hannover, House of Insurance, Welfengarten 1, 30167 Hannover, Germany. E-mail: leon-roman.rinke@insurance.uni-hannover.de

1 Introduction

Random graph models often impose exchangeability to express the idea that vertex labels carry no information. For an infinite adjacency matrix $A = (A_{ij})_{i,j\geq 1}$, joint exchangeability means invariance in distribution under simultaneous permutations of rows and columns.

The Aldous–Hoover–Kallenberg representation, due to Hoover (1979), Aldous (1981), and Kallenberg (1989), characterizes this symmetry as a mixture of graphon models: in the ergodic case, an exchangeable graph can be generated by assigning independent latent variables U_i to the vertices and then connecting vertices i and j independently with probability $W(U_i, U_j)$. The symmetric measurable function W is the corresponding graphon.

This classical symmetry has a structural limitation. Infinite jointly exchangeable graphs are almost surely either empty or dense (see, e.g., Orbanz and Roy (2015)), and hence do not describe sparse graphs whose edge count grows linearly in the number of vertices — the regime of many real-world networks (Newman, 2018). The graphex framework of Veitch and Roy (2015) and Caron and Fox (2017) resolves this by changing the object of exchangeability: it imposes exchangeability not on the adjacency matrix, but on a random measure that encodes the edges. This yields sparse infinite graphs, but the resulting symmetry is no longer ordinary matrix exchangeability.

The present paper takes an alternative approach by employing a point-process construction of spatial random graphs. It demonstrates that classical matrix exchangeability is recovered on bounded observation windows. A *graph process* is generated from two ingredients: a locally finite point process on a latent Polish space S and a symmetric connection kernel $W : S^2 \to [0, 1]$. Vertices are encoded by the atoms of the point process, and, given their locations (S_i), edges between distinct vertices are drawn conditionally independently as Bernoulli variables with parameters $W(S_i, S_j)$. The global graph may be sparse and non-exchangeable, but its bounded restrictions can still determine ordinary jointly exchangeable adjacency laws.

The local viewpoint is also natural in applications. Vertices are typically observed through a finite spatial, temporal, or relational window, while the underlying population is idealized as large or unbounded. Sparsity is then a global scaling feature, whereas exchangeability is often plausible only after restricting attention to a bounded frame of observation in which labels carry no information. At the same time, the latent space may carry geometric content: distance can encode similarity, exposure, or interaction cost, and the volume growth of the space affects the number of potential neighbours. This makes it natural to ask for models in which local exchangeability, sparsity, and geometric structure are treated within the same framework.

This local viewpoint separates two questions. The first is structural: when does a point-process graph admit a local form of classical matrix exchangeability? The answer developed below is bounded exchangeability, under which every nontrivial bounded restriction determines an associated infinite jointly exchangeable adjacency matrix, unique in distribution, and hence a local graphon representation. The second question is geometric: when the latent space carries intrinsic notions of volume and distance, how do these geometric quantities determine the limiting graph statistics?

We address the second question on Riemannian manifolds. The vertex process is defined from Riemannian volume, and the connection rule is defined from intrinsic distance. In the threshold model, the manifold analogue of the random plane network of Gilbert (1961), two vertices are connected exactly when their Riemannian distance is at most a fixed radius t. On every homogeneous infinite-volume Riemannian manifold this gives a sparse, geometrically interpretable class of graph processes whose limiting degree law is controlled by the volume V_t of a ball of radius t, while the limiting clustering coefficient carries additional geometric information through p_t, the probability that two independent uniform points of such a ball lie within distance t. The limiting degree law depends on δV_t, whereas the limiting clustering coefficient also depends on p_t. Consequently, after matching δV_t and hence the limiting degree

distribution, hyperbolic space has a strictly smaller limiting clustering coefficient than Euclidean space.

1.1 Contributions

(i) **Graph processes and graphex exchangeability.** We formulate graph processes as point-process-driven random graphs: the vertices are encoded by the atoms of a locally finite point process on a latent Polish space, and edges between distinct vertices are drawn conditionally independently from a symmetric connection kernel W. Under the standing assumptions of the graphex framework, we prove that for a Poisson graph process with diffuse σ-finite intensity measure, its edge measure — which records only present edges — is jointly exchangeable in the random-measure sense if and only if its connection kernel is essentially constant. Hence nonconstant connection kernels are incompatible with graphex-type exchangeability within this class.

(ii) **Bounded exchangeability and local graphon representations.** We introduce *bounded exchangeability* as a local form of classical matrix exchangeability. For every nontrivial bounded observation window B, a boundedly exchangeable graph process determines an infinite jointly exchangeable adjacency matrix A^B, unique in distribution. This associated exchangeable law has an Aldous–Hoover–Kallenberg representation with local graphon $W_B = \widetilde{W} \circ (f_B \otimes f_B)$, where f_B samples from the local vertex distribution on B and $\widetilde{W}(x,y) \coloneqq \mathbf{1}\{x \neq y\}W(x,y)$ is the zero-diagonal version of W. We also describe how the local graphons and exchangeable matrices transform when observation windows are enlarged, restricted, or partitioned into disjoint subwindows.

(iii) **Intrinsic graph processes on Riemannian manifolds.** We single out a geometrically natural class of graph processes on complete connected d-dimensional Riemannian manifolds (M, g), taking the vertex intensity to be λ_g for the Riemannian volume measure λ_g and using the intrinsic distance d_g in the connection rule. For distance-based kernels $W(x,y) = h(d_g(x,y))$, we prove a scale-invariance theorem.

(iv) **Sparsity, degree distribution, clustering, and curvature.** For the threshold model on homogeneous infinite-volume manifolds, we prove asymptotic sparsity and identify the limiting degree and clustering statistics. The empirical degree distribution converges pointwise to a Poisson law determined by the expected number of points in a ball of connection radius. The clustering coefficient converges in probability to a limit that depends both on this expected local number of points and on the geometry of such balls. More precisely, the remaining geometric contribution is governed by how likely two independently chosen points in a connection ball are themselves close enough to be connected. Thus the limiting degree law is fixed by the expected number of points in a connection ball, whereas the clustering limit also reflects the shape of the underlying space. In the Euclidean–hyperbolic comparison, the degree law can be matched by matching these expected local point counts, but the clustering limit remains curvature-dependent; in particular, hyperbolic space yields strictly lower limiting clustering than Euclidean space.

1.2 Literature

The probabilistic study of random graphs began in earnest with the foundational models of Erdős and Rényi (1959) and Gilbert (1959), in which edges are formed independently between a fixed number of vertices. A spatial counterpart was introduced shortly thereafter by Gilbert (1961), whose random plane network places vertices according to a Poisson point process in the Euclidean plane and connects pairs of points lying within a fixed distance. This model initiated a large body of work on random geometric graphs and continuum percolation; see,

for example, Meester and Roy (1996) and Penrose (2003) for comprehensive accounts. For the random geometric graph, limit theorems for length-power functionals, including the edge count, were obtained by Reitzner et al. (2017), and Grygierek and Thäle (2020) derived a quantitative central limit theorem for edge counts in the high-dimensional regime. More general distance-dependent connection rules are treated in the random connection model, where pairs of Poisson points are connected independently with probability depending on their separation (Meester and Roy, 1996; Meester et al., 1997); central limit theorems for its counts of finite components were established by Last et al. (2021). Limit theory for related Poisson-based graphs extends beyond the Euclidean setting: Sambale et al. (2025) prove quantitative central limit theorems for the nearest neighbour embracing graph in Euclidean and hyperbolic space, and Rosen et al. (2025) study the degree of the origin and edge-length functionals of the hyperbolic radial spanning tree; in both cases the hyperbolic results are contrasted with their Euclidean counterparts. The threshold graph processes studied in this paper may be viewed as intrinsic, manifold-valued analogues of these spatial random graph models.

A second line of work relevant to the present paper concerns exchangeability. The probabilistic theory of exchangeable arrays was developed by Hoover (1979), Aldous (1981), and Kallenberg (1989). In the setting of graphs, this theory is closely related to the graphon framework. Graphons were introduced by Lovász and Szegedy (2006) as limit objects for convergent sequences of dense graphs, with convergence expressed in terms of subgraph densities. Borgs et al. (2006) and Borgs et al. (2008) placed this theory on a metric footing through the cut metric, while Lovász and Szegedy (2010) further developed the topological structure of the space of graphons. Lovász (2012) provides a comprehensive account of the resulting theory of large dense graph limits. Diaconis and Janson (2008) made explicit the connection between graph limits and the Aldous–Hoover–Kallenberg representation of jointly exchangeable random arrays. From this perspective, an infinite jointly exchangeable adjacency matrix can be represented via a graphon. This representation is powerful but intrinsically dense: a jointly exchangeable infinite random graph is almost surely either empty or dense; see Orbanz and Roy (2015).

The graphex framework provides one of the principal resolutions of this tension between exchangeability and sparsity. Building on the representation theory for exchangeable random measures due to Kallenberg (1990), Caron and Fox (2017) and Veitch and Roy (2015) constructed sparse random graphs from exchangeable random measures on $\mathbb{R}_+^2$. Veitch and Roy (2015) identified the graphex as the canonical representation for such models and established integrability conditions ensuring local finiteness. Subsequent work developed a corresponding sparse graph limit theory: Borgs et al. (2018) introduced graphon processes and studied convergence, Janson (2022) developed graphex convergence, Borgs et al. (2019) studied sampling consistency, and Veitch and Roy (2019) addressed estimation. Different sparsity regimes are considered by Caron et al. (2023). The present paper is complementary to this literature. Rather than imposing exchangeability on an edge measure and discarding isolated vertices, we instead ask how classical matrix exchangeability may be recovered on bounded restrictions of the underlying point process.

Another related area is latent-space network modelling. In the model of Hoff et al. (2002), vertices are assigned unobserved positions in a Euclidean space and connection probabilities depend on latent distances or projections. This approach was extended by Handcock et al. (2007) to include clustering and by Salter-Townshend and McCormick (2017) to multiview network data. The stochastic blockmodel of Holland et al. (1983) may be viewed as a latent-space model with a discrete latent space. Non-Euclidean latent spaces have also received substantial attention. Hyperbolic latent-space models were proposed by Krioukov et al. (2010) and Boguñá et al. (2010), motivated by the emergence of heterogeneous degrees and strong clustering in negatively curved spaces. Rigorous results on degree sequences and clustering in random hyperbolic graphs were obtained by Gugelmann et al. (2012). Further work on hyperbolic network geometry includes the geometric renormalisation procedure of García-Pérez et al. (2018). Spherical

latent spaces have been studied, for example, by McCormick and Zheng (2015) in the context of aggregated relational data. For surveys on geometric and latent-space network models, see Smith et al. (2019) and Boguñá et al. (2021).

Inference for latent geometries has also been studied extensively. The problem of estimating the curvature of a latent space is addressed by Lubold et al. (2023) and Wilkins-Reeves and McCormick (2022), while the inverse problem of embedding an observed network into a latent space is discussed, for example, by Raftery et al. (2012), Shalizi and Asta (2024), and Rastelli et al. (2024). These works are primarily statistical in nature. By contrast, the present paper focuses on structural and asymptotic properties of graph processes generated from intrinsic Riemannian volume and distance.

The present paper develops a point-process framework for random graphs on general latent spaces and studies the form of matrix exchangeability induced by bounded restrictions. It then specialises this framework to homogeneous Riemannian manifolds and derives the corresponding limiting degree and clustering statistics, including their dependence on the underlying geometry.

1.3 Outline

The paper is organised as follows. Section 2 introduces the graph-process framework: random adjacency measures directed by locally finite point processes and connected by measurable kernels. We discuss i.i.d. representations and compare graph processes with the graphon and graphex models, showing in particular that a Poisson graph process is jointly exchangeable as a random measure only when its connection kernel is essentially constant. Section 3 introduces *bounded exchangeability*: we show that every bounded restriction of a boundedly exchangeable graph process determines an infinite jointly exchangeable adjacency matrix, unique in distribution, and hence admits a local Aldous–Hoover–Kallenberg representation by a graphon; we also describe how these local graphons transform as the observation window changes.

Section 4 specialises the construction to Riemannian manifolds, using intrinsic distance and volume to define a geometrically natural class of graph processes. For this class, we prove a scale-invariance theorem describing how simultaneous rescaling of the metric, vertex density, and connection rule leaves the induced graph distribution unchanged. The principal example is the threshold model, a manifold-valued analogue of Gilbert's random plane network. Section 5 studies this model on homogeneous infinite-volume manifolds: we establish asymptotic sparsity and derive the limiting degree distribution and clustering coefficient. The asymptotic results for Euclidean and hyperbolic space show that, after matching the limiting degree distribution, the limiting clustering coefficients differ. The section concludes with simulations that illustrate the proved sparsity, degree, and clustering limits and the Euclidean–hyperbolic comparison.

2 Graph Processes

In this section we set up the graph-process framework used throughout the paper. The vertices are identified with the atoms of a locally finite point process on a Polish latent space S, and, conditionally on these latent locations, edges between distinct vertices are drawn independently according to a symmetric connection kernel $W : S^2 \to [0, 1]$. In this sense, graph processes extend the conditional-independence structure of the graphon model beyond latent probability spaces, allowing in particular locally finite point processes with infinitely many atoms. We then compare this framework with the graphon and graphex constructions. The graphon model is recovered on probability spaces under i.i.d. sampling, whereas the graphex comparison shows that, under the standing assumptions of Section 2.3, a Poisson graph process is jointly exchangeable in the random-measure sense only when its connection kernel is essentially constant.

2.1 Foundations

Broadly speaking, a graph consists of a set of points together with a set of connections between pairs of these points. The points are called vertices or nodes, and the connections are called edges. We assume throughout that edges are undirected and unweighted, and no two vertices are joined by more than one edge. Formally, a graph is a pair $G = (v(G), e(G))$, where $v(G)$ is its (countable) vertex set and

$$e(G) \subseteq \big\{\{x, y\} \mid x, y \in v(G)\big\}$$

is its edge set. Thus, edges are regarded as unordered pairs of vertices, and an edge of the form $\{x\}$ represents a loop at x. Unless otherwise stated we exclude loops, in which case the graph is called simple. When the vertices are enumerated by integers, a graph may be represented by an adjacency matrix — a symmetric $\{0,1\}$-valued matrix $A = (A_{ij})_{i,j\geq 1}$ with $A_{ij} = 1$ if and only if vertices i and j are adjacent; in particular, $A_{ii} = 0$ for simple graphs. This representation depends on the chosen enumeration.

For a general measure space $(S, \mathcal{S}, \mu)$ and $B \in \mathcal{S}$, we let $\mu|_B$ denote the restriction of the measure μ to B, that is, $\mu|_B(C) := \mu(B \cap C)$ for all $C \in \mathcal{S}$. Following Kallenberg (2017), we work with a fixed complete separable metric space (S, d) and write $\mathcal{B}(S)$ for its Borel σ-algebra. We denote by $\widehat{\mathcal{B}}(S)$ the ring of Borel sets that are bounded with respect to d. A measure μ on S is locally finite if $\mu(B) < \infty$ for every $B \in \widehat{\mathcal{B}}(S)$; we write $\mathcal{M}(S)$ for the space of locally finite measures on S, a Polish space under the corresponding vague topology. We call $\mu \in \mathcal{M}(S)$ a simple measure if it is a countable sum of distinct unit point masses.

All random elements in this paper are defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$, assumed rich enough to support the independent $\mathrm{U}[0,1]$ random variables used below. When needed, we use Kallenberg (2021, Corollary 8.18) to realise distributionally specified random elements on this probability space so that the corresponding representation holds almost surely. Unless stated otherwise, $U_\emptyset$, (U_i) and (U_{ij}) denote independent arrays of i.i.d. $\mathrm{U}[0,1]$ random variables with $U_{ij} = U_{ji}$, independent of all other randomness. We write $X \perp\!\!\!\perp Y$ when the random elements X and Y are independent. Finally, we define $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$, $\overline{\mathbb{N}} := \mathbb{N} \cup \{\infty\}$, $\overline{\mathbb{N}}_0 := \mathbb{N}_0 \cup \{\infty\}$, and $\mathbb{R}_+ := [0, \infty)$. The terminology below closely follows Veitch and Roy (2015).

Definition 1. *An adjacency measure on S is a simple measure ξ on $S^2 \times \{0,1\}$ such that*

(i) there is a simple measure ν on S with $\xi(C \times \{0,1\}) = \nu^{\otimes 2}(C)$ for every $C \in \mathcal{B}(S)^{\otimes 2}$,

(ii) $(\theta \otimes \mathrm{id})\,\xi = \xi$, where $\theta(x, y) = (y, x)$ and id *denotes the identity function on $\{0,1\}$.*

For $B \in \mathcal{B}(S)$, the B-restriction of ξ is the adjacency measure $\xi|_{B\times B\times\{0,1\}}$, which we denote by $\xi|_B$ for brevity.

Here, local finiteness on $S^2 \times \{0,1\}$ is understood with respect to any p-product metric of the fixed metric on S and the discrete metric on $\{0,1\}$, the choice being immaterial.

Let $\mathcal{A}(S)$ denote the class of adjacency measures on S. The set of atoms of $\xi \in \mathcal{A}(S)$ is

$$\mathrm{At}(\xi) := \big\{(x, y, a) \in S^2 \times \{0,1\} : \xi(\{(x, y, a)\}) = 1\big\}.$$

Projecting onto the first coordinate, we obtain $V(\xi) := \pi_1(\mathrm{At}(\xi))$, the *vertex set* of ξ. Observe that ν in Definition 1 is uniquely determined and equal to

$$\nu_\xi := \sum_{x \in V(\xi)} \delta_x.$$

We then say that ξ is *directed* by ν_ξ; we simply write ν when ξ can be inferred from the context. By Definition 1, for each $x, y \in V(\xi)$ there is exactly one $a_\xi(x, y) \in \{0,1\}$ such that

$(x, y, a_\xi(x, y)) \in \mathrm{At}(\xi)$. An enumeration $V(\xi) = \{s_i \mid i \leq n\}$ with $n \in \overline{\mathbb{N}}_0$, together with $a_{ij} \coloneqq a_\xi(s_i, s_j)$, yields

$$\xi = \sum_{i,j \leq n} \delta_{(s_i, s_j, a_{ij})}.$$

We call the $\{0,1\}$-valued symmetric array $(a_{ij})_{i,j \leq n}$ the *adjacency matrix* of the representation; it depends on the chosen enumeration of $V(\xi)$.

We now formalise a correspondence between adjacency measures and graphs.

Definition 2. *Let G be an undirected, unweighted graph without multiple edges, possibly with loops, whose vertex set $v(G) \subseteq S$ is locally finite. The adjacency measure ξ_G associated with G is given by*

$$\xi_G \coloneqq \sum_{x,y \in v(G)} \delta_{(x, y, \mathbf{1}\{\{x,y\} \in e(G)\})}.$$

It is readily confirmed that this is indeed an adjacency measure. Contrary to the convention of Veitch and Roy (2015), our definition retains vertices of degree zero. Consequently, two graphs G and G' have the same adjacency measure if and only if they have the same vertex and edge sets, hence are equal. The associated graph of an adjacency measure, recovering G from its measure, is made precise in Definition 3.

Definition 3. *Let ξ be an adjacency measure. The graph G_ξ associated with ξ has vertex set $v(G_\xi) \coloneqq V(\xi)$ and edge set*

$$e(G_\xi) \coloneqq \big\{\{x, y\} \mid x, y \in V(\xi) \text{ and } a_\xi(x, y) = 1\big\}.$$

We now introduce randomness.

Definition 4. *A random adjacency measure on S is a random measure ξ on $S^2 \times \{0,1\}$ such that $\mathbb{P}(\xi \in \mathcal{A}(S)) = 1$.*

By Kallenberg (2017, Lemma 1.6), there exists a random variable $N \in \overline{\mathbb{N}}_0$ and a.s. distinct random elements $(S_i)_{i \leq N}$ of S such that the point process ν_ξ on S satisfies

$$\nu_\xi \overset{a.s.}{=} \sum_{i \leq N} \delta_{S_i}.$$

Setting $A_{ij} \coloneqq a_\xi(S_i, S_j)$, we obtain

$$\xi \overset{a.s.}{=} \sum_{i,j \leq N} \delta_{(S_i, S_j, A_{ij})}. \tag{1}$$

Conversely, any locally finite sum (1), involving a random variable $N \in \overline{\mathbb{N}}_0$, a.s. distinct random elements $(S_i)_{i \leq N}$ of S, and a $\{0,1\}$-valued symmetric random array $A \coloneqq (A_{ij})_{i,j \leq N}$, is a random adjacency measure.

Our proposed random graph model rests on the correspondence between graphs and adjacency measures: its vertices are the atoms of a point process, and its edges are drawn conditionally independently according to a connection kernel, made precise in Definition 6 below.

Definition 5. *A connection kernel on S is a symmetric measurable function $W : S^2 \to [0,1]$.*

A connection kernel on $[0,1]$ is called a *graphon* in the graph-limit literature, see for example Borgs et al. (2008). To reflect the interplay of graphs and point processes, we call our random graph model a *graph process*.

Definition 6. *A random adjacency measure ξ on S is a graph process on S if it admits a representation* (1) *in which, for some connection kernel W on S,*

$$A_{ii} = 0, \qquad A_{ij} \mid \big(N, (S_k)\big) \overset{\text{ind.}}{\sim} \text{Bernoulli}\big(W(S_i, S_j)\big), \qquad i < j,$$

the remaining entries being given by symmetry, $A_{ji} = A_{ij}$. We call (1) *an i.i.d. representation of ξ if, in addition, the (S_k) are i.i.d. according to some probability measure $\hat{\mu}$ on S and $N \perp\!\!\!\perp (S_k)$.*

We then say that ξ is *connected* by W. Unlike the directing random measure ν, the connecting kernel W is not uniquely determined by ξ: several kernels may connect the same graph process. This non-uniqueness is revisited in Section 3, where the local graphons associated with bounded restrictions are shown to be unique only up to measure-preserving transformations. Note also that every graph process has a simple associated graph, irrespective of the values of W on the diagonal of S^2.

A point process ν on S with $\nu \overset{d}{=} \sum_{i \le N} \delta_{S_i}$, for a random variable $N \in \mathbb{N}_0$ and an array $(S_i) \perp\!\!\!\perp N$ of i.i.d. random elements of S with common distribution $\hat{\mu}$, is a *mixed binomial process* based on N and $\hat{\mu}$; note that $N < \infty$ a.s., so such a process is a.s. finite. It is a *Poisson process* based on $\mu \in \mathcal{M}(S)$, the *intensity measure* of ν, if $\nu(B_1), \dots, \nu(B_n)$ are independent and Poisson distributed with means $\mu(B_1), \dots, \mu(B_n)$ for all disjoint $B_1, \dots, B_n \in \widehat{\mathcal{B}}(S)$; when $\mu(S) = \infty$ this process has infinitely many atoms, though it remains locally finite. It is a *mixed Poisson process* based on a random variable $\delta \ge 0$ and μ if, conditionally on δ, it is Poisson based on $\delta\mu$. A graph process directed by a (mixed) Poisson or mixed binomial process is called a *(mixed) Poisson* or *mixed binomial graph process*, respectively. Finally, a mixed Poisson process based on μ, respectively a mixed binomial process based on $\hat{\mu}$, is simple if and only if μ, respectively $\hat{\mu}$, is diffuse (Kallenberg, 2017, Lemma 3.6).

All measures in this paper are locally finite. When the latent space has infinite measure, we therefore work with bounded restrictions, on which the directing process is a.s. finite. Such restrictions remain graph processes; in fact this holds for every Borel set.

Proposition 7. *Let ξ be a graph process on S and $B \in \mathcal{B}(S)$. If ξ is directed by ν and connected by W, then $\xi|_B$ is again a graph process, directed by $\nu|_B$ and connected by W.*

Proof. Let ξ be represented by (1). Restriction to B retains exactly the atoms indexed by $I_B := \{\, i \le N : S_i \in B \,\}$, so $\xi|_B$ is a random adjacency measure directed by $\nu|_B = \sum_{i \in I_B} \delta_{S_i}$. Since I_B is a deterministic function of N and (S_i), conditioning on N and (S_i) leaves the retained off-diagonal marks $(A_{ij})_{i,j \in I_B, i<j}$ independent with $A_{ij} \sim \text{Bernoulli}(W(S_i, S_j))$, while $A_{ii} = 0$; hence $\xi|_B$ is connected by W. □

In Section 3 we introduce a local notion of exchangeability, resting on graph processes whose vertices are locally i.i.d. samples from a probability measure. It is motivated by the following two lemmas.

Lemma 8 (Kallenberg, 2017, Theorem 3.4)**.** *Let ν be a mixed Poisson or binomial process on S based on μ, and fix any $B \in \widehat{\mathcal{B}}(S)$. Then $\nu|_B$ is a mixed binomial process based on $\mu|_B$.*

In a slight abuse of notation, the base $\mu|_B$ is understood to be normalised to a probability measure.

Lemma 9 (Kallenberg, 2017, Theorem 3.7)**.** *Let ν be a point process on S, and fix any $B_1, B_2, \dots \in \widehat{\mathcal{B}}(S)$ with $B_n \uparrow S$. Then ν is a mixed Poisson or binomial process if and only if $\nu|_{B_n}$ is a mixed binomial process for every $n \in \mathbb{N}$.*

These let us establish the local existence of i.i.d. representations for mixed Poisson and mixed binomial graph processes.

Proposition 10. *The B-restriction of a graph process ξ on S admits an i.i.d. representation for every $B \in \widehat{\mathcal{B}}(S)$ if and only if ξ is directed by a mixed Poisson or mixed binomial point process.*

Proof. Let ξ be directed by ν; by Proposition 7, $\xi|_B$ is directed by $\nu|_B$. Suppose $\xi|_B$ admits an i.i.d. representation $\xi|_B \overset{a.s.}{=} \sum_{i,j\le N} \delta_{(S_i,S_j,A_{ij})}$. Since B is bounded, $N = \nu(B) < \infty$ a.s., and by Definition 6 the locations (S_i) are i.i.d. with $N \perp\!\!\!\perp (S_i)$; hence $\nu|_B \overset{a.s.}{=} \sum_{i\le N} \delta_{S_i}$ is a mixed binomial process. Fix bounded sets $B_n \uparrow S$; the same argument applied to each B_n shows that $\nu|_{B_n}$ is mixed binomial for every n, so by Lemma 9 ν is a mixed Poisson or mixed binomial process.

Conversely, suppose ν is mixed Poisson or mixed binomial. By Proposition 7 and Lemma 8, $\xi|_B$ is directed by the mixed binomial process $\nu|_B \overset{a.s.}{=} \sum_{i\le N} \delta_{S_i}$, with $(S_i) \perp\!\!\!\perp N$ i.i.d. from some probability law on B. Since ξ is a graph process it is connected by some kernel W, and by Proposition 7 so is $\xi|_B$. Define

$$\widetilde{A}_{ii} \coloneqq 0, \qquad \widetilde{A}_{ij} \coloneqq \mathbf{1}\{U_{ij} \le W(S_i, S_j)\}, \qquad i < j,$$

the remaining entries being given by symmetry, $\widetilde{A}_{ji} = \widetilde{A}_{ij}$, and set $\widetilde{\xi} \coloneqq \sum_{i,j\le N} \delta_{(S_i,S_j,\widetilde{A}_{ij})}$. Then $\widetilde{\xi}$ is a graph process with an i.i.d. representation, and $\widetilde{\xi} \overset{d}{=} \xi|_B$ by Definition 6. By the remark at the beginning of the section, this distributional identity may be realised almost surely, giving an i.i.d. representation of $\xi|_B$. □

A local i.i.d. representation does not entail a global one. Indeed, let ξ be a Poisson graph process on S directed by ν with intensity measure μ satisfying $\mu(S) = \infty$. A global i.i.d. representation would make ν a mixed binomial process, hence a.s. finite; but a Poisson process with intensity measure μ satisfying $\mu(S) = \infty$ has infinitely many atoms almost surely, a contradiction. A unit-rate Poisson graph process on $\mathbb{R}$ is one such example.

2.2 Comparison to graphon models

A random matrix $(X_{ij})_{i,j\ge1}$ is *jointly exchangeable* if, for all $n \ge 1$ and all permutations π of $\{1, \dots, n\}$,

$$(X_{ij})_{i,j\le n} \overset{d}{=} (X_{\pi(i)\pi(j)})_{i,j\le n}. \tag{2}$$

This symmetry underpins an array analogue of de Finetti's theorem, the Aldous–Hoover–Kallenberg theorem, due to Hoover (1979), Aldous (1981), and Kallenberg (1989), which provides a functional representation of such arrays.

A random graph is *(jointly) exchangeable* if it admits a jointly exchangeable adjacency matrix. Recall that the adjacency matrix of a random graph is a symmetric $\{0,1\}$-valued random matrix with a vanishing diagonal. The Aldous–Hoover–Kallenberg theorem then asserts that an infinite jointly exchangeable adjacency matrix $A = (A_{ij})_{i,j\in\mathbb{N}}$ admits a representation

$$A_{ij} \overset{a.s.}{=} \mathbf{1}\{U_{ij} \le w(U_\emptyset, U_i, U_j)\}, \qquad i, j \in \mathbb{N}, \tag{3}$$

where $w : [0,1]^3 \to [0,1]$ is measurable and symmetric in the second and third argument. Because A has a vanishing diagonal, the representing function may, and henceforth will, be chosen so that

$$w(u_0, u, u) = 0 \qquad \text{for all } u_0, u \in [0,1].$$

Conversely, every array with a representation (3) involving such a zero-diagonal representative is a jointly exchangeable adjacency matrix. The array is *ergodic* when w does not depend on $U_\emptyset$, in which case (3) reduces to

$$A_{ij} \overset{a.s.}{=} \mathbf{1}\{U_{ij} \le W(U_i, U_j)\}, \qquad i, j \in \mathbb{N} \tag{4}$$

for a graphon $W : [0,1]^2 \to [0,1]$, chosen so that $W(u,u) = 0$ for every $u \in [0,1]$. The general case is a mixture of such ergodic laws.

For the zero-diagonal representative W fixed above, the ergodic case (4) admits the following sampling description. Sample i.i.d. random variables $U_1, U_2, \cdots \sim \mathrm{U}[0,1]$. Set $A_{ii} = 0$ and draw

$$A_{ij} \mid (U_k) \overset{\text{ind.}}{\sim} \text{Bernoulli}\big(W(U_i, U_j)\big), \qquad i < j,$$

the remaining entries being given by symmetry, $A_{ji} = A_{ij}$. The sampled A then has the law of the array in (4). We call the associated random graph with vertex set $\mathbb{N}$ the *graphon model* based on W. We obtain the *finite* graphon model on $n \in \mathbb{N}$ vertices based on W by restricting A to its leading principal $n \times n$ submatrix.

The graphon W in (4) is not uniquely determined by the law of A: for example, replacing W by $W \circ (T \otimes T)$ for any measure-preserving transformation T of $[0,1]$ leaves the law of A unchanged. This non-uniqueness is revisited in Section 3.

The following example frames the graphon model as a graph process. Note that only the finite graphon model defines a graph process, for otherwise the directing measure would a.s. have infinitely many atoms in $[0,1]$, violating local finiteness.

Example 11. *Let $W : [0,1]^2 \to [0,1]$ be a graphon such that $W(u,u) = 0$ for every $u \in [0,1]$. Sample i.i.d. random variables $U_1, U_2, \cdots \sim \mathrm{U}[0,1]$. Set $A_{ii} = 0$ and draw*

$$A_{ij} \mid (U_k) \overset{\text{ind.}}{\sim} \text{Bernoulli}\big(W(U_i, U_j)\big), \qquad i < j,$$

the remaining entries being given by symmetry, $A_{ji} = A_{ij}$. For $n \in \mathbb{N}$, define a graph process on $[0,1]$ by

$$\xi \coloneqq \sum_{i,j \le n} \delta_{(U_i, U_j, A_{ij})}.$$

Indeed, this is an i.i.d. representation by Definition 6. Its adjacency matrix $(A_{ij})_{i,j \le n}$ is jointly exchangeable and realises the finite graphon model on n vertices based on W.

By Orbanz and Roy (2015), an infinite exchangeable graph is almost surely either empty or dense, where *dense* means that the number of edges e_n among the first n vertices satisfies $e_n \geq c\, n^2$ for some constant $c > 0$ independent of n. By contrast, we later show that graph processes can produce *sparse* graphs, with $e_n = O(n)$.

2.3 Comparison to graphex models

The *graphex model*, introduced by Caron and Fox (2017) and Veitch and Roy (2015) building on the theory of exchangeable random measures on $\mathbb{R}_+^2$ developed by Kallenberg (1990), generalises the graphon framework to produce sparse infinite graphs. Following Veitch and Roy (2015), the distribution of such a graph is characterised by a *graphex*, a triple (I, S, W) consisting of a nonnegative real $I \in \mathbb{R}_+$, an integrable function $S : \mathbb{R}_+ \to \mathbb{R}_+$, and a symmetric measurable function $W : \mathbb{R}_+^2 \to [0,1]$ satisfying weak integrability conditions; these govern, respectively, the isolated edges, the star components, and the principal connection structure of the graph. As Veitch and Roy (2015) do, and as later emphasised by Janson (2022), we restrict to the case $I = S = 0$, in which the graph is determined by W alone; we then refer to W itself as the graphex. In this construction, a unit-rate Poisson process $\Pi = \{(\theta_i, \vartheta_i)\}_{i \geq 1}$ on $\mathbb{R}_+^2$ provides the vertices, where $\theta_i \in \mathbb{R}_+$ is a time stamp and $\vartheta_i \in \mathbb{R}_+$ a latent connectivity parameter. Conditionally on Π, each pair of distinct vertices is connected independently, i.e.

$$Z_{ii} = 0, \qquad Z_{ij} \mid \Pi \overset{\text{ind.}}{\sim} \text{Bernoulli}\big(W(\vartheta_i, \vartheta_j)\big), \qquad i < j, \tag{5}$$

the remaining entries being given by symmetry, $Z_{ji} = Z_{ij}$; the edge $\{\theta_i, \theta_j\}$ being present if and only if $Z_{ij} = 1$. The edge random measure $\sum_{Z_{ij}=1} \delta_{(\theta_i, \theta_j)}$ on $\mathbb{R}_+^2$ is jointly exchangeable

in the sense of Kallenberg (1990), and is locally finite under the integrability conditions on W established by Veitch and Roy (2015). We emphasise that the adjacency matrix $Z = (Z_{ij})_{i,j\in\mathbb{N}}$ need not be jointly exchangeable as an infinite array.

Throughout this subsection we show that graph processes are in general distinct from graphex models, additionally assuming that S is uncountable and endowed with a σ-finite diffuse Borel measure μ such that $\mu(S) > 0$. For each $d \in \mathbb{N}$, we write λ_d for Lebesgue measure on $\mathbb{R}^d$, and abbreviate $\lambda \coloneqq \lambda_1$. In particular, $\lambda_2 = \lambda^{\otimes 2}$.

Definition 12. *A measurable map $T : S \to S$ is a μ-preserving transformation of S if its pushforward $T\mu \coloneqq \mu \circ T^{-1}$ equals μ, that is, $\mu(T^{-1}(B)) = \mu(B)$ for all $B \in \mathcal{B}(S)$.*

Definition 13. *A random measure η on S^2 is jointly μ-exchangeable if $(T \otimes T)\,\eta \overset{d}{=} \eta$ for every μ-preserving transformation T of S, where $T \otimes T : S^2 \to S^2$, $(x, y) \mapsto (Tx, Ty)$, and $(T \otimes T)\eta$ denotes the pushforward of η.*

Recall that a measurable (with respect to the Borel σ-algebra) bijection from S to some other topological space $\tilde{S}$ with measurable inverse is called a *Borel isomorphism*. The next result is known, but we could not locate a precise statement in the literature.

Proposition 14. *Let $I \coloneqq [0, \mu(S))$, with the convention $I = \mathbb{R}_+$ when $\mu(S) = \infty$. Then there is a Borel isomorphism $\phi : S \to I$ with $\phi\mu = \lambda|_I$, and for any such ϕ, a random measure η on S^2 is jointly μ-exchangeable if and only if $(\phi \otimes \phi)\eta$ on I^2 is jointly $\lambda|_I$-exchangeable.*

We use the following lemma of Kechris (1995, Theorem 17.41).

Lemma 15. *If $\mu(S) = 1$, there exists a Borel isomorphism $\phi : S \to [0, 1)$ such that $\phi\mu = \lambda|_{[0,1)}$.*

Now we can prove Proposition 14.

Proof of Proposition 14. If $\mu(S) < \infty$, normalising, applying Lemma 15, and rescaling gives the required Borel isomorphism, so assume $\mu(S) = \infty$. Set $B_0 \coloneqq \emptyset$ and $\beta_0 \coloneqq 0$. Choose $B_n \uparrow S$ in $\mathcal{B}(S)$ with $\beta_n \coloneqq \mu(B_n)$ finite and strictly increasing, and set $D_n \coloneqq B_n \setminus B_{n-1}$. Each increment D_n has mass $\beta_n - \beta_{n-1} > 0$; normalising, applying Lemma 15, and rescaling yields a Borel isomorphism $\phi_n : D_n \to [\beta_{n-1}, \beta_n)$ with $\phi_n(\mu|_{D_n}) = \lambda|_{[\beta_{n-1},\beta_n)}$. The increments D_n partition S and their images partition $I \coloneqq \mathbb{R}_+$, so the ϕ_n glue to a single Borel isomorphism $\phi : S \to I$ such that $\phi|_{D_n} = \phi_n$, with $\phi\mu = \lambda|_I$.

It remains to show that η is jointly μ-exchangeable if and only if $(\phi \otimes \phi)\eta$ is jointly $\lambda|_I$-exchangeable. Suppose the former, and let $T : I \to I$ be $\lambda|_I$-preserving. Then $\phi^{-1} \circ T \circ \phi$ is μ-preserving, since

$$(\phi^{-1} \circ T \circ \phi)\mu = \phi^{-1}\big(T(\phi\mu)\big) = \phi^{-1}\big(T\lambda|_I\big) = \phi^{-1}\lambda|_I = \mu.$$

By hypothesis, $\big((\phi^{-1} \circ T \circ \phi) \otimes (\phi^{-1} \circ T \circ \phi)\big)\eta \overset{d}{=} \eta$; pushing forward by $\phi \otimes \phi$ and using

$$(\phi \otimes \phi) \circ \big((\phi^{-1} \circ T \circ \phi) \otimes (\phi^{-1} \circ T \circ \phi)\big) = (T \otimes T) \circ (\phi \otimes \phi),$$

we obtain $(T \otimes T)(\phi \otimes \phi)\eta \overset{d}{=} (\phi \otimes \phi)\eta$. The converse is analogous, exchanging the roles of ϕ and ϕ^{-1}. □

In the graphex literature, a random graph is identified with a random measure on $\mathbb{R}_+^2$ whose coordinates are vertex labels and whose point masses mark the edges, following Caron and Fox (2017), Veitch and Roy (2015), and Borgs et al. (2018). Under the graphex construction this edge measure is jointly λ-exchangeable in the sense of Definition 13 with $S = \mathbb{R}_+$ and $\mu = \lambda$, as established by Kallenberg (1990) and Veitch and Roy (2015). In contrast to Definition 2, this

encoding records only present edges, so isolated (degree-zero) vertices leave no trace. Adopting the same convention on a general latent space, we define the *edge measure* of a random adjacency measure ξ on S by

$$E_\xi(C) \coloneqq \xi(C \times \{1\}), \quad C \in \mathcal{B}(S)^{\otimes 2}.$$

Under a measurable enumeration (1), this random measure on S^2 has the representation

$$E_\xi \overset{a.s.}{=} \sum_{i,j \le N} A_{ij}\, \delta_{(S_i, S_j)}. \tag{6}$$

Abusing notation, we write ξ for E_ξ when the intended interpretation is clear from the context.

For a graph process in particular, the edge measure is not in general jointly μ-exchangeable, unlike the graphex edge measure, as Theorem 16 makes precise.

Theorem 16. *Let S be uncountable, and let ξ be the edge measure* (6) *of a Poisson graph process on S with diffuse σ-finite intensity measure μ, connected by a kernel W. Then ξ is jointly μ-exchangeable if and only if W is $\mu^{\otimes 2}$-a.e. constant.*

To prove this, the following lemma is useful.

Lemma 17. *Let $W : [0,1]^2 \to [0,1]$ be measurable, and suppose that for every $\lambda|_{[0,1]}$-preserving bijection T of $[0,1]$ the identity $W \circ (T \otimes T) = W$ holds $\lambda^{\otimes 2}$-a.e. Then W is $\lambda^{\otimes 2}$-a.e. constant.*

We defer the short proof of this Lemma to the appendix.

Proof of Theorem 16. If W is $\mu^{\otimes 2}$-a.e. constant the claim is immediate. Conversely, suppose ξ is jointly μ-exchangeable. By Proposition 14 we may transport ξ along a measure-preserving Borel isomorphism ϕ and assume that ξ is a jointly $\lambda|_I$-exchangeable Poisson graph process on $I = [0, \mu(S))$ with intensity measure $\lambda|_I$ connected by $W : I^2 \to [0,1]$. It then suffices to show that W is $\lambda^{\otimes 2}$-a.e. constant on I^2.

Fix a bounded subinterval $J \subseteq I$ and a $\lambda|_J$-preserving bijection T of J, extended by the identity to a $\lambda|_I$-preserving bijection of I. For bounded measurable f, g supported in J, the multivariate Mecke equation (Last and Penrose, 2017, Theorem 4.4) gives

$$\mathbb{E}\left[\int f(x)g(y)\, \xi(dx, dy)\right] = \int_{J^2} f(x)g(y)\, W(x,y)\, dx\, dy,$$

using that the diagonal adjacency marks of ξ are zero by Definition 6; the right-hand side is finite as $\lambda(J) < \infty$ and $W \le 1$. By joint exchangeability $(T \otimes T)\xi \overset{d}{=} \xi$, so

$$\mathbb{E}\left[\int f(x)g(y)\, \xi(dx, dy)\right] = \mathbb{E}\left[\int f(x)g(y)\, (T \otimes T)\xi(dx, dy)\right] = \mathbb{E}\left[\int f(Tx)g(Ty)\, \xi(dx, dy)\right].$$

As $f \circ T$ and $g \circ T$ are again supported in J, the Mecke identity applies to the right-hand side and yields

$$\int_{J^2} f(x)g(y)\, W(x,y)\, dx\, dy = \int_{J^2} f(Tx)g(Ty)\, W(x,y)\, dx\, dy.$$

Composing the right-hand side with $T^{-1} \otimes T^{-1}$, which also preserves $\lambda^{\otimes 2}$, gives

$$\int_{J^2} f(x)g(y)\, W(x,y)\, dx\, dy = \int_{J^2} f(x)g(y)\, W\big(T^{-1}x, T^{-1}y\big)\, dx\, dy$$

for all such f, g, and so

$$W = W \circ \big(T^{-1} \otimes T^{-1}\big) \qquad \lambda^{\otimes 2}\text{-a.e. on } J^2.$$

Rescaling J to $[0,1]$ and applying Lemma 17, the restriction of W to J^2 equals a constant c_J $\lambda^{\otimes 2}$-a.e. Choosing $J_n \uparrow I$ with $\lambda(J_n) > 0$, the overlaps $J_n^2 \subseteq J_{n+1}^2$ have positive measure, so all c_{J_n} coincide in a common value c, giving $W = c$ $\lambda^{\otimes 2}$-a.e. on I^2. □

Note that ξ is jointly μ-exchangeable if and only if it is jointly $\delta\mu$-exchangeable for every $\delta > 0$, since a map preserves μ exactly when it preserves $\delta\mu$. The following example illustrates the essential idea behind the proof of Theorem 16, namely that joint exchangeability fails when the invariance under measure-preserving transformations does not pass to the connection kernel W.

Example 18. *Let ξ be the edge measure (6) of a graph process on $\mathbb{R}$ directed by a unit-rate Poisson point process and connected by $W : (x, y) \mapsto \mathbf{1}\{\|x - y\| \leq 1\}$. Then ξ is not jointly λ-exchangeable as a random measure on $\mathbb{R}^2$. To see this, let T be the λ-preserving transformation that swaps the intervals $(1, 2]$ and $(4, 5]$. Then $\mathbb{E}[\xi((1, 2] \times (2, 3])] > 0$, whereas $(T \otimes T)\xi\big((1, 2] \times (2, 3]\big) = \xi((4, 5] \times (2, 3]) \overset{a.s.}{=} 0$, since every $x \in (4, 5]$ and $y \in (2, 3]$ satisfy $\|x - y\| > 1$.*

3 Exchangeability

Recall that the class of infinite random graphs with jointly exchangeable adjacency matrices is completely characterised by mixtures of graphon models. In this section we introduce a new notion of exchangeability for graph processes, which we call *bounded exchangeability*. For bounded $B \subseteq S$, we show that under a mild support condition on the directing process, made precise in Definition 20, the restriction $\xi|_B$ to vertices in B determines an infinite jointly exchangeable matrix A^B, unique in distribution. This holds even when the adjacency matrix of the global graph is not jointly exchangeable. As a consequence, boundedly exchangeable graph processes can locally be represented as graphon models. This involves a local graphon $W_B : [0, 1]^2 \to [0, 1]$, which, up to measure-preserving transformations, is the composition of the zero-diagonal modification

$$\widetilde{W}(x, y) \coloneqq \mathbf{1}\{x \neq y\}W(x, y), \qquad \text{for all } x, y \in S$$

of the connection kernel $W : S^2 \to [0, 1]$ with a local sampling function $f_B : [0, 1] \to B$. Finally, we establish relations between the local graphons as the observation window changes.

3.1 Boundedly exchangeable graph processes

For a real-valued matrix $A = (A_{ij})_{i,j\geq 1}$ and a permutation π of $\mathbb{N}$, define the permuted matrix $A \circ \pi \coloneqq (A_{\pi(i)\pi(j)})_{i,j\geq 1}$. For $n \in \mathbb{N}$, let $A[n] \coloneqq (A_{ij})_{i,j\leq n}$ denote the leading principal $n \times n$ submatrix of A. We call a permutation π of $\mathbb{N}$ *finitely supported* if $\pi(i) = i$ for all but finitely many $i \in \mathbb{N}$.

First, we prove a lemma concerning random matrices of random finite dimension, establishing necessary and sufficient conditions under which the dimension-conditional distributions arise as the laws of the leading principal submatrices of an infinite jointly exchangeable matrix, which is then unique in distribution. Consider the space of finite-dimensional real-valued square matrices $\mathcal{R} \coloneqq \bigsqcup_{n=0}^{\infty} \mathbb{R}^{n\times n}$. A random element $\mathbf{A} = (A, N)$ in $\mathcal{R}$ consists of a random variable $N \in \mathbb{N}_0$ and a random matrix $A \in \mathbb{R}^{N\times N}$.

Lemma 19. *Let $\mathbf{A} = (A, N)$ be a random element in $\mathcal{R}$ such that $\mathbb{P}(N = n) > 0$ for every $n \in \mathbb{N}$. Then $\mathbf{A}$ satisfies*

(i) (Finite exchangeability) For all $n \geq 1$ and all permutations π of $\{1, \ldots, n\}$:

$$A \mid (N = n) \overset{d}{=} A \circ \pi \mid (N = n)$$

(ii) (Compatibility) For all $n \geq 1$ and all $m \geq n$:

$$A \mid (N = n) \overset{d}{=} A[n] \mid (N = m)$$

if and only if there exist an infinite jointly exchangeable matrix $A^{(\infty)} \in \mathbb{R}^{\mathbb{N}\times\mathbb{N}}$ such that for all $n \geq 1$:

$$A \mid (N = n) \stackrel{d}{=} A^{(\infty)}[n] \tag{7}$$

In that case, $A^{(\infty)}$ is unique in distribution.

Proof. The converse is trivial, so we only prove the forward direction. By the compatibility of conditional laws and full support on $\mathbb{N}$, Kolmogorov's extension theorem applies. Hence, there exists a random matrix $A^{(\infty)} \in \mathbb{R}^{\mathbb{N}\times\mathbb{N}}$, which is unique in distribution, such that (7) holds for all $n \geq 1$. Let π be a finitely supported permutation of $\mathbb{N}$ and choose n_0 with $\pi(i) = i$ for all $i > n_0$. For every $n \geq n_0$ the permutation π restricts to a permutation of $\{1, \ldots, n\}$, so

$$\left(A^{(\infty)} \circ \pi\right)[n] = \left(A^{(\infty)}[n]\right) \circ \pi \stackrel{d}{=} A^{(\infty)}[n],$$

the equality of laws holding by (7) together with condition (i). Since this holds for all $n \geq n_0$ and π is arbitrary, it follows that $A^{(\infty)}$ is jointly exchangeable. □

The idea is to apply Lemma 19 to the adjacency matrix of a graph process. For the adjacency matrix of an i.i.d. representation, the first two conditions follow from the i.i.d. structure of the locations and the conditional independence of the marks; in light of Proposition 10 we therefore make the following definition.

Definition 20. *We call a graph process ξ on S boundedly exchangeable if it is directed by a mixed Poisson or mixed binomial process ν and, for every $B \in \widehat{\mathcal{B}}(S)$ such that $\mathbb{P}(\nu(B) = 0) < 1$, we have $\mathbb{P}(\nu(B) = n) > 0$ for all $n \in \mathbb{N}$.*

Note that ξ is automatically boundedly exchangeable when ν is a mixed Poisson process, since the count $\nu(B)$ then charges every $n \in \mathbb{N}$ whenever $\mathbb{P}(\nu(B) = 0) < 1$. For a mixed binomial process this need not hold, which is why the support condition is imposed.

Theorem 21. *Let ξ be a boundedly exchangeable graph process on S directed by ν. For every $B \in \widehat{\mathcal{B}}(S)$ where $\mathbb{P}(\nu(B) = 0) < 1$ there exists an infinite jointly exchangeable matrix $A^B \in \mathbb{R}^{\mathbb{N}\times\mathbb{N}}$ such that the adjacency matrix (A, N) of any i.i.d. representation of $\xi|_B$ satisfies for all $n \geq 1$:*

$$A \mid (N = n) \stackrel{d}{=} A^B[n] \tag{8}$$

Moreover, the distribution of A^B is unique.

Proof. Let $B \in \widehat{\mathcal{B}}(S)$ with $\mathbb{P}(\nu(B) = 0) < 1$. By bounded exchangeability, $\mathbb{P}(\nu(B) = n) > 0$ for every $n \in \mathbb{N}$, so the conditional laws below are well defined. By Proposition 10, $\xi|_B$ admits an i.i.d. representation

$$\xi|_B \stackrel{a.s.}{=} \sum_{i,j\leq N} \delta_{(S_i, S_j, A_{ij})},$$

where (S_i) are i.i.d. random elements of B, independent of $N \stackrel{a.s.}{=} \nu(B)$. If $W : S^2 \to [0, 1]$ is a connection kernel for this representation, then

$$A \mid (N = n) \stackrel{d}{=} \big(\mathbf{1}\{i \neq j\}\mathbf{1}\{U_{ij} \leq W(S_i, S_j)\}\big)_{i,j\leq n}.$$

Thus Lemma 19 applies and yields an infinite jointly exchangeable matrix A^B satisfying (8). It may be realized as

$$A^B_{ij} := \mathbf{1}\{i \neq j\}\mathbf{1}\{U_{ij} \leq W(S_i, S_j)\}, \qquad i, j \in \mathbb{N}.$$

It remains to show that the law of A^B is independent of the chosen i.i.d. representation. Fix $n \in \mathbb{N}$ and work conditionally on $\{N = n\}$. Denote by $\mathcal{A}_n(B)$ the standard Borel space of adjacency measures on B with n vertices, let

$$\mathcal{Y}_n := \big\{(s_1, \dots, s_n) \in B^n \mid s_i \neq s_j \text{ for } i \neq j\big\} \times \big\{(a_{ij})_{i,j\leq n} \in \{0,1\}^{n\times n} \mid a_{ii} = 0, a_{ij} = a_{ji}\big\},$$

also standard Borel, and define the measurable map

$$\Phi_n : \mathcal{Y}_n \to \mathcal{A}_n(B), \qquad \Phi_n\big((s_i)_{i\leq n}, (a_{ij})_{i,j\leq n}\big) := \sum_{i,j\leq n} \delta_{(s_i, s_j, a_{ij})}.$$

The group $\mathfrak{S}_n$ of permutations of $\{1, \dots, n\}$ acts on $\mathcal{Y}_n$ by

$$\pi\big((s_i)_{i\leq n}, (a_{ij})_{i,j\leq n}\big) := \big((s_{\pi(i)})_{i\leq n}, (a_{\pi(i)\pi(j)})_{i,j\leq n}\big).$$

The labelled pair $L_n := \big((S_i)_{i\leq n}, (A_{ij})_{i,j\leq n}\big)$ is invariant in distribution under this action, Φ_n is invariant, and each fibre $\Phi_n^{-1}\{\zeta\}$ consists of the $n!$ enumerations of the vertex set of $\zeta \in \mathcal{A}_n(B)$. The law μ_n of L_n pushes forward under Φ_n to $\nu_n := \mathcal{L}(\xi|_B \mid N = n)$.

By Kallenberg (2021, Theorem 8.5), there exists a probability kernel K_n from $\mathcal{A}_n(B)$ to $\mathcal{Y}_n$ such that

(i) $K_n(\zeta, \Phi_n^{-1}(\zeta)) = 1$ for ν_n-a.e. $\zeta \in \mathcal{A}_n(B)$,

(ii) $\mu_n(E) = \int_{\mathcal{A}_n(B)} K_n(\zeta, E)\, \nu_n(d\zeta)$ for every $E \in \mathcal{B}(\mathcal{Y}_n)$,

and any two probability kernels satisfying (i)–(ii) agree for ν_n-a.e. ζ. Note that, for every $\pi \in \mathfrak{S}_n$,

$$K_n^\pi(\zeta, E) := K_n(\zeta, \pi^{-1}E),$$

is another probability kernel satisfying (i)–(ii); this follows from finite exchangeability of L_n and invariance of Φ_n. By uniqueness, $K_n^\pi(\zeta, \cdot) = K_n(\zeta, \cdot)$ for ν_n-a.e. ζ. Since $K_n(\zeta, \cdot)$ is a probability measure on the finite set $\Phi_n^{-1}(\zeta)$ for ν_n-a.e. ζ, and $\mathfrak{S}_n$ acts transitively on this set, we conclude that $K_n(\zeta, \cdot)$ is the uniform measure on the fibre $\Phi_n^{-1}(\zeta)$ for ν_n-a.e. ζ.

If $\big((S_i')_{i\leq n}, (A_{ij}')_{i,j\leq n}\big)$ is obtained from another i.i.d. representation, then $N' \overset{a.s.}{=} N$ and the same argument yields a probability kernel K_n' from $\mathcal{A}_n(B)$ to $\mathcal{Y}_n$ such that $K_n'(\zeta, \cdot)$ is the uniform measure on the fibre $\Phi_n^{-1}(\zeta)$ for ν_n-a.e. ζ. Integrating both with respect to ν_n, (ii) yields

$$A \mid (N = n) \overset{d}{=} A' \mid (N' = n)$$

for every $n \in \mathbb{N}$. By the uniqueness clause of Lemma 19, the law of A^B is the same for both representations. □

We call the unique distribution determined by (8) the *(jointly) exchangeable distribution associated with ξ on B*. Any random matrix having this distribution is called a *(jointly) exchangeable matrix associated with ξ on B*. In particular, every associated matrix A^B has a vanishing diagonal almost surely; equivalently,

$$\mathbb{P}\big(A_{ii}^B = 0 \text{ for all } i \in \mathbb{N}\big) = 1.$$

3.2 Aldous–Hoover–Kallenberg representations

Recall that infinite jointly exchangeable arrays are characterised by their Aldous–Hoover–Kallenberg representation, which we now discuss in relation to boundedly exchangeable graph processes; we will see that the bounded restrictions realize the ergodic case of this representation. Throughout this subsection, let ξ be a boundedly exchangeable graph process on S directed by ν and connected by a connection kernel $W : S^2 \to [0,1]$. Also, let $\widetilde{W}(x,y) := \mathbf{1}\{x \neq y\}W(x,y)$ for every $x, y \in S$ be the zero-diagonal modification of W. Thus $\widetilde{W} = W$ off the diagonal of S^2, while $\widetilde{W}(x,x) = 0$ for every $x \in S$.

We use the following transfer lemma.

Lemma 22 (Kallenberg, 2021, Lemma 4.22)**.** *Let K be a probability kernel from a measurable space X to a standard Borel space Y. Then there exists a measurable function $f : X \times [0,1] \to Y$ such that if U is* $\mathrm{U}[0,1]$*, then $f(x,U)$ has distribution $K(x,\cdot)$ for every $x \in X$.*

First, we describe a localization procedure by which the zero-diagonal modification $\widetilde{W}$ of the connection kernel, composed with a local sampling function, gives rise to local graphons, recovering the classical ergodic graphon model on each bounded region.

Proposition 23. *For any $B \in \widehat{\mathcal{B}}(S)$ with $\mathbb{P}(\nu(B) = 0) < 1$, there exists a graphon $W_B : [0,1]^2 \to [0,1]$ such that every exchangeable matrix A^B associated with ξ on B admits a representation*

$$A^B_{ij} \overset{a.s.}{=} \mathbf{1}\{U_{ij} \leq W_B(U_i, U_j)\}, \qquad i,j \in \mathbb{N}. \tag{9}$$

Proof. By the construction in Theorem 21 there is an array (S^B_i) of i.i.d. random elements of B such that the matrix defined by

$$A^B_{ij} \coloneqq \mathbf{1}\{i \neq j\}\mathbf{1}\big\{U_{ij} \leq W(S^B_i, S^B_j)\big\}, \qquad i,j \in \mathbb{N},$$

is an exchangeable matrix associated with ξ on B. Since the S^B_i are almost surely distinct, this matrix satisfies

$$A^B_{ij} \overset{a.s.}{=} \mathbf{1}\big\{U_{ij} \leq \widetilde{W}(S^B_i, S^B_j)\big\}, \qquad i,j \in \mathbb{N}.$$

As a Borel subset of the Polish space S, the set B is a standard Borel space. If η_B is the marginal law of the (S^B_i), then by Lemma 22 there exists a measurable $f_B : [0,1] \to B$ with $f_B(U) \sim \eta_B$ for $U \sim \mathrm{U}[0,1]$; hence there exists a representation $(S^B_i) \overset{a.s.}{=} f_B(U_i)$. Substituting this yields

$$A^B_{ij} \overset{a.s.}{=} \mathbf{1}\big\{U_{ij} \leq \widetilde{W}(f_B(U_i), f_B(U_j))\big\}, \qquad i,j \in \mathbb{N},$$

which is (9) with $W_B \coloneqq \widetilde{W} \circ (f_B \otimes f_B)$. Indeed, W_B is a graphon, being symmetric and measurable as W is. Since $\tilde{A}^B \overset{d}{=} A^B$ for any other exchangeable matrix $\tilde{A}^B$ associated with ξ on B, the same representation holds for $\tilde{A}^B$. □

We call W_B a *local graphon* for ξ on B. Observe that (9) is the ergodic Aldous–Hoover–Kallenberg representation of an infinite jointly exchangeable adjacency matrix; in other words, the exchangeable distribution associated with ξ on B is a graphon model based on W_B. Similarly, we call f_B a *sampling function* for ξ on B. Its defining property, that $f_B(U) \sim \eta_B$ for $U \sim \mathrm{U}[0,1]$, depends only on the marginal law η_B of the (S^B_i) and hence not on the connection kernel W; in particular a single sampling function localizes every kernel connecting ξ. The function f_B itself is generally not unique.

Like the sampling functions, local graphons are generally not unique; however, any two are related through measure-preserving transformations.

Proposition 24. *For any $B \in \widehat{\mathcal{B}}(S)$ with $\mathbb{P}(\nu(B) = 0) < 1$, let W_B and W'_B be two local graphons for ξ on B. Then there exist $\lambda|_{[0,1]}$-preserving transformations $T, T' : [0,1] \to [0,1]$ such that*

$$W_B \circ (T \otimes T) = W'_B \circ (T' \otimes T') \quad \lambda^{\otimes 2}\text{-a.e.} \tag{10}$$

Proof. By definition, the infinite jointly exchangeable matrices generated by W_B and W'_B via (9) are equal in distribution. Applying Kallenberg (2005, Theorem 7.28) to their off-diagonal laws yields (10). □

To understand how local graphons and exchangeable distributions behave under changing frames of reference, it suffices to find suitable sampling functions. The following two results outline how sampling functions for different frames of reference may be obtained from one another via conditional sampling.

Proposition 25. *Let $B, C \in \widehat{\mathcal{B}}(S)$ be disjoint such that $\mathbb{P}(\nu(B) = 0) < 1$ and $\mathbb{P}(\nu(C) = 0) < 1$, and let f_B and f_C be sampling functions for ξ on B and C, respectively. Then, for some $p \in (0, 1)$,*

$$f_{B\cup C}(x) \coloneqq \begin{cases} f_B\left(\frac{x}{p}\right), & \text{for } x \in [0, p) \\ f_C\left(\frac{x-p}{1-p}\right), & \text{for } x \in [p, 1] \end{cases}$$

is a sampling function for ξ on $B \cup C$.

Proof. Since $B \cup C \in \widehat{\mathcal{B}}(S)$, Proposition 10 gives an i.i.d. representation of $\xi|_{B\cup C}$; let (S_i) be its vertex array and η their marginal law, well defined as $\mathbb{P}(\nu(B \cup C) = 0) < 1$. Note that $\xi|_B = (\xi|_{B\cup C})|_B$, and conditioning the i.i.d. vertices on B gives vertex marginal $\eta(B)^{-1}\eta|_B =: \eta_B$ for $\xi|_B$, which is the law sampled by f_B; similarly $\eta_C \coloneqq \eta(C)^{-1}\eta|_C$ is sampled by f_C. As B, C partition the support of η, we have the mixture $\eta = \eta(B)\,\eta_B + \eta(C)\,\eta_C$. It suffices to show $f_{B\cup C}(U) \sim \eta$ for $U \sim \mathrm{U}[0, 1]$. Set $p \coloneqq \eta(B)$. Then $p \in (0, 1)$: if $p = 0$ then $\eta(B) = 0$, so $\xi|_B$ would have no vertices a.s., contradicting $\mathbb{P}(\nu(B) = 0) < 1$; the case $p = 1$ contradicts $\mathbb{P}(\nu(C) = 0) < 1$ symmetrically. The maps $x \mapsto x/p$ and $x \mapsto (x - p)/(1 - p)$ send $[0, p)$ and $[p, 1]$ uniformly onto $[0, 1]$, so $f_{B\cup C}(U) = f_B(U/p) \sim \eta_B$ on $\{U < p\}$ and $f_{B\cup C}(U) = f_C((U - p)/(1 - p)) \sim \eta_C$ on $\{U \geq p\}$. Hence $f_{B\cup C}(U) \sim p\,\eta_B + (1 - p)\eta_C = \eta$. □

Proposition 26. *Let $B, C \in \widehat{\mathcal{B}}(S)$ be disjoint such that $\mathbb{P}(\nu(B) = 0) < 1$, and let $f_{B\cup C}$ be sampling function for ξ on $B \cup C$. Then, defining $T_B : [0, 1] \to [0, 1]$ by*

$$T_B(x) \coloneqq \inf\left\{ y \in [0, 1] \;\middle|\; \frac{\lambda\left(f_{B\cup C}^{-1}(B) \cap [0, y]\right)}{\lambda\left(f_{B\cup C}^{-1}(B)\right)} \geq x \right\},$$

the composition $f_{B\cup C} \circ T_B$ is a sampling function for ξ on B.

Proof. Let η be the law of $f_{B\cup C}(U)$ for $U \sim \mathrm{U}[0, 1]$; then $f_{B\cup C}$ pushes $\lambda|_{[0,1]}$ to η, so $\lambda(f_{B\cup C}^{-1}(B)) = \eta(B) > 0$, the positivity from $\mathbb{P}(\nu(B) = 0) < 1$. A map f_B is a sampling function for ξ on B if and only if $f_B(U) \sim \eta(B)^{-1}\eta|_B$. Now T_B is the quantile function of the uniform law on $f_{B\cup C}^{-1}(B)$, so $T_B(U)$ is uniform on $f_{B\cup C}^{-1}(B)$; pushing this forward by $f_{B\cup C}$ yields η conditioned on B, namely $\eta(B)^{-1}\eta|_B$. Hence $f_{B\cup C} \circ T_B$ is a sampling function for ξ on B. □

Next, we show that a jointly exchangeable matrix $A^{B\cup C}$ associated with $\xi|_{B\cup C}$ is a composition of two jointly exchangeable matrices, A^B and A^C, associated with $\xi|_B$ and $\xi|_C$, respectively, and a separately exchangeable (see Kallenberg (2005)) matrix A^{BC} describing the interactions.

Proposition 27. *Let $B, C \in \widehat{\mathcal{B}}(S)$ be disjoint such that $\mathbb{P}(\nu(B) = 0) < 1$ and $\mathbb{P}(\nu(C) = 0) < 1$, and let $A^{B\cup C}$ be a jointly exchangeable matrix associated with ξ on $B \cup C$. Then there exist a random partition of $\mathbb{N}$ into subsequences (b_i) and (c_i) such that a.s.*

$$A^B_{ij} \coloneqq A^{B\cup C}_{b_i b_j},$$
$$A^C_{ij} \coloneqq A^{B\cup C}_{c_i c_j}$$

are jointly exchangeable matrices associated with ξ on B and C, respectively. Moreover, the matrix $A^{BC}_{ij} \coloneqq A^{B\cup C}_{b_i c_j}$ is separately exchangeable.

Proof. Let $f_{B\cup C}$ be a sampling function for ξ on $B \cup C$. Then there exists a representation

$$A_{ij}^{B\cup C} \overset{a.s.}{=} \mathbf{1}\big\{U_{ij} \le \widetilde{W}(f_{B\cup C}(U_i), f_{B\cup C}(U_j))\big\}, \qquad i,j \in \mathbb{N}.$$

For $k \in \mathbb{N}$, let b_k be the index of the k-th element of (U_i) lying in $f_{B\cup C}^{-1}(B)$, and define c_k analogously with $f_{B\cup C}^{-1}(C)$. Since $f_{B\cup C}$ maps into $B \cup C$, these preimages partition $[0,1]$ up to a null set, so (b_i) and (c_i) partition $\mathbb{N}$ and are each a.s. infinite. The subsequence of (U_i) in $f_{B\cup C}^{-1}(B)$ is i.i.d. uniform on $f_{B\cup C}^{-1}(B)$, so $(U_{b_i}) \overset{d}{=} (T_B(U_i))$; as $(U_{ij}) \perp (U_i)$ and (b_i) is (U_i)-measurable, the marks $(U_{b_i b_j})$ remain i.i.d. uniform. Hence, Proposition 26 shows that

$$A_{ij}^{B} \overset{a.s.}{=} \mathbf{1}\big\{U_{b_i b_j} \le \widetilde{W}(f_{B\cup C}(U_{b_i}), f_{B\cup C}(U_{b_j}))\big\}, \qquad i,j \in \mathbb{N}$$

is a jointly exchangeable matrix associated with ξ on B; the argument for A^C is analogous. Finally, (b_i) and (c_j) select disjoint, hence independent, sub-arrays of (U_i), so

$$(f_{B\cup C}(U_{b_i}), f_{B\cup C}(U_{c_j})) \overset{d}{=} (T_B(U_i), T_C(U'_j))$$

for an independent i.i.d. array (U'_j); as $(U_{b_i c_j})$ is independent of both groups, A^{BC} is separately exchangeable. □

For more background on separately exchangeable arrays, see Kallenberg (2005).

4 Graph Processes on Riemannian Manifolds

We now specialise to graph processes on a connected, geodesically complete d-dimensional Riemannian manifold (M, g), on which the Riemannian volume measure λ_g and the Riemannian distance d_g furnish natural notions of volume and distance. These single out a geometrically natural class of graph processes: the vertices form a Poisson process with intensity measure $\delta\lambda_g$, where $\delta > 0$, and the connection kernel is distance-based,

$$W(x,y) = h(d_g(x,y)), \quad x,y \in M,$$

where $h\colon \mathbb{R}_+ \to [0,1]$ is measurable and non-increasing. This models the assumption that nearby distinct vertices are more likely to be connected, and yields a manifold-valued analogue of the random connection model (Meester et al., 1997). For this class we prove a scale-invariance theorem (Theorem 29): for every $c > 0$, the simultaneous rescaling

$$g' = c^2 g, \qquad \delta' = c^{-d}\delta, \qquad h'(r) = h(r/c)$$

leaves the law of the induced graph process invariant. This theorem lets us normalise one of the parameters without loss of generality. The special case

$$h(r) = \mathbf{1}_{[0,t]}(r), \quad r \in \mathbb{R}_+,$$

is called the *threshold model* on M with density $\delta > 0$ and threshold $t > 0$. Here, two distinct vertices are adjacent exactly when $d_g(x,y) \le t$; this is the manifold-valued analogue of the random plane network of Gilbert (1961) and the object of study in the remaining sections.

We begin with a few general observations. As is standard, manifolds are taken to be second-countable and Hausdorff. Any two points of a connected smooth manifold can be joined by a piecewise regular curve of finite length, so for $x, y \in M$ the *Riemannian distance*

$$d_g(x,y) \coloneqq \inf\big\{\text{length}(\gamma) : \gamma \text{ a piecewise regular curve from } x \text{ to } y\big\}$$

on M induced by g is finite and well-defined. Indeed, d_g is a metric on M whose induced topology coincides with the manifold topology. Being second-countable, M is separable, so (M, d_g) is a

separable metric space; geodesic completeness, via the Hopf–Rinow theorem, renders it complete and yields that closed and d_g-bounded subsets of M are compact. Hence (M, d_g) is a complete separable metric space, and we may take $\widehat{\mathcal{B}}(M)$ to be the ring of d_g-bounded Borel sets. For details we refer to Lee (2018).

The next lemma constructs the *Riemannian volume measure* λ_g on M, the intrinsic volume measure induced by g, which we use to define uniform distributions on Borel subsets of M of finite volume. These in turn single out a geometrically natural class of graph processes on M.

Lemma 28. *There exists a unique Borel measure λ_g on M such that for every chart (V, φ) of M, with $\varphi = (x^1, \dots, x^d)$, and every $B \in \mathcal{B}(M)$ with $B \subseteq V$,*

$$\lambda_g(B) = \int_{\varphi(B)} \varphi_* \sqrt{|G|}\, d\lambda_d, \tag{11}$$

where $G \coloneqq \det[g_{jk}]$ is the determinant of the metric components $g_{jk} \coloneqq g\big(\frac{\partial}{\partial x^j}, \frac{\partial}{\partial x^k}\big)$ in the chart (V, φ), and $\varphi_* \sqrt{|G|} = \sqrt{|G|} \circ \varphi^{-1}$ denotes the pushforward of $\sqrt{|G|}$ on V along φ. Moreover, λ_g is diffuse, σ-finite, and satisfies $\lambda_g(B) < \infty$ for every $B \in \widehat{\mathcal{B}}(M)$.

Proof. Existence and uniqueness of λ_g are proved in Chapter XII.1 of Amann and Escher (2008). By Theorem XII.1.7(iii), λ_g is diffuse. By XII.1.7(ii), λ_g is Radon, so every $x \in M$ has an open neighborhood V_x with $\lambda_g(V_x) < \infty$. If $K \subseteq M$ is compact, finitely many V_x cover K, so $\lambda_g(K) < \infty$. By XII.1.7(i), M is σ-compact, so λ_g is σ-finite. Finally, if $B \in \widehat{\mathcal{B}}(M)$, then B is d_g-bounded, hence so is $\overline{B}$; being closed and bounded, $\overline{B}$ is compact by Hopf–Rinow, and therefore $\lambda_g(B) \le \lambda_g(\overline{B}) < \infty$. □

For $B \in \mathcal{B}(M)$ with $0 < \lambda_g(B) < \infty$, we call a random element of M with distribution $\lambda_g(B)^{-1}\, \lambda_g|_B$ *g-uniformly distributed* on B. For details we refer to Amann and Escher (2008).

As described at the beginning of the section, let ξ be a Poisson graph process on M with intensity measure $\delta\lambda_g$ for some $\delta > 0$, connected by the kernel $W\colon (x, y) \mapsto h(d_g(x, y))$. This setup possesses a scale invariance, which we now make precise.

Theorem 29 (Scale invariance)**.** *Let ξ be the Poisson graph process defined above. For any $c > 0$, let ξ' be obtained from ξ by the transformation*

$$g' = c^2 g, \qquad \delta' = c^{-d}\delta, \qquad h'(r) = h(r/c).$$

Then $\xi' \overset{d}{=} \xi$.

Proof. Consider $B \in \widehat{\mathcal{B}}(M)$. By Proposition 7 and Lemma 8, $\xi|_B$ is directed by $\sum_{i \le N} \delta_{X_i}$ with $N \sim \text{Poisson}(\delta\lambda_g(B))$ and $(X_i) \perp\!\!\!\perp N$ i.i.d. g-uniform on B, and is connected by $W = h \circ d_g$; likewise $\xi'|_B$ is directed by $\sum_{i \le N'} \delta_{X'_i}$ with $N' \sim \text{Poisson}(\delta'\lambda_{g'}(B))$ and $(X'_i) \perp\!\!\!\perp N'$ i.i.d. g'-uniform points, and connected by $h' \circ d_{g'}$.

Since $g' = c^2 g$, the metric components scale as $g'_{jk} = c^2 g_{jk}$, so the determinant in Lemma 28 satisfies $G' = c^{2d} G$ and $\sqrt{|G'|} = c^d \sqrt{|G|}$. Hence, by (11), $\lambda_{g'} = c^d \lambda_g$ on any chart (V, φ), and hence on $\mathcal{B}(M)$. In particular the g- and g'-uniform laws on B coincide, and

$$\delta'\lambda_{g'}(B) = c^{-d}\delta \cdot c^d \lambda_g(B) = \delta\lambda_g(B),$$

so the two directing measures share their count and location laws, hence agree in distribution.

For the kernels, $g' = c^2 g$ scales lengths by c, so $d_{g'} = c\, d_g$ and

$$h' \circ d_{g'} = h'(c\, d_g) = h(d_g) = h \circ d_g$$

by $h'(r) = h(r/c)$. Thus $\xi|_B$ and $\xi'|_B$ share their directing law and connection kernel, so $\xi|_B \overset{d}{=} \xi'|_B$. As $B \in \widehat{\mathcal{B}}(M)$ was arbitrary, $\xi' \overset{d}{=} \xi$. □

For the remainder of the paper we focus on the threshold model on (M, g). Because its edges are determined by the vertex locations, much of the subsequent analysis simplifies. The threshold model also shows that boundedly exchangeable graph processes need not be globally exchangeable, which is the content of the next result.

Proposition 30. *Let (M, g) satisfy $\lambda_g(M) = \infty$. A threshold model ξ on M with density $\delta > 0$ and threshold $t > 0$ admits no representation* (1) *whose adjacency matrix is jointly exchangeable.*

Proof. Since $\lambda_g(M) = \infty$, the directing Poisson process has infinitely many atoms a.s., so any representation (1) has an infinite adjacency matrix. Suppose such a representation has (A_{ij}) jointly exchangeable, and set $X_j := A_{1,j+1}$ for $j \geq 1$. Permutations of $\mathbb{N}$ fixing 1 act on A as joint permutations and send the first row to a permutation of itself, so $(X_j)_{j\geq 1}$ is an infinite exchangeable $\{0,1\}$-sequence. By de Finetti's theorem there is a random $\Theta \in [0,1]$ such that, conditionally on Θ, the X_j are i.i.d. Bernoulli(Θ).

The degree of vertex 1 is $D_1 = \sum_{j\geq 1} X_j$, which is finite a.s. by local finiteness. Conditionally on $\Theta = \theta$ with $\theta > 0$, the X_j are i.i.d. with $\sum_j \mathbb{P}(X_j = 1) = \infty$, so the second Borel–Cantelli lemma gives $D_1 = \infty$ a.s., a contradiction. Hence $\Theta = 0$ a.s., so $\mathbb{P}(A_{12} = 1) = 0$; by joint exchangeability every pair maps to $(1, 2)$ under a permutation, whence $\mathbb{P}(A_{ij} = 1) = 0$ for all $i \neq j$. By countability the graph is a.s. edgeless.

But for any $x_0 \in M$, the ball $B_{t/2}(x_0)$ has positive volume, so with positive probability it contains at least two atoms — which are then within distance t and hence adjacent. This contradicts a.s. edgelessness, so no jointly exchangeable representation exists. □

5 Applications

In this section we treat examples of the threshold model on a homogeneous Riemannian manifold (M, g) with $\lambda_g(M) = \infty$. On such a manifold we first show that the threshold model produces sparse graphs (Proposition 31) — a regime that graphon models, by the dichotomy recalled in Section 2.2, cannot attain. We then derive the limiting degree distribution and clustering coefficient of the infinite random graph G associated with a threshold model ξ. Euclidean and hyperbolic space are subsequently treated in detail: by the Killing–Hopf theorem, the complete, simply connected Riemannian manifolds of constant sectional curvature are, up to isometry and in each dimension, the sphere, Euclidean space, and hyperbolic space, and the infinite-volume assumption excludes the compact spherical case. Our examples concentrate on the 2-dimensional case, since exact formulas become unwieldy in higher dimensions. In particular, we show that while the degree law is fixed by the expected number of points in a connection ball, curvature shifts the clustering coefficient independently of the degree distribution.

Here, a connected Riemannian manifold is called *homogeneous* if its group of isometries acts transitively on it; this implies geodesic completeness (see, e.g., Lee, 2018), so we remain in the setting of Section 4. Since isometries preserve λ_g and map metric balls onto metric balls, the volume of the closed ball $B_t(x_0) \subseteq M$ centred at $x_0 \in M$ with radius $t > 0$ does not depend on x_0, and we write $V_t := \lambda_g(B_t(x_0))$. For the same reason, the probability that two independent g-uniform points in $B_t(x_0)$ lie at distance at most t from each other is independent of x_0; for reasons that will become apparent, we call it the *clustering tendency* of (M, g) at radius t, denoted by p_t.

Throughout this section, let ξ be a threshold model on (M, g) with density $\delta > 0$ and threshold $t > 0$, and let G be the random graph associated with ξ; as before, $v(G)$ and $e(G)$ denote its vertex and edge set. Since $\lambda_g(M) = \infty$, the graph G is almost surely infinite, and it is not immediately obvious how graph statistics such as the degree distribution or the clustering coefficient should be defined for it. One natural approach is to evaluate the statistic of interest

along a sequence of finite induced subgraphs of G and to pass to the limit; the limit, however, could a priori depend on the chosen sequence.

We therefore fix a sequence $E_1, E_2, \dots \in \mathcal{B}(M)$ with $\lambda_g(E_n) < \infty$ for all $n \in \mathbb{N}$ and $E_n \uparrow M$, and show that restricting G to the sets E_n yields limits in probability that do not depend on the choice of this sequence. We may impose the additional growth condition

$$\sum_{n=1}^{\infty} \frac{1}{\lambda_g(E_n)} < \infty, \tag{12}$$

under which our limits hold almost surely rather than merely in probability; it is satisfied, e.g., by balls $E_n = B_n(x_0)$ in $\mathbb{R}^d$ for $d \geq 2$.

5.1 Asymptotic sparsity

We begin by showing that G is asymptotically sparse: more precisely, the number of edges among the vertices in a growing observation window grows at most linearly in the number of such vertices. The proof relies on the univariate and multivariate Mecke equations (Last and Penrose, 2017, Theorems 4.1 and 4.4).

Proposition 31. *Set* $v_n \coloneqq \#\{v(G) \cap E_n\}$ *and* $e_n \coloneqq \#\big\{\{x,y\} \in e(G) \mid x, y \in E_n\big\}$. *Then, for every* $\varepsilon > 0$,

$$\lim_{n\to\infty} \mathbb{P}\left(e_n > \left(\frac{\delta V_t}{2} + \varepsilon \right) v_n \right) = 0.$$

If (12) *holds, then almost surely*

$$\limsup_{n\to\infty} \frac{e_n}{v_n} \leq \frac{\delta V_t}{2}.$$

Proof. Let ξ be directed by $\nu = \sum_{i=1}^{\infty} \delta_{S_i}$. The degree of any vertex S_i is $\nu(B_t(S_i)) - 1$. By the handshaking lemma,

$$e_n \leq \frac{1}{2} \sum_{S_i \in E_n} \big(\nu(B_t(S_i)) - 1\big) =: \frac{1}{2} h_n.$$

By the Mecke equation, in which the atom added at x cancels the -1,

$$\mathbb{E}[h_n] = \delta \int_{E_n} \mathbb{E}[\nu(B_t(x))]\, \lambda_g(dx) = \delta^2 V_t \lambda_g(E_n),$$

using that $\nu(B_t(x)) \sim \mathrm{Poisson}(\delta V_t)$ for all x. Next, the multivariate Mecke equation gives

$$\begin{aligned}
&\mathbb{E}\Big[h_n^2 - \sum_{S_i \in E_n} \big(\nu(B_t(S_i)) - 1\big)^2\Big] \\
&\quad = \mathbb{E}\Big[\sum_{\substack{S_i, S_j \in E_n \\ S_i \neq S_j}} \big(\nu(B_t(S_i)) - 1\big)\big(\nu(B_t(S_j)) - 1\big)\Big] \\
&\quad = \delta^2 \int_{\{d_g(x,y) \leq t\}} \mathbb{E}\big[(\nu(B_t(x)) + 1)(\nu(B_t(y)) + 1)\big]\, \lambda_g^{\otimes 2}(dx, dy) \\
&\qquad + \delta^2 \int_{\{d_g(x,y) > t\}} \mathbb{E}\big[\nu(B_t(x))\nu(B_t(y))\big]\, \lambda_g^{\otimes 2}(dx, dy) \\
&\quad = \delta^2 \int_{\{d_g(x,y) > 2t\}} \mathbb{E}[\nu(B_t(x))]\mathbb{E}[\nu(B_t(y))]\, \lambda_g^{\otimes 2}(dx, dy) + O(\lambda_g(E_n)) \\
&\quad = \mathbb{E}[h_n]^2 + O(\lambda_g(E_n)),
\end{aligned}$$

where all domains of integration are intersected with E_n^2. For the third equality, note that $B_t(x)$ and $B_t(y)$ are disjoint when $d_g(x,y) > 2t$, so $\nu(B_t(x))$ and $\nu(B_t(y))$ are independent; the complementary region $\{(x,y) \in E_n^2 : d_g(x,y) \le 2t\}$ has $\lambda_g^{\otimes 2}$-measure at most $V_{2t}\lambda_g(E_n)$ by homogeneity, while the integrands are bounded by a constant depending only on δV_t. Since

$$\mathbb{E}\Big[\sum_{S_i \in E_n} \big(\nu(B_t(S_i)) - 1\big)^2\Big] = \delta \int_{E_n} \mathbb{E}[\nu(B_t(x))^2]\, \lambda_g(dx) = O(\lambda_g(E_n)),$$

we obtain $\mathrm{Var}(h_n) = O(\lambda_g(E_n))$. As $E_n \uparrow M$ and $\lambda_g(M) = \infty$, we have $\lambda_g(E_n) \to \infty$, so Chebyshev's inequality yields $h_n/\lambda_g(E_n) \xrightarrow{\mathbb{P}} \delta^2 V_t$. Since $v_n = \nu(E_n) \sim \mathrm{Poisson}(\delta\lambda_g(E_n))$, the same argument gives $v_n/\lambda_g(E_n) \xrightarrow{\mathbb{P}} \delta$, and hence $\mathbb{P}\big(h_n > (\delta V_t + 2\varepsilon)v_n\big) \to 0$ for every $\varepsilon > 0$. As $e_n \le h_n/2$,

$$\mathbb{P}\Big(e_n > \big(\frac{\delta V_t}{2} + \varepsilon\big)v_n\Big) \le \mathbb{P}\big(h_n > (\delta V_t + 2\varepsilon)\, v_n\big) \longrightarrow 0.$$

Now suppose that (12) holds. The variance estimates above and Chebyshev's inequality imply, for every $\eta > 0$,

$$\sum_{n=1}^{\infty} \mathbb{P}\left(\left|\frac{h_n}{\lambda_g(E_n)} - \delta^2 V_t\right| > \eta\right) < \infty \quad \text{and} \quad \sum_{n=1}^{\infty} \mathbb{P}\left(\left|\frac{v_n}{\lambda_g(E_n)} - \delta\right| > \eta\right) < \infty.$$

Thus, by the first Borel–Cantelli lemma,

$$\frac{h_n}{\lambda_g(E_n)} \longrightarrow \delta^2 V_t \quad \text{and} \quad \frac{v_n}{\lambda_g(E_n)} \longrightarrow \delta \qquad \text{almost surely.}$$

In particular, $v_n > 0$ eventually almost surely, and therefore

$$\limsup_{n\to\infty} \frac{e_n}{v_n} \le \lim_{n\to\infty} \frac{h_n}{2v_n} = \frac{\delta V_t}{2} \qquad \text{almost surely.}$$

□

5.2 Asymptotic degree distribution

The *degree* of a vertex $x \in v(G)$ is $\deg(x) := \#\{y \in v(G) \mid \{x,y\} \in e(G)\}$; in the threshold model, $\deg(x) = \nu(B_t(x)) - 1$ for every vertex x, which is almost surely finite since ν is locally finite.

Using the same technique as in the proof of Proposition 31, we now derive the asymptotic degree distribution of G.

Proposition 32. *For $k \in \mathbb{N}_0$, write*

$$\widehat{p}_n(k) := \frac{1}{\nu(E_n)} \sum_{x \in v(G) \cap E_n} \mathbf{1}\{\deg(x) = k\}$$

on the event $\{\nu(E_n) > 0\}$, and set $\widehat{p}_n(k) := 0$ otherwise. Then

$$\widehat{p}_n(k) \xrightarrow{\mathbb{P}} e^{-\delta V_t} \frac{(\delta V_t)^k}{k!}.$$

If (12) holds, then this convergence holds almost surely.

Proof. Let ξ be directed by $\nu = \sum_{i=1}^{\infty} \delta_{S_i}$. A vertex S_i has degree k if and only if $\nu(B_t(S_i)) = k+1$, so the number of degree-k vertices in E_n is

$$H_n(k) \coloneqq \sum_{S_i \in E_n} \mathbf{1}\{\nu(B_t(S_i)) = k+1\}.$$

By the Mecke equation, in which the atom added at x accounts for the shift from $k+1$ to k,

$$\mathbb{E}[H_n(k)] = \delta \int_{E_n} \mathbb{P}\big(\nu(B_t(x)) = k\big)\, \lambda_g(dx) = \delta \lambda_g(E_n) e^{-\delta V_t} \frac{(\delta V_t)^k}{k!}.$$

As in the proof of Proposition 31, the multivariate Mecke equation and the bound on the region $\{d_g(x,y) \le 2t\}$ yield $\operatorname{Var}(H_n(k)) = O\big(\lambda_g(E_n)\big)$. Hence $H_n(k)/\lambda_g(E_n) \xrightarrow{\mathbb{P}} \delta e^{-\delta V_t}(\delta V_t)^k/k!$. Since $\nu(E_n)/\lambda_g(E_n) \xrightarrow{\mathbb{P}} \delta$ and $\mathbb{P}\big(\nu(E_n) = 0\big) \longrightarrow 0$, the first assertion follows.

Now suppose that (12) holds. By the variance bound above, Chebyshev's inequality and the same Borel–Cantelli argument as in the proof of Proposition 31 yield

$$\frac{H_n(k)}{\lambda_g(E_n)} \longrightarrow \delta e^{-\delta V_t} \frac{(\delta V_t)^k}{k!} \quad \text{and} \quad \frac{\nu(E_n)}{\lambda_g(E_n)} \longrightarrow \delta \qquad \text{almost surely.}$$

Hence $\nu(E_n) > 0$ eventually almost surely, and therefore

$$\widehat{p}_n(k) = \frac{H_n(k)/\lambda_g(E_n)}{\nu(E_n)/\lambda_g(E_n)} \longrightarrow e^{-\delta V_t} \frac{(\delta V_t)^k}{k!} \qquad \text{almost surely.}$$

□

Corollary 33. *Let*

$$\widehat{\mu}_n \coloneqq \sum_{k=0}^{\infty} \widehat{p}_n(k)\, \delta_k$$

denote the empirical degree distribution on $\mathbb{N}_0$, where $\widehat{p}_n(k)$ is defined as in Proposition 32, and let $\pi \coloneqq \mathrm{Poisson}(\delta V_t)$. Then

$$d_{\mathrm{TV}}(\widehat{\mu}_n, \pi) \xrightarrow{\mathbb{P}} 0.$$

If, in addition, (12) *holds, then this convergence holds almost surely.*

Proof. By Proposition 32,

$$\widehat{p}_n(k) \xrightarrow{\mathbb{P}} \pi(\{k\}) =: \pi_k$$

for every $k \in \mathbb{N}_0$. Since $\widehat{\mu}_n$ and π are probability measures on the countable space $\mathbb{N}_0$, pointwise convergence of their masses implies convergence in total variation in probability.

Now suppose that (12) holds. Then $\widehat{p}_n(k) \to \pi_k$ almost surely, separately for every $k \in \mathbb{N}_0$. Since $\mathbb{N}_0$ is countable, these convergences hold simultaneously outside a single null set. Since $\widehat{\mu}_n$ is eventually a probability measure almost surely, Kallenberg (2021, Lemma 1.34) yields

$$d_{\mathrm{TV}}(\widehat{\mu}_n, \pi) \longrightarrow 0 \qquad \text{almost surely.}$$

□

5.3 Asymptotic clustering coefficient

Note that the limiting degree distribution depends on the manifold, the density, and the threshold only through the product δV_t: on any homogeneous manifold, every Poisson law arises as the limiting degree distribution for a suitable choice of δ. The degree distribution alone therefore cannot distinguish the underlying geometries. This raises the question of whether different manifolds can produce asymptotically different graphs at all; to answer it affirmatively, we turn to the clustering coefficient of G.

The *local clustering coefficient* of a vertex $x \in v(G)$ is the proportion of pairs of neighbours of x that are themselves adjacent: writing $\tau(x)$ for the number of triangles in G containing x,

$$C(x) \coloneqq \frac{2\tau(x)}{\deg(x)(\deg(x)-1)},$$

with the convention that $C(x) \coloneqq 0$ if $\deg(x) \le 1$. The *clustering coefficient* of a nonempty finite set of vertices is the average of the local clustering coefficients of its members.

Proposition 34. *Write*

$$\widehat{C}_n \coloneqq \frac{1}{\nu(E_n)} \sum_{x \in v(G) \cap E_n} C(x)$$

on the event $\{\nu(E_n) > 0\}$*, and* $\widehat{C}_n \coloneqq 0$ *otherwise. Then*

$$\widehat{C}_n \xrightarrow{\mathbb{P}} \big(1 - e^{-\delta V_t}(1+\delta V_t)\big)\, p_t.$$

If, in addition, (12) *holds, then this convergence holds almost surely.*

Proof. Let ξ be directed by $\nu = \sum_{i=1}^{\infty} \delta_{S_i}$, and set

$$T_n \coloneqq \sum_{S_i \in E_n} C(S_i).$$

By homogeneity and the Mecke equation,

$$\mathbb{E}[T_n] = \delta \lambda_g(E_n) \mathbb{E}^x[C(x)],$$

where $\mathbb{E}^x$ denotes expectation for the process directed by $\nu + \delta_x$, and $x \in M$ is arbitrary. Writing $N_x \coloneqq \nu(B_t(x))$, the multivariate Mecke equation gives

$$\begin{aligned}
\mathbb{E}^x[C(x)] &= \mathbb{E}\Big[\sum_{\substack{S_i, S_j \in B_t(x) \\ S_i \neq S_j}} \frac{\mathbf{1}\{d_g(S_i,S_j) \le t\}}{N_x(N_x-1)} \mathbf{1}\{N_x \ge 2\} \Big] \\
&= \delta^2 \mathbb{E}\Big[\frac{1}{(N_x+2)(N_x+1)} \Big] \int_{B_t(x)^2} \mathbf{1}\{d_g(y,z) \le t\}\, \lambda_g(dy) \lambda_g(dz).
\end{aligned}$$

Since $N_x \sim \mathrm{Poisson}(\delta V_t)$,

$$\mathbb{E}\left[\frac{1}{(N_x+2)(N_x+1)} \right] = \frac{1 - e^{-\delta V_t}(1+\delta V_t)}{(\delta V_t)^2},$$

while the integral equals $V_t^2 p_t$. Hence, with

$$c_t \coloneqq \big(1 - e^{-\delta V_t}(1+\delta V_t)\big) p_t,$$

we have $\mathbb{E}[T_n] = \delta\lambda_g(E_n)c_t$. As in the proof of Proposition 31, the multivariate Mecke equation and the bound on the region $\{d_g(x,y) \le 2t\}$ yield $\mathrm{Var}(T_n) = O(\lambda_g(E_n))$. Thus $T_n/\lambda_g(E_n) \xrightarrow{\mathbb{P}} \delta c_t$. Together with $\nu(E_n)/\lambda_g(E_n) \xrightarrow{\mathbb{P}} \delta$ and $\mathbb{P}(\nu(E_n) = 0) \to 0$, this gives $\widehat{C}_n = T_n/\nu(E_n) \xrightarrow{\mathbb{P}} c_t$.

If (12) holds, the variance bound, Chebyshev's inequality, and the same Borel–Cantelli argument as in the proof of Proposition 31 give $T_n/\lambda_g(E_n) \to \delta c_t$ almost surely. Since that proof also gives $\nu(E_n)/\lambda_g(E_n) \to \delta$ almost surely, it follows that $\widehat{C}_n \to c_t$ almost surely. □

The limiting average clustering coefficient thus factorizes into $1 - e^{-\delta V_t}(1 + \delta V_t)$ — the probability that a Poisson(δV_t) variable is at least 2, and hence a function of the limiting degree distribution alone — and the purely geometric factor p_t. Two manifolds with matched δV_t therefore share their limiting degree distribution, yet may differ in their limiting clustering. Subsections 5.4 and 5.5 show that Euclidean and hyperbolic space do: their clustering tendencies differ at every threshold.

5.4 Euclidean Space

Attribute	Threshold model on $\mathbb{R}^d$
Distance	$d_{\bar{g}}(x,y) = \|x-y\|_2$
Euclidean volume element	$\lambda_d(dy) = r^{d-1}\,dr\,d\omega$
Ball volume	$\overline{V}_t^{\,d} = \dfrac{\pi^{d/2}}{\Gamma(d/2+1)}t^d$
Clustering tendency	$\overline{p}_d = 1 - \dfrac{3\sqrt{3}}{4\pi}$ for $d=2$
Sparse	Yes
Limiting degree law	Poisson($\delta\overline{V}_t^{\,d}$)
Limiting clustering coefficient	$\big(1 - e^{-\delta\overline{V}_t^{\,d}}(1+\delta\overline{V}_t^{\,d})\big)\overline{p}_d$
Globally exchangeable	No
Boundedly exchangeable	Yes
Associated matrix on $B_R(0)$	$A_{ij}^{B_R(0)} = \mathbf{1}\{i\neq j\}\mathbf{1}\{\|S_i - S_j\|_2 \le t\}$, with (S_i) i.i.d. uniform on $B_R(0)$
Sampling function for $B_R(0)$	$f_{B_R(0)}(u) = R\sqrt{u_1}\big(\cos(2\pi u_2), \sin(2\pi u_2)\big)$ for $d=2$
Local graphon for $B_R(0)$	$W_{B_R(0)}(u,v) = \mathbf{1}\{u\neq v\}\mathbf{1}\{\|f_{B_R(0)}(u) - f_{B_R(0)}(v)\|_2 \le t\}$

Table 1: Key attributes of the threshold model on d-dimensional Euclidean space with threshold $t > 0$ and density $\delta > 0$. Here u_1 and u_2 are obtained from $u \in [0,1]$ as described below.

The simplest example of our setting is d-dimensional Euclidean space $(\mathbb{R}^d, \overline{g})$, where $\overline{g}$ denotes the Euclidean metric; it has constant sectional curvature zero. The Riemannian distance induced by $\overline{g}$ is

$$d_{\overline{g}}(x,y) = \|x-y\|_2 = \sqrt{\sum_{i=1}^{d}(x_i - y_i)^2}, \qquad x, y \in \mathbb{R}^d.$$

Since $\sqrt{|G|} \equiv 1$ in the identity chart, the Riemannian volume measure on $\mathbb{R}^d$ induced by $\overline{g}$ is the Lebesgue measure λ_d. In spherical polar coordinates around the origin (see Appendix 7.2), the volume element takes the form

$$d\lambda_d = r^{d-1}\,dr\,d\omega,$$

where $d\omega$ denotes the surface-area element of the unit sphere S^{d-1}, with total mass $A(S^{d-1}) = 2\pi^{d/2}/\Gamma(d/2)$. Integrating, the volume of the closed ball $B_t(x)$ is

$$\overline{V}_t^{\,d} := \frac{A(S^{d-1})}{d}\,t^d = \frac{\pi^{d/2}}{\Gamma(d/2+1)}\,t^d,$$

independently of x.

By a dilation argument, the clustering tendency of d-dimensional Euclidean space at radius t does not depend on t; we denote it by $\overline{p}_d$. For $d = 2$, one can show

$$\overline{p}_2 = \frac{1}{\pi}\int_0^1 \Big(4x\arccos\big(\frac{x}{2}\big) - x^2\sqrt{4-x^2}\Big)\,dx = 1 - \frac{3\sqrt{3}}{4\pi} \approx 0.5865.$$

Consider the threshold model ξ on $\mathbb{R}^d$ with density $\delta > 0$ and threshold $t > 0$: points are placed according to a Poisson process with intensity measure $\delta\lambda_d$, and two distinct points are adjacent exactly when their distance is at most t. As noted after Proposition 10, ξ admits no global i.i.d. representation, and by Proposition 30 no representation of ξ has a jointly exchangeable adjacency matrix.

However, since its directing process is Poisson, ξ is boundedly exchangeable (see the remark following Definition 20). Hence every closed ball $B_R(x) \subseteq \mathbb{R}^d$ with $R > 0$ identifies an infinite jointly exchangeable matrix $A^{B_R(x)}$ via Theorem 21; the adjacency matrix of any i.i.d. representation of $\xi|_{B_R(x)}$ may be regarded as a submatrix of $A^{B_R(x)}$ of random finite dimension. Explicitly, since ξ is connected by the threshold kernel,

$$A_{ij}^{B_R(x)} \overset{a.s.}{=} \mathbf{1}\{i \neq j\}\mathbf{1}\{\|S_i - S_j\| \leq t\}, \qquad i, j \in \mathbb{N},$$

for an array (S_i) of i.i.d. uniformly distributed random elements of $B_R(x)$. By homogeneity, the distribution of $A^{B_R(x)}$ does not depend on x, and we set $x = 0$ without loss of generality.

By Proposition 23, $A^{B_R(0)}$ admits an Aldous–Hoover–Kallenberg representation (4) with a local graphon $W_{B_R(0)}\colon [0,1]^2 \to [0,1]$. To obtain one explicitly, it suffices to choose a sampling function for ξ on $B_R(0)$ — here, any measurable $f_{B_R(0)}\colon [0,1] \to B_R(0)$ such that $f_{B_R(0)}(U)$ is uniformly distributed on $B_R(0)$ when $U \sim \mathrm{U}[0,1]$ — and to set

$$W_{B_R(0)}(u,v) \coloneqq \mathbf{1}\{u \neq v\}\mathbf{1}\big\{\big\|f_{B_R(0)}(u) - f_{B_R(0)}(v)\big\| \leq t\big\};$$

different choices of the sampling function yield local graphons that are equivalent in the sense of Proposition 24. Sampling functions can be constructed by inverse-transform sampling in suitable coordinates, but the resulting formulas become unwieldy as the dimension increases. We therefore give an explicit construction only for $d = 2$.

Let $U \sim \mathrm{U}[0,1]$ and write

$$U = (0.a_1a_2a_3\ldots)_2$$

for the binary expansion chosen to contain infinitely many zeros. Split the binary digits into their odd and even subsequences,

$$U_1 \coloneqq (0.a_1a_3a_5\ldots)_2, \qquad U_2 \coloneqq (0.a_2a_4a_6\ldots)_2.$$

Then (U_1, U_2) is uniformly distributed on $[0,1]^2$. Consequently,

$$f_{B_R(0)}(U) \coloneqq R\sqrt{U_1}\big(\cos(2\pi U_2), \sin(2\pi U_2)\big)$$

is uniformly distributed on $B_R(0)$.

For $d = 2$, we illustrate Proposition 34 by simulation. We set $\delta = 1$ and $t = 1$, so that $\delta\overline{V}_t^2 = \pi$ and, by Corollary 33, the empirical degree distribution converges to the Poisson law with mean $\pi \approx 3.14$. The limiting clustering coefficient of Proposition 34 is $\big(1 - e^{-\pi}(1+\pi)\big)\,\overline{p}_2 \approx 0.482$. Taking $E_n = B_n(0)$, for which $\sum_n \lambda_2(B_n(0))^{-1} = \pi^{-1}\sum_n n^{-2} < \infty$, the convergence is almost sure by (12). We sampled 1000 independent realizations of G restricted to $B_{31}(0)$ and computed the clustering coefficient $\widehat{C}_R$ of the vertices in $B_R(0)$ for a range of radii $R \leq 30$; this ensures that the local clustering coefficients entering $\widehat{C}_R$ are free of boundary truncation. Figure 1 shows the concentration of $\widehat{C}_R$ around the limit, together with a single realization illustrating the almost-sure convergence.

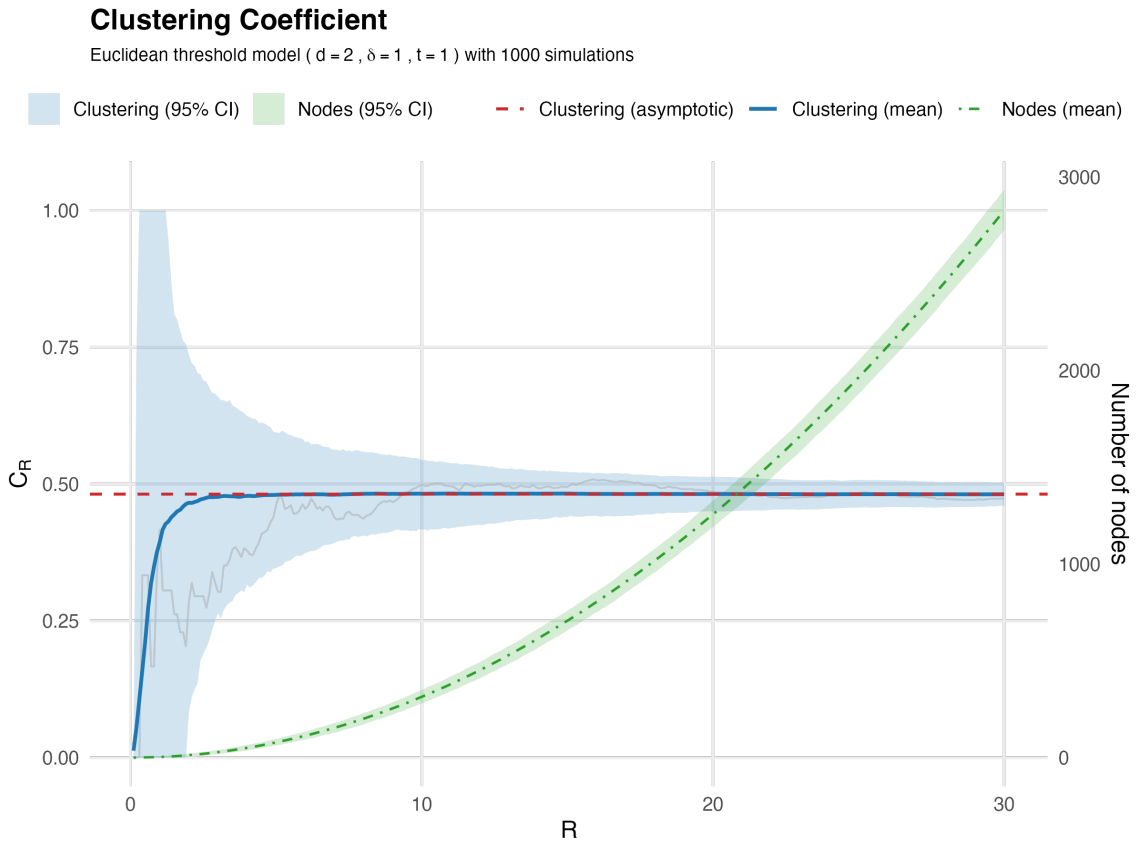

Figure 1: Empirical clustering coefficient $\widehat{C}_R$ of the 2-dimensional Euclidean threshold model with $\delta = 1$ and $t = 1$. Dark blue: mean of $\widehat{C}_R$ over 1000 realizations; light-blue band: empirical 2.5% and 97.5% quantiles; grey: one randomly selected realization; dashed red: the asymptotic limit ≈ 0.482 (Proposition 34). Dark green and light-green band (right axis): mean and corresponding quantiles of the number of vertices in $B_R(0)$.

Both limits above depend on the parameters (δ, t) only through the product $\delta\overline{V}_t^2$: the limiting degree law is Poisson$(\delta\overline{V}_t^2)$, and since the clustering tendency $\overline{p}_2$ does not depend on the threshold, the limiting average clustering coefficient is a fixed function of the same quantity. Within Euclidean space, matching the degree distribution therefore also fixes the clustering limit. In hyperbolic space, by contrast, the clustering tendency varies with the curvature, which decouples the two limits; we turn to this next.

5.5 Hyperbolic Space

Our second example is d-dimensional hyperbolic space, the complete simply connected manifold of constant negative sectional curvature. By Hilbert's theorem, 2-dimensional hyperbolic space cannot be isometrically embedded in 3-dimensional Euclidean space, making it more difficult to visualize.

For $\kappa > 0$, let $\mathbb{H}^d_\kappa$ be the set of points $(x_0, x_1, \ldots, x_d) \in \mathbb{R}^{d+1}$ with $x_0 > 0$ satisfying the relation

$$-x_0^2 + \sum_{i=1}^{d} x_i^2 = -\frac{1}{\kappa^2},$$

that is, the upper sheet of a two-sheeted hyperboloid in $\mathbb{R}^{d+1}$. The pullback of the Minkowski metric $\overline{q} = -\mathrm{d}x_0^2 + \sum_{i=1}^{d} \mathrm{d}x_i^2$ on $\mathbb{R}^{d+1}$ under the inclusion $\iota\colon \mathbb{H}^d_\kappa \to \mathbb{R}^{d+1}$ is positive definite on $\mathbb{H}^d_\kappa$ and hence induces a Riemannian metric $\breve{g} \coloneqq \iota^*\overline{q}$. We call the Riemannian manifold $(\mathbb{H}^d_\kappa, \breve{g})$

Attribute	Threshold model on $\mathbb{H}^d_\kappa$
Distance	$d_{\breve{g}}(x,y) = \frac{1}{\kappa}\operatorname{arcosh}\left(-\kappa^2\left(-x_0y_0 + \sum_{i=1}^d x_iy_i\right)\right)$
Hyperbolic volume element	$\lambda_{\breve{g}}(dy) = \frac{1}{\kappa^{d-1}}\sinh^{d-1}(\kappa r)\,dr\,d\omega$
Ball volume	$\breve{V}_t^{\,d}(\kappa) = \frac{2\pi^{d/2}}{\kappa^{d-1}\Gamma(d/2)}\int_0^t \sinh^{d-1}(\kappa r)\,dr$
Clustering tendency	$\breve{p}_d(\kappa t)$, see (13) for $d=2$
Sparse	Yes
Limiting degree law	Poisson$(\delta \breve{V}_t^{\,d}(\kappa))$
Limiting clustering coefficient	$\left(1 - e^{-\delta \breve{V}_t^{\,d}(\kappa)}(1+\delta\breve{V}_t^{\,d}(\kappa))\right)\breve{p}_d(\kappa t)$
Globally exchangeable	No
Boundedly exchangeable	Yes
Associated matrix on $B_R(o_\kappa)$	$A_{ij}^{B_R(o_\kappa)} = \mathbf{1}\{i\neq j\}\mathbf{1}\{d_{\breve{g}}(S_i,S_j)\le t\}$, with (S_i) i.i.d. uniform on $B_R(o_\kappa)$
Sampling function for $B_R(o_\kappa)$	$f_{B_R(o_\kappa)}(u)$, see (14) for $d=2$
Local graphon for $B_R(o_\kappa)$	$W_{B_R(o_\kappa)}(u,v) = \mathbf{1}\{u\neq v\}\mathbf{1}\{d_{\breve{g}}(f_{B_R(o_\kappa)}(u), f_{B_R(o_\kappa)}(v))\le t\}$

Table 2: Key attributes of the threshold model on d-dimensional hyperbolic space with curvature parameter $\kappa > 0$, threshold $t > 0$, and density $\delta > 0$.

the d-dimensional hyperbolic space of radius $1/\kappa$. It is homogeneous with $\lambda_{\breve{g}}(\mathbb{H}^d_\kappa) = \infty$ and has constant sectional curvature $-\kappa^2$, so it falls within the scope of this section. See Lee (2018) for details.

For $x, y \in \mathbb{H}^d_\kappa$, the Riemannian distance induced by $\breve{g}$ is given by

$$d_{\breve{g}}(x,y) = \frac{1}{\kappa}\operatorname{arcosh}\left(-\kappa^2\left(-x_0y_0 + \sum_{i=1}^d x_iy_i\right)\right).$$

Again, it is useful to switch to spherical polar coordinates (see Appendix 7.3), in which the Riemannian volume measure induced by $\breve{g}$ satisfies

$$\mathrm{d}\lambda_{\breve{g}} = \frac{1}{\kappa^{d-1}}\sinh^{d-1}(\kappa r)\,\mathrm{d}r\,\mathrm{d}\omega.$$

Integrating, the volume of the closed ball $B_t(x) \subseteq \mathbb{H}^d_\kappa$ is

$$\breve{V}_t^d(\kappa) \coloneqq \frac{2\pi^{d/2}}{\kappa^{d-1}\Gamma(d/2)}\int_0^t \sinh^{d-1}(\kappa r)\,\mathrm{d}r,$$

independently of x; for $d=2$ this reduces to $\breve{V}_t^2(\kappa) = 2\pi(\cosh(\kappa t)-1)/\kappa^2$.

By a rescaling argument, the clustering tendency of $\mathbb{H}^d_\kappa$ at radius t depends only on the product $\alpha \coloneqq \kappa t$, so we denote it by $\breve{p}_d(\alpha)$. For $d=2$, one can show that

$$\breve{p}_2(\alpha) = \int_0^1 \frac{2\alpha\sinh(\alpha x)}{\pi(\cosh(\alpha)-1)^2}\Bigg[\cosh(\alpha)\arccos\left(\coth(\alpha)\tanh\left(\frac{\alpha x}{2}\right)\right) - \arctan\left(\frac{\sqrt{\cosh^2(\alpha)-\cosh^2(\frac{\alpha x}{2})}}{\sinh(\frac{\alpha x}{2})}\right)\Bigg]\mathrm{d}x. \tag{13}$$

The hyperbolic clustering tendency satisfies the following, proved in Appendix 7.4.

Proposition 35. *Let $d \ge 2$. The Euclidean and hyperbolic clustering tendencies satisfy:*

(i) $\breve{p}_d(\alpha) < \overline{p}_d$ *for all* $\alpha > 0$,

(ii) $\alpha \mapsto \breve{p}_d(\alpha)$ *is strictly decreasing on* $(0, \infty)$,

(iii) $\breve{p}_2(\alpha) = \overline{p}_2 - \big(\sqrt{3}/(16\pi)\big)\alpha^2 + O(\alpha^4)$ *as* $\alpha \to 0^+$,

(iv) $\breve{p}_2(\alpha) = (8/\pi)e^{-\alpha/2} + o(e^{-\alpha/2})$ *as* $\alpha \to \infty$.

Consider the threshold model ξ on $\mathbb{H}^d_\kappa$ with density $\delta > 0$ and threshold $t > 0$: points are placed according to a Poisson process with intensity measure $\delta\lambda_{\breve{g}}$, and two distinct points are adjacent exactly when their distance is at most t. As noted after Proposition 10, ξ admits no global i.i.d. representation, and by Proposition 30 no representation of ξ has a jointly exchangeable adjacency matrix.

Instead, since its directing process is Poisson, ξ is boundedly exchangeable (see the remark following Definition 20). Again, every closed ball $B_R(x) \subseteq \mathbb{H}^d_\kappa$ with $R > 0$ identifies an infinite jointly exchangeable matrix $A^{B_R(x)}$ via Theorem 21, and the adjacency matrix of any i.i.d. representation of $\xi|_{B_R(x)}$ may be regarded as a submatrix of $A^{B_R(x)}$ of random finite dimension. Explicitly, since ξ is connected by the threshold kernel,

$$A^{B_R(x)}_{ij} \overset{a.s.}{=} \mathbf{1}\{i \neq j\}\mathbf{1}\{d_{\breve{g}}(S_i, S_j) \leq t\}, \qquad i, j \in \mathbb{N},$$

for an array (S_i) of i.i.d. uniformly distributed random elements of $B_R(x)$. By homogeneity, the distribution of $A^{B_R(x)}$ does not depend on x, so it suffices to consider the hyperbolic origin $o_\kappa := (\kappa^{-1}, 0, \ldots, 0)$.

By Proposition 23, $A^{B_R(o_\kappa)}$ admits an Aldous–Hoover–Kallenberg representation (4) with a local graphon $W_{B_R(o_\kappa)} : [0,1]^2 \to [0,1]$. Again, it suffices to choose a sampling function for ξ on $B_R(o_\kappa)$ — here, any measurable $f_{B_R(o_\kappa)} : [0,1] \to B_R(o_\kappa)$ such that $f_{B_R(o_\kappa)}(U)$ is uniformly distributed on $B_R(o_\kappa)$ when $U \sim \mathrm{U}[0,1]$ — and to set

$$W_{B_R(o_\kappa)}(u, v) := \mathbf{1}\{u \neq v\}\mathbf{1}\{d_{\breve{g}}(f_{B_R(o_\kappa)}(u), f_{B_R(o_\kappa)}(v)) \leq t\};$$

different choices of the sampling function yield local graphons that are equivalent in the sense of Proposition 24. Given the Riemannian volume measure in spherical polar coordinates, we construct a sampling function $f_{B_R(o_\kappa)}$ explicitly for $d = 2$. With (U_1, U_2) obtained via the same binary expansion of $U \sim \mathrm{U}[0,1]$ as before,

$$f_{B_R(o_\kappa)}(U) := \frac{1}{\kappa}\begin{pmatrix} U_1(\cosh(\kappa R) - 1) + 1 \\ \sqrt{(U_1(\cosh(\kappa R) - 1) + 1)^2 - 1}\,\cos(2\pi U_2) \\ \sqrt{(U_1(\cosh(\kappa R) - 1) + 1)^2 - 1}\,\sin(2\pi U_2) \end{pmatrix} \tag{14}$$

is uniformly distributed on $B_R(o_\kappa)$.

For $d = 2$, we again illustrate Proposition 34 by simulation, now on $\mathbb{H}^2_2$. To compare against the Euclidean threshold model, we set $\delta = 1$ and $t = \operatorname{arcosh}(3)/2$, so that $\delta\breve{V}^2_t(2) = \pi$ and, by Corollary 33, the empirical degree distribution converges to the Poisson law with mean $\pi \approx 3.14$, as in the Euclidean case. The limiting clustering coefficient of Proposition 34 is $\big(1 - e^{-\pi}(1 + \pi)\big)\,\breve{p}_2(\operatorname{arcosh}(3)) \approx 0.404$, strictly below the Euclidean value ≈ 0.482. Taking $E_n = B_n(0)$, condition (12) holds since $\lambda_{\breve{g}}(B_n(o_\kappa))$ grows exponentially in n, so the convergence is almost sure. We sampled 1000 independent realizations of G restricted to $B_{4.4}(o_\kappa)$ and computed the clustering coefficient $\widehat{C}_R$ of the vertices in $B_R(o_\kappa)$ for a range of radii $R \leq 3.5$; this ensures that the local clustering coefficients entering $\widehat{C}_R$ are free of boundary truncation. Figure 2 shows the concentration of $\widehat{C}_R$ around the limit, together with a single realization illustrating the almost-sure convergence.

Although G again has a Poisson degree distribution with mean π, it exhibits strictly lower asymptotic clustering than its Euclidean counterpart. This illustrates Proposition 35(i): the

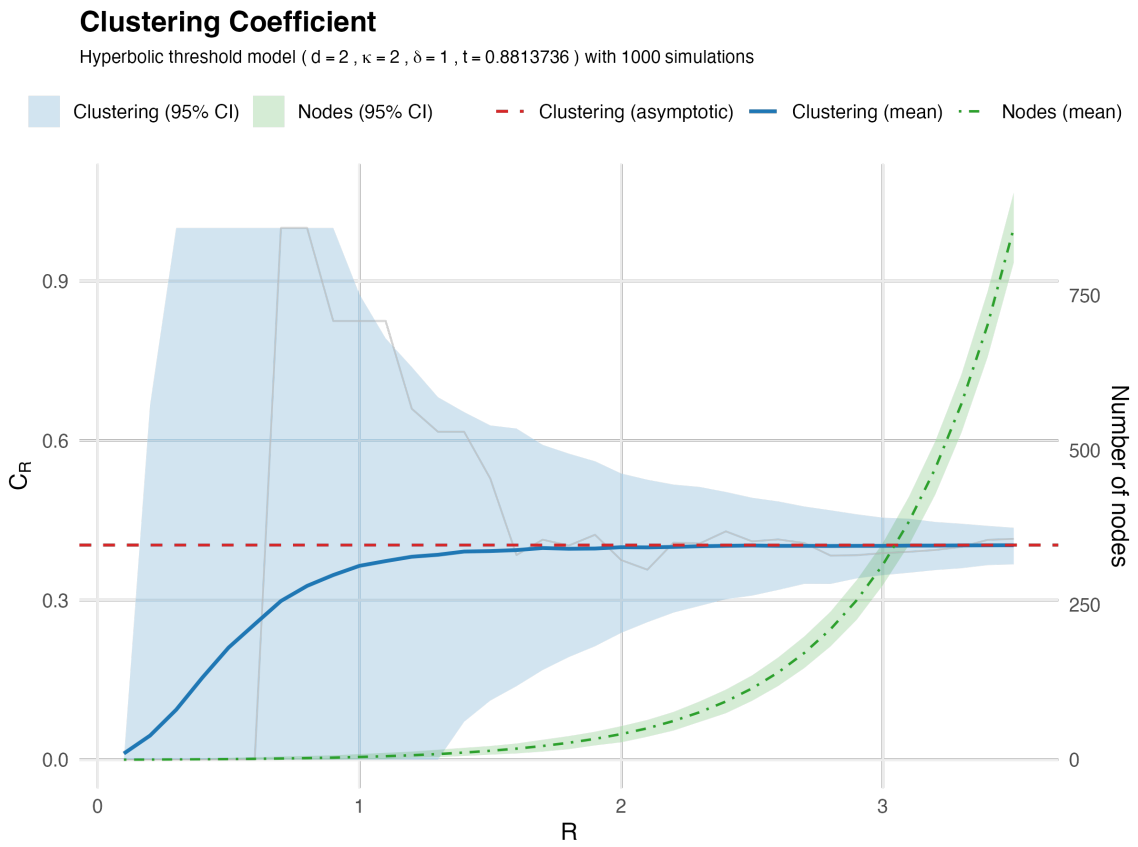

Figure 2: Empirical clustering coefficient $\widehat{C}_R$ of the 2-dimensional hyperbolic threshold model with $\delta = 1$ and $t = \operatorname{arcosh}(3)/2$. Dark blue: mean of $\widehat{C}_R$ over 1000 realizations; light-blue band: empirical 2.5% and 97.5% quantiles; grey: one randomly selected realization; dashed red: the asymptotic limit ≈ 0.404 (Proposition 34). Dark green and light-green band (right axis): mean and corresponding quantiles of the number of vertices in $B_R(o_\kappa)$.

hyperbolic clustering tendency is strictly below the Euclidean one, $\breve{p}_2(\alpha) < \overline{p}_2$ for all $\alpha > 0$ (see Figure 3). Hyperbolic space therefore tunes the asymptotic clustering coefficient independently of the asymptotic degree distribution: at a matched degree distribution, the hyperbolic threshold model always produces a strictly lower asymptotic clustering coefficient.

6 Conclusion

This paper studies graph processes — a class of random graph models in which vertices are identified with the points of a point process on a latent Polish space, and edges between distinct vertices are drawn conditionally independently via a connection kernel. Boundedly exchangeable graph processes recover classical matrix exchangeability on bounded observation windows: every nontrivial bounded restriction determines an infinite jointly exchangeable adjacency matrix, unique in distribution, and hence a local graphon representation. The global graph, by contrast, may be sparse and admit no jointly exchangeable representation. The threshold model on Riemannian manifolds serves as a geometrically natural example, demonstrating that the geometry of the latent space is directly reflected in the macroscopic properties of the network. Once the limiting degree distribution is matched across manifolds by a joint choice of density and threshold, curvature acts as an additional geometric parameter affecting the limiting clustering coefficient: at the matched degree law, hyperbolic space produces strictly smaller limiting clustering than Euclidean space. This makes graph processes a flexible and geometrically

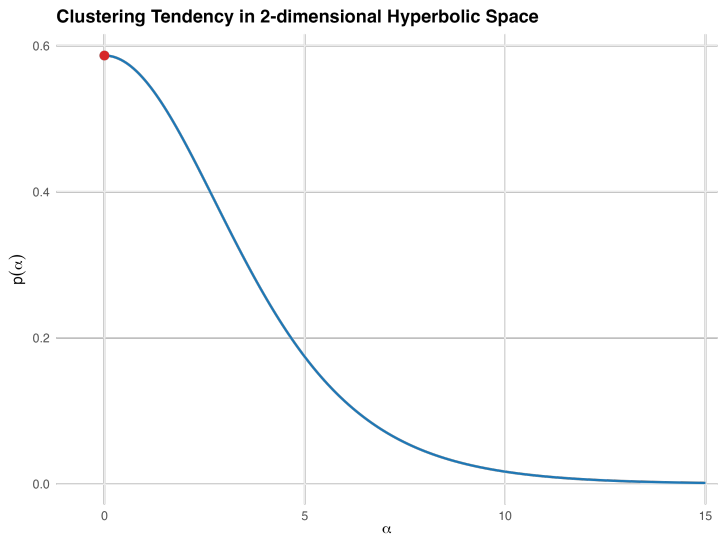

Figure 3: The clustering tendency $\breve{p}_2(\alpha)$ of 2-dimensional hyperbolic space, plotted as a function of $\alpha := \kappa t$ (see (13)). As $\alpha \to 0^+$ it converges to its Euclidean counterpart $\overline{p}_2 = 1 - 3\sqrt{3}/(4\pi) \approx 0.5865$, marked by the red dot, with quadratic approach by Proposition 35(iii); as $\alpha \to \infty$ it decays to zero, exponentially by Proposition 35(iv).

interpretable tool for modelling large sparse networks.

Several natural avenues for future research present themselves. On the modelling side, the asymptotic analysis of Section 5 could be extended from the threshold model to general distance-based kernels $W(x, y) = h(d_g(x, y))$, whose connection probabilities decay smoothly with distance rather than cutting off sharply, or to manifolds of locally varying curvature. The latter would forgo the homogeneity underlying our basepoint-free degree and clustering statistics, but could let these structural properties — such as clustering or degree heterogeneity — vary across different parts of the graph. On the statistical side, the explicit dependence of the limiting degree and clustering statistics on curvature suggests using them to infer model parameters, and the geometry of the latent space itself, from observed network data — a significant step towards practical applicability.

Acknowledgements

The authors acknowledge the use of the AI-assisted tools ChatGPT (OpenAI; several versions of GPT-5; accessed via chatgpt.com) and Claude (Anthropic; Claude 4.8 Opus and Claude 5 Sonnet, accessed via claude.ai).

These tools were used to suggest improvements to wording and grammar, to flag possible minor errors or omissions in the mathematical exposition for subsequent independent review, and to assist with LaTeX table formatting and code used to format plots. The authors evaluated every suggestion and independently verified all mathematical statements, proofs, references, numerical values, tables, and figures.

All scientific content, interpretations, and conclusions were fully created, verified, and approved by the authors. The authors take full responsibility for the content of the manuscript.

Funding information

Leon-Roman Rinke's position at Leibniz Universität Hannover was fully financed through the dual PhD programme of the House of Insurance, which is supported by the insurance industry.

This financing entailed no obligations regarding the research, and the supporting institutions had no role in the design or conduct of the research, its analysis or interpretation, or the preparation of this manuscript.

Competing interests

There are no competing interests to declare which arose during the preparation or publication process of this article.

References

Aldous, David J. (1981). "Representations for partially exchangeable arrays of random variables". In: *Journal of Multivariate Analysis* 11 (4), pp. 581–598.

Amann, Herbert and Joachim Escher (2008). *Analysis III*. 2nd ed. Birkhäuser Basel.

Boguñá, Marián, Ivan Bonamassa, Manlio De Domenico, Shlomo Havlin, Dmitri Krioukov, and M. Ángeles Serrano (2021). "Network geometry". In: *Nature Reviews Physics* 3, pp. 114–135.

Boguñá, Marián, Fragkiskos Papadopoulos, and Dmitri Krioukov (2010). "Sustaining the Internet with hyperbolic mapping". In: *Nature Communications* 1, 62.

Borgs, Christian, Jennifer T. Chayes, Henry Cohn, and Nina Holden (2018). "Sparse exchangeable graphs and their limits via graphon processes". In: *Journal of Machine Learning Research* 18 (210), pp. 1–71.

Borgs, Christian, Jennifer T. Chayes, Henry Cohn, and Victor Veitch (2019). "Sampling perspectives on sparse exchangeable graphs". In: *The Annals of Probability* 47 (5), pp. 2754–2800.

Borgs, Christian, Jennifer T. Chayes, László M. Lovász, Vera T. Sós, Balázs Szegedy, and Katalin Vesztergombi (2006). "Graph limits and parameter testing". In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, pp. 261–270.

Borgs, Christian, Jennifer T. Chayes, László M. Lovász, Vera T. Sós, and Katalin Vesztergombi (2008). "Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing". In: *Advances in Mathematics* 219 (6), pp. 1801–1851.

Caron, François and Emily B. Fox (2017). "Sparse Graphs Using Exchangeable Random Measures". In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 79 (5), pp. 1295–1366.

Caron, François, Francesca Panero, and Judith Rousseau (2023). "On sparsity, power-law, and clustering properties of graphex processes". In: *Advances in Applied Probability* 55 (4), pp. 1211–1253.

Diaconis, Persi and Svante Janson (2008). "Graph limits and exchangeable random graphs". In: *Rendiconti di Matematica e delle sue Applicazioni* 28 (1), pp. 33–61.

Erdős, Paul and Alfréd Rényi (1959). "On random graphs I." In: *Publicationes Mathematicae Debrecen* 6, pp. 290–297.

García-Pérez, Guillermo, Marián Boguñá, and M. Ángeles Serrano (2018). "Multiscale unfolding of real networks by geometric renormalization". In: *Nature Physics* 14, pp. 583–589.

Gilbert, Edgar N. (1959). "Random graphs". In: *The Annals of Mathematical Statistics* 30 (4), pp. 1141–1144.

— (1961). "Random Plane Networks". In: *Journal of the Society for Industrial and Applied Mathematics* 9 (4), pp. 533–543.

Grygierek, Jens and Christoph Thäle (2020). "Gaussian fluctuations for edge counts in high-dimensional random geometric graphs". In: *Statistics & Probability Letters* 158, 108674.

Gugelmann, Luca, Konstantinos Panagiotou, and Ueli Peter (2012). "Random Hyperbolic Graphs: Degree Sequence and Clustering". In: *Automata, Languages, and Programming*. Ed. by Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer. Vol. 7392. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 573–585.

Handcock, Mark S., Adrian E. Raftery, and Jeremy M. Tantrum (2007). "Model-based clustering for social networks". In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 170 (2), pp. 301–354.

Hoff, Peter D., Adrian E. Raftery, and Mark S. Handcock (2002). "Latent Space Approaches to Social Network Analysis". In: *Journal of the American Statistical Association* 97 (460), pp. 1090–1098.

Holland, Paul W., Kathryn B. Laskey, and Samuel Leinhardt (1983). "Stochastic blockmodels: First steps". In: *Social Networks* 5 (2), pp. 109–137.

Hoover, Douglas N. (1979). *Relations on Probability Spaces and Arrays of Random Variables*. Preprint. Princeton, NJ: Institute for Advanced Study.

Janson, Svante (2022). "On convergence for graphexes". In: *European Journal of Combinatorics* 104, 103549.

Kallenberg, Olav (1989). "On the representation theorem for exchangeable arrays". In: *Journal of Multivariate Analysis* 30 (1), pp. 137–154.

— (1990). "Exchangeable random measures in the plane". In: *Journal of Theoretical Probability* 3, pp. 81–136.

— (2005). *Probabilistic Symmetries and Invariance Principles*. Probability and its Applications. New York, NY: Springer.

— (2017). *Random Measures, Theory and Applications*. Vol. 77. Probability Theory and Stochastic Modelling. Springer Cham.

— (2021). *Foundations of Modern Probability*. 3rd ed. Vol. 99. Probability Theory and Stochastic Modelling. Springer International Publishing.

Kechris, Alexander S. (1995). *Classical Descriptive Set Theory*. Graduate Texts in Mathematics. Springer New York, NY.

Krioukov, Dmitri, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá (2010). "Hyperbolic geometry of complex networks". In: *Physical Review E* 82, 036106.

Last, Günter, Franz Nestmann, and Matthias Schulte (2021). "The random connection model and functions of edge-marked Poisson processes: Second order properties and normal approximation". In: *The Annals of Applied Probability* 31 (1), pp. 128–168.

Last, Günter and Mathew Penrose (2017). *Lectures on the Poisson Process*. Vol. 7. Institute of Mathematical Statistics Textbooks. Cambridge University Press.

Lee, John M. (2018). *Introduction to Riemannian Manifolds*. 2nd ed. Vol. 176. Graduate Texts in Mathematics. Springer Cham.

Lovász, László M. (2012). *Large Networks and Graph Limits*. Vol. 60. Colloquium Publications. Providence, RI: American Mathematical Society.

Lovász, László M. and Balázs Szegedy (2006). "Limits of dense graph sequences". In: *Journal of Combinatorial Theory, Series B* 96 (6), pp. 933–957.

— (2010). "Regularity Partitions and The Topology of Graphons". In: *An Irregular Mind: Szemerédi is 70*. Vol. 21. Bolyai Society Mathematical Studies. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 415–446.

Lubold, Shane, Arun G. Chandrasekhar, and Tyler H. McCormick (2023). "Identifying the latent space geometry of network models through analysis of curvature". In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 85 (2), pp. 240–292.

McCormick, Tyler H. and Tian Zheng (2015). "Latent surface models for networks using aggregated relational data". In: *Journal of the American Statistical Association* 110 (512), pp. 1684–1695.

Meester, Ronald, Mathew D. Penrose, and Anish Sarkar (1997). "The random connection model in high dimensions". In: *Statistics & Probability Letters* 35 (2), pp. 145–153.
Meester, Ronald and Rahul Roy (1996). *Continuum Percolation*. Vol. 119. Cambridge Tracts in Mathematics. Cambridge University Press.
Newman, Mark (2018). *Networks*. 2nd ed. Oxford University Press.
Orbanz, Peter and Daniel M. Roy (2015). "Bayesian Models of Graphs, Arrays and Other Exchangeable Random Structures". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37 (2), pp. 437–461.
Penrose, Mathew (2003). *Random Geometric Graphs*. Vol. 5. Oxford Studies in Probability. Oxford: Oxford University Press.
Raftery, Adrian E., Xiaoyue Niu, Peter D. Hoff, and Ka Yee Yeung (2012). "Fast inference for the latent space network model using a case-control approximate likelihood". In: *Journal of Computational and Graphical Statistics* 21 (4), pp. 901–919.
Rastelli, Riccardo, Florian Maire, and Nial Friel (2024). "Computationally efficient inference for latent position network models". In: *Electronic Journal of Statistics* 18 (1), pp. 2531–2570.
Reitzner, Matthias, Matthias Schulte, and Christoph Thäle (2017). "Limit theory for the Gilbert graph". In: *Advances in Applied Mathematics* 88, pp. 26–61.
Rosen, Daniel, Matthias Schulte, Christoph Thäle, and Vanessa Trapp (2025). "The radial spanning tree in hyperbolic space: Degree and edge-length". In: *Electronic Journal of Probability* 30, pp. 1–35.
Salter-Townshend, Michael and Tyler H. McCormick (2017). "Latent space models for multiview network data". In: *The Annals of Applied Statistics* 11 (3), pp. 1217–1244.
Sambale, Holger, Christoph Thäle, and Tara Trauthwein (2025). "Central limit theorems for the nearest neighbour embracing graph in Euclidean and hyperbolic space". In: *Stochastic Processes and their Applications* 188, 104671.
Shalizi, Cosma and Dena Asta (2024). "Consistency of maximum likelihood for continuous-space network models I". In: *Electronic Journal of Statistics* 18 (1), pp. 335–354.
Smith, Anna L., Dena M. Asta, and Catherine A. Calder (2019). "The Geometry of Continuous Latent Space Models for Network Data". In: *Statistical Science* 34 (3), pp. 428–453.
Veitch, Victor and Daniel M. Roy (2015). "The class of random graphs arising from exchangeable random measures". In: *arXiv preprint arXiv:1512.03099*.
— (2019). "Sampling and estimation for (sparse) exchangeable graphs". In: *The Annals of Statistics* 47 (6), pp. 3274–3299.
Wilkins-Reeves, Steven and Tyler H. McCormick (2022). "Asymptotically Normal Estimation of Local Latent Network Curvature". In: *arXiv preprint arXiv:2211.11673*.

7 Appendix

7.1 Proof of Lemma 17

Proof of Lemma 17. For $m \in \mathbb{N}$, write $\mathcal{P}_m$ for the σ-algebra generated by the dyadic partition of $[0,1]$ into $N_m \coloneqq 2^m$ intervals $\{I_{m,1}, \dots, I_{m,N_m}\}$ of equal length. The process

$$W_m \coloneqq \mathbb{E}\big[W \mid \mathcal{P}_m^{\otimes 2}\big].$$

is a martingale with respect to the filtration $(\mathcal{P}_m^{\otimes 2})_{m\in\mathbb{N}}$. Observe that W_m is constant on every rectangle $I_{m,i} \times I_{m,j}$. Since W is bounded and the dyadic rectangles generate $\mathcal{B}([0,1]^2)$, martingale convergence (see, e.g., Kallenberg, 2021, Theorem 9.24) yields

$$\|W_m - W\|_1 \longrightarrow 0.$$

Every permutation σ of $\{1, \dots, N_m\}$ is induced, up to changes at finitely many endpoints, by a λ-preserving bijection T_σ satisfying

$$T_\sigma(I_{m,i}) = I_{m,\sigma(i)}.$$

The assumed invariance of W therefore implies that the averages of W over $I_{m,i} \times I_{m,j}$ and $I_{m,\sigma(i)} \times I_{m,\sigma(j)}$ coincide. Since the symmetric group acts transitively on the diagonal pairs (i,i) and on the off-diagonal pairs (i,j), $i \neq j$, there exist $\alpha_m, \beta_m \in [0,1]$ such that

$$W_m = \alpha_m \quad \text{off } \Delta_m, \qquad W_m = \beta_m \quad \text{on } \Delta_m,$$

where

$$\Delta_m \coloneqq \bigcup_{i=1}^{N_m} I_{m,i}^2, \qquad \lambda^{\otimes 2}(\Delta_m) = \frac{1}{N_m}.$$

Let $c \coloneqq \int_{[0,1]^2} W \, d\lambda^{\otimes 2}$. Since conditional expectation preserves integrals,

$$c = \left(1 - \frac{1}{N_m}\right) \alpha_m + \frac{1}{N_m} \beta_m.$$

Hence $|\alpha_m - c| \leq N_m^{-1}$, while $\|W_m - \alpha_m\|_1 \leq N_m^{-1}$. Consequently,

$$\|W_m - c\|_1 \leq \frac{2}{N_m} \longrightarrow 0.$$

Together with $\|W_m - W\|_1 \to 0$, the triangle inequality gives $\|W - c\|_1 = 0$. Thus

$$W = c \qquad \lambda^{\otimes 2}\text{-a.e.}$$

□

7.2 Spherical polar coordinates in euclidean space

In spherical polar coordinates, a point $(x_1, \ldots, x_d) \in \mathbb{R}^d$ $(d \geq 2)$ is given by

$$\begin{aligned}
x_1 &= r\cos(\theta_1) \\
x_2 &= r\sin(\theta_1)\cos(\theta_2) \\
x_3 &= r\sin(\theta_1)\sin(\theta_2)\cos(\theta_3) \\
&\vdots \\
x_{d-1} &= r\sin(\theta_1)\cdots\sin(\theta_{d-2})\cos(\theta_{d-1}) \\
x_d &= r\sin(\theta_1)\cdots\sin(\theta_{d-2})\sin(\theta_{d-1})
\end{aligned}$$

for $r \in (0,\infty)$, $\theta_1, \ldots, \theta_{d-2} \in (0,\pi)$, and $\theta_{d-1} \in (0, 2\pi)$, with the convention that for $d = 2$ the only angle is $\theta_1 \in (0, 2\pi)$. This parametrizes $\mathbb{R}^d$ up to a λ_d-null set, namely the half-hyperplane $\{x_d = 0,\ x_{d-1} \geq 0\}$.

In these coordinates, the volume element of the Lebesgue measure λ_d on $\mathbb{R}^d$ satisfies

$$\mathrm{d}\lambda_d = r^{d-1} \sin^{d-2}(\theta_1) \sin^{d-3}(\theta_2) \cdots \sin(\theta_{d-2}) \, \mathrm{d}r \, \mathrm{d}\theta_1 \cdots \mathrm{d}\theta_{d-1}.$$

We write

$$\mathrm{d}\omega \coloneqq \sin^{d-2}(\theta_1) \sin^{d-3}(\theta_2) \cdots \sin(\theta_{d-2}) \, \mathrm{d}\theta_1 \cdots \mathrm{d}\theta_{d-1},$$

the surface area element of the unit sphere S^{d-1}, with total mass $\int_{S^{d-1}} \mathrm{d}\omega = 2\pi^{d/2}/\Gamma(d/2)$, to obtain

$$\mathrm{d}\lambda_d = r^{d-1} \, \mathrm{d}r \, \mathrm{d}\omega.$$

7.3 Spherical polar coordinates in hyperbolic space

We proceed similarly to Euclidean space. Define a point $(x_0, x_1, \dots, x_d) \in \mathbb{H}^d_\kappa$ $(d \geq 2)$ by

$$\begin{aligned}
x_0 &= \frac{1}{\kappa}\cosh(\kappa r) \\
x_1 &= \frac{1}{\kappa}\sinh(\kappa r)\cos(\theta_1) \\
x_2 &= \frac{1}{\kappa}\sinh(\kappa r)\sin(\theta_1)\cos(\theta_2) \\
&\vdots \\
x_{d-1} &= \frac{1}{\kappa}\sinh(\kappa r)\sin(\theta_1)\cdots\sin(\theta_{d-2})\cos(\theta_{d-1}) \\
x_d &= \frac{1}{\kappa}\sinh(\kappa r)\sin(\theta_1)\cdots\sin(\theta_{d-2})\sin(\theta_{d-1})
\end{aligned}$$

for $r \in (0, \infty)$, $\theta_1, \dots, \theta_{d-2} \in (0, \pi)$, and $\theta_{d-1} \in (0, 2\pi)$, with the convention that for $d = 2$ the only angle is $\theta_1 \in (0, 2\pi)$. This parametrizes $\mathbb{H}^d_\kappa$ up to a $\lambda_{\breve{g}}$-null set, namely $\{x \in \mathbb{H}^d_\kappa : x_d = 0,\ x_{d-1} \geq 0\}$.

In these coordinates, the volume element of the Riemannian volume measure $\lambda_{\breve{g}}$ on $\mathbb{H}^d_\kappa$ satisfies

$$\mathrm{d}\lambda_{\breve{g}} = \frac{1}{\kappa^{d-1}}\sinh^{d-1}(\kappa r)\sin^{d-2}(\theta_1)\sin^{d-3}(\theta_2)\cdots\sin(\theta_{d-2})\,\mathrm{d}r\,\mathrm{d}\theta_1\cdots\mathrm{d}\theta_{d-1}.$$

With $\mathrm{d}\omega$ the surface area element of the unit sphere S^{d-1} from Appendix 7.2 this becomes

$$\mathrm{d}\lambda_{\breve{g}} = \frac{1}{\kappa^{d-1}}\sinh^{d-1}(\kappa r)\,\mathrm{d}r\,\mathrm{d}\omega.$$

7.4 Proof of Proposition 35

Proof of Proposition 35. We start with the first part. By scale invariance, we may normalize the threshold to 1. Use spherical polar coordinates around the centre of the ball, and write the radial coordinates of two sampled points as $r, s \in [0, 1]$ and the smaller angle between their radial directions as $\theta \in [0, \pi]$. In the Euclidean ball, the radial distribution function is $F_0(r) = r^d$. In the hyperbolic ball, it is

$$F_\alpha(r) = \frac{\int_0^r \sinh^{d-1}(\alpha u)\,du}{\int_0^1 \sinh^{d-1}(\alpha u)\,du}.$$

Since $u \mapsto \sinh(\alpha u)/u$ is strictly increasing on $(0, \infty)$,

$$F_\alpha(r) = \frac{\int_0^r u^{d-1}\left(\frac{\sinh(\alpha u)}{u}\right)^{d-1} du}{\int_0^1 u^{d-1}\left(\frac{\sinh(\alpha u)}{u}\right)^{d-1} du} \leq \frac{\int_0^r u^{d-1}\,du}{\int_0^1 u^{d-1}\,du} = r^d = F_0(r),$$

with strict inequality for $r \in (0, 1)$ and $\alpha > 0$.

Let $q_0(r, s)$ and $q_\alpha(r, s)$ be the conditional probabilities, over the angle, that two points at radii r, s are at distance at most 1 in the Euclidean and hyperbolic balls, respectively. We claim that $q_\alpha(r, s) \leq q_0(r, s)$. When $r + s \leq 1$, this follows trivially from the triangle inequality. Otherwise, let γ_0 and γ_α be the Euclidean and hyperbolic angles opposite a side of length 1 in triangles with remaining side lengths r, s. Then $d_0 \leq 1$ is equivalent to $\theta \leq \gamma_0$, and $d_\alpha \leq 1$ is

equivalent to $\theta \le \gamma_\alpha$. It remains to compare the two critical angles. Set $\sigma = (r+s+1)/2$. The Euclidean and hyperbolic laws of cosines give

$$\sin^2\frac{\gamma_0}{2} = \frac{(\sigma-r)(\sigma-s)}{rs}, \qquad \sin^2\frac{\gamma_\alpha}{2} = \frac{\sinh(\alpha(\sigma-r))\sinh(\alpha(\sigma-s))}{\sinh(\alpha r)\sinh(\alpha s)}.$$

Since $r, s \le 1$ and $r+s>1$, we have $0<\sigma-r\le s$ and $0<\sigma-s\le r$. Again using that $y\mapsto \sinh(\alpha y)/y$ is strictly increasing, we obtain

$$\frac{\sinh(\alpha(\sigma-r))}{\sigma-r} \le \frac{\sinh(\alpha s)}{s}, \qquad \frac{\sinh(\alpha(\sigma-s))}{\sigma-s} \le \frac{\sinh(\alpha r)}{r}.$$

Multiplying these inequalities yields

$$\sin^2\frac{\gamma_\alpha}{2} \le \sin^2\frac{\gamma_0}{2}.$$

Both angles lie in $[0,\pi]$, so $\gamma_\alpha \le \gamma_0$. Since the distribution of θ is the same in both cases, this proves $q_\alpha(r,s) \le q_0(r,s)$.

Couple the radial variables by independent $U, V \sim \mathrm{U}[0,1]$, setting

$$R_0 = F_0^{-1}(U), \quad S_0 = F_0^{-1}(V), \qquad R_\alpha = F_\alpha^{-1}(U), \quad S_\alpha = F_\alpha^{-1}(V).$$

Since $F_\alpha < F_0$ on $(0,1)$, we have $R_\alpha > R_0$ and $S_\alpha > S_0$ almost surely. Therefore

$$\breve{p}_d(\alpha) = \mathbb{E}[q_\alpha(R_\alpha,S_\alpha)] \le \mathbb{E}[q_0(R_\alpha,S_\alpha)] \le \mathbb{E}[q_0(R_0,S_0)] = \overline{p}_d.$$

The second inequality is strict because $\mathbb{P}(R_0+S_0>1)>0$, and on this event q_0 is strictly decreasing in both arguments while $R_\alpha > R_0$ and $S_\alpha > S_0$ almost surely. Hence $\breve{p}_d(\alpha) < \overline{p}_d$. Using the same line of reasoning one shows that $\breve{p}_d(\alpha') < \breve{p}_d(\alpha)$ when $\alpha' > \alpha$.

For the third claim, write

$$\breve{p}_2(\alpha) = \int_0^1 P_\alpha(x)\big(A_\alpha(x) - B_\alpha(x)\big)\,dx,$$

where

$$P_\alpha(x) \coloneqq \frac{2\alpha\sinh(\alpha x)}{\pi(\cosh\alpha - 1)^2},$$

$$A_\alpha(x) \coloneqq \cosh(\alpha)\arccos\Big(\coth(\alpha)\tanh\Big(\frac{\alpha x}{2}\Big)\Big),$$

and

$$B_\alpha(x) \coloneqq \arctan\left(\frac{\sqrt{\cosh^2(\alpha)-\cosh^2(\alpha x/2)}}{\sinh(\alpha x/2)}\right).$$

We expand each term as $\alpha \to 0^+$. For $P_\alpha(x)$ we obtain

$$P_\alpha(x) = \frac{8x}{\pi\alpha^2} + \frac{4x(x^2-1)}{3\pi} + O(\alpha^2).$$

For $A_\alpha(x)$ we obtain

$$A_\alpha(x) = \arccos\Big(\frac{x}{2}\Big) + \alpha^2 A_2(x) + \alpha^4 A_4(x) + O(\alpha^6),$$

where

$$A_2(x) = \frac{1}{2}\arccos\Big(\frac{x}{2}\Big) - \frac{x\sqrt{4-x^2}}{12},$$

and

$$A_4(x) = \frac{7x^3\sqrt{4-x^2} - 52x\sqrt{4-x^2} + 60\arccos(x/2)}{1440}.$$

For $B_\alpha(x)$ we obtain

$$B_\alpha(x) = \arccos\Big(\frac{x}{2}\Big) + \alpha^2 B_2(x) + \alpha^4 B_4(x) + O(\alpha^6),$$

where

$$B_2(x) = \frac{x\sqrt{4-x^2}}{24},$$

and

$$B_4(x) = -\frac{x(x^2+14)\sqrt{4-x^2}}{2880}.$$

It follows that

$$P_\alpha(x)\big(A_\alpha(x) - B_\alpha(x)\big) = F_0(x) + \alpha^2 F_2(x) + O(\alpha^4),$$

uniformly on $[0,1]$, with

$$F_0(x) = \frac{1}{\pi}\left(4x\arccos\Big(\frac{x}{2}\Big) - x^2\sqrt{4-x^2}\right),$$

and

$$F_2(x) = \frac{16x^3\arccos(x/2) - 8x\arccos(x/2) - 3x^4\sqrt{4-x^2} - 2x^2\sqrt{4-x^2}}{24\pi}.$$

We now integrate these terms. Observe that

$$\int_0^1 F_0(x)\,dx = \overline{p}_2.$$

For the correction term, direct integration gives

$$\int_0^1 F_2(x)\,dx = -\frac{\sqrt{3}}{16\pi},$$

so

$$\breve{p}_2(\alpha) = \overline{p}_2 - \frac{\sqrt{3}}{16\pi}\alpha^2 + O(\alpha^4) \qquad \text{as } \alpha \to 0^+.$$

For the final part, note that $0 \le B_\alpha(x) \le \pi/2$, so

$$0 \le \int_0^1 \frac{2\alpha\sinh(\alpha x)}{\pi(\cosh(\alpha)-1)^2} B_\alpha(x)\,dx \le \frac{1}{\cosh(\alpha)-1} = O(e^{-\alpha}).$$

It remains to analyse the A_α-term. Setting $y = \alpha(1-x)$ gives

$$I_\alpha \coloneqq \int_0^1 \frac{2\alpha\sinh(\alpha x)}{\pi(\cosh(\alpha)-1)^2} A_\alpha(x)\,dx = \int_0^\alpha \frac{2\sinh(\alpha-y)}{\pi(\cosh(\alpha)-1)^2} A_\alpha\Big(1-\frac{y}{\alpha}\Big)\,dy.$$

For fixed $y \ge 0$, as $\alpha \to \infty$,

$$\frac{2\sinh(\alpha-y)\cosh(\alpha)}{\pi(\cosh(\alpha)-1)^2} = \frac{2}{\pi}e^{-y} + O(e^{-\alpha}).$$

Moreover,

$$\coth(\alpha)\tanh\left(\frac{\alpha-y}{2}\right) = 1 - 2e^{y-\alpha} + O(e^{-2\alpha}),$$

and therefore

$$\arccos\left(\coth(\alpha)\tanh\left(\frac{\alpha-y}{2}\right)\right) = 2e^{(y-\alpha)/2} + O(e^{-\alpha}).$$

Consequently,

$$e^{\alpha/2}\frac{2\sinh(\alpha-y)}{\pi(\cosh(\alpha)-1)^2}A_\alpha\Big(1-\frac{y}{\alpha}\Big) \to \frac{4}{\pi}e^{-y/2}.$$

pointwise on $y \geq 0$. By dominated convergence,

$$\lim_{\alpha\to\infty} e^{\alpha/2}I_\alpha = \int_0^\infty \frac{4}{\pi}e^{-y/2}\,dy = \frac{8}{\pi}.$$

Combining this with the estimate for the B_α-term gives

$$\breve{p}_2(\alpha) = \frac{8}{\pi}e^{-\alpha/2} + o(e^{-\alpha/2}) \qquad \text{as } \alpha \to \infty.$$

□