POSET LIMITS AND EXCHANGEABLE RANDOM POSETS

ABSTRACT. We develop a theory of limits of finite posets in close analogy to the recent theory of graph limits. In particular, we study representations of the limits by functions of two variables on a probability space, and connections to exchangeable random infinite posets.

1. Introduction and main results

A deep theory of limit objects of (finite) graphs has in recent years been created by Lovász and Szegedy [19] and Borgs, Chayes, Lovász, Sós and Vesztergombi [6, 7], and further developed in a series of papers by these and other authors. It is shown by Diaconis and Janson [10] that the theory is closely connected with the Aldous–Hoover theory of representations of exchangeable arrays of random variables, further developed and described in detail by Kallenberg [18]; the connection is through exchangeable random infinite graphs. (See also Tao [24] and Austin [2].)

The basic ideas of the graph limit theory extend to other structures too; note that the Aldous–Hoover theory as stated by Kallenberg [18] includes both multi-dimensional arrays (corresponding to hypergraphs) and some different symmetry conditions (or lack thereof). For bipartite graphs and digraphs (i.e., directed graphs), some details are given by Diaconis and Janson [10]. For hypergraphs, an extension is given by Elek and Szegedy [11]; see also [10] (where no details are given) and Tao [24] and Austin [2].

It seems possible that some future version of the theory will be formulated in a general way that includes all these cases as well as others. While waiting for such a theory, it is interesting to study further structures. Brightwell and Georgiou [8] have initiated the study of limits of finite *posets* (i.e., partially ordered sets). In the present paper, we develop this theory further.

The theory for posets can be developed in analogy with the theory for graph limits, but it can also be obtained as a special case of the theory for digraphs. We will in this paper use both views. Our main results are parallel to results for graph limits.

In this paper, all posets (and graphs) are assumed to be non-empty. They are usually finite, but we will sometimes use infinite posets as well. If $(P, <)$ is a poset, we call P its *ground set*; we also say that $(P, <)$ is a poset on P. For simplicity, we often use the same notation for a poset and its ground set when there is no danger of confusion. Sometimes we write $<_P$ for the partial

---

Date: 3 February, 2009; revised 6 April 2011; references updated 3 June 2011.
2000 *Mathematics Subject Classification.* 06A06; 05C99, 60C05.

order and P° for the ground set of a poset P. We let $\mathcal{P}$ denote the set of unlabelled finite posets. (For this and other definitions, see also Sections 2–3 where more details are given.)

We may regard a poset $(P, <)$ as a digraph, with vertex set P and a directed edge $i \to j$ if and only if $i < j$ for all $i, j \in P$. (In particular, the digraph is loopless.) The poset and the digraph determine each other uniquely, so we may identify a poset with the corresponding digraph, but note that not every digraph is a poset. Hence, we can regard $\mathcal{P}$ as a subset of the set $\mathcal{D}$ of unlabelled finite digraphs.

A *poset homomorphism* $Q \to P$ is a map $\varphi : Q^\circ \to P^\circ$ between the ground sets such that $x <_Q y \implies \varphi(x) <_P \varphi(y)$. We say that Q is a *subposet* of P, and write $Q \subseteq P$, if $Q^\circ \subseteq P^\circ$ and $x <_Q y \implies x <_P y$, i.e., if the identity map $Q \to P$ is a poset homomorphism. We say that Q is an *induced subposet* of P if further $x <_Q y \iff x <_P y$ for all $x, y \in Q^\circ$. If P is a poset and A is a subset of its ground set P°, then $P|_A$ denotes the restriction of P to A, i.e., A with the order $<_P$ inherited from P. Thus, Q is an induced subposet of P if and only if Q equals $P|_A$ for some (non-empty) $A \subseteq P^\circ$. Note that these definitions agree with the corresponding definitions for digraphs, so we may identify posets with digraphs as above without problems.

In analogy with the graph case in [19; 6], we define the functional $t(Q, P)$ for finite posets as the proportion of all maps $Q \to P$ that are poset homomorphisms. We similarly also define $t_{\mathrm{inj}}(Q, P)$ as the proportion of all injective maps $Q \to P$ that are poset homomorphisms and $t_{\mathrm{ind}}(Q, P)$ as the proportion of all injective maps $\varphi : Q \to P$ such that $x <_Q y \iff \varphi(x) <_P \varphi(y)$ (i.e., φ is an isomorphism onto an induced subposet of P). (t_{ind} is denoted λ in [8].)

We say that a sequence (P_n) of finite posets with $|P_n| \to \infty$ *converges*, if $t(Q, P_n)$ converges for every finite poset Q. (All unspecified limits in this paper are as $n \to \infty$.) It is easy to see that this is equivalent to convergence of $t_{\mathrm{inj}}(Q, P_n)$ for every Q, or of $t_{\mathrm{ind}}(Q, P_n)$ for every Q. For completeness, we also say that a sequence (P_n) of finite posets with $|P_n| \not\to \infty$ converges if it is eventually constant.

If a sequence of posets converges in this sense, what is its limit? Exactly as for graph limits [19; 6; 10], we may define limit objects in several different, equivalent, ways. One possibility is to define the limit objects as equivalence classes of convergent sequences, where two convergent sequences (P_n) and (P'_n) are defined to be equivalent if the combined sequence $(P_1, P'_1, P_2, P'_2, \dots)$ converges. This is similar to the standard construction of the completion of a metric space using Cauchy sequences. In fact, it is easy to define a metric on $\mathcal{P}$ such that the Cauchy sequences are exactly the convergent sequences, and then the poset limits are exactly the elements of the completion, except the ones that correspond to elements of $\mathcal{P}$. A simple way to construct such a metric is to use one of the embedding in Theorem 3.1 of $\mathcal{P}$ into a compact metric space. Equivalently, and this is the method that we find technically most convenient, we choose one of these embeddings, for example $\hat{\tau}^+ : \mathcal{P} \to [0,1]^{\mathcal{P}^+} = [0,1]^{\mathcal{P}} \times [0,1]$ defined in Section 3, identify $\mathcal{P}$ and its image $\hat{\tau}^+(\mathcal{P})$, and let $\overline{\mathcal{P}}$ be its closure in $[0,1]^{\mathcal{P}^+}$. We define $\mathcal{P}_\infty := \overline{\mathcal{P}} \setminus \mathcal{P}$; this is the set of *poset limits*. Thus $\overline{\mathcal{P}} = \mathcal{P} \cup \mathcal{P}_\infty$ is the set of posets and poset limits; this is the natural setting for convergence of posets to a poset limit.

Note that $\overline{\mathcal{P}}$ is a compact metric space, because $[0,1]^{\mathcal{P}^+}$ is. Further, $\mathcal{P}$ is an open dense subset of $\overline{\mathcal{P}}$, and thus $\mathcal{P}_\infty$ is a closed subset and thus itself a compact metric space. (We should emphasize that we think of $\overline{\mathcal{P}}$ as an abstract completion of $\mathcal{P}$ rather than a specific construction of such a completion. The construction just given should be seen as one convenient construction, but any construction, and any equivalent metric, is valid for our purposes.)

It follows from this construction that the functionals $t(Q,\cdot)$, $t_{\mathrm{inj}}(Q,\cdot)$ and $t_{\mathrm{ind}}(Q,\cdot)$ extends by continuity to $\overline{\mathcal{P}}$ for every $Q \in \mathcal{P}$, and that $P_n \to \Pi \in \mathcal{P}_\infty$ if and only if $|P_n| \to \infty$ and $t(Q,P_n) \to t(Q,\Pi)$ for every $Q \in \mathcal{P}$. As a consequence, a poset limit $\Pi \in \mathcal{P}_\infty$ is determined by $t(Q,\Pi)$, $Q \in \mathcal{P}$.

Just as for graph limits, this construction is convenient for the definition and existence of poset limits, but a more concrete representation is desirable. We will study two such representations, by *kernels* and by *exchangeable random posets*.

For graph limits, Lovász and Szegedy [19] gave an important (non-unique) representation by symmetric functions $W : [0,1]^2 \to [0,1]$ (or, more generally, $W : \mathcal{S}^2 \to [0,1]$ for a probability space $\mathcal{S}$), see also [6; 10]. (These functions are called *graphons* [6].) There is also a (more complicated) version of this for digraphs, see [10] and Section 10 below. A similar construction for poset limits is as follows, cf. [8].

Definition 1.1. An *ordered probability space* $(\mathcal{S},\mathcal{F},\mu,\prec)$ is a probability space $(\mathcal{S},\mathcal{F},\mu)$ equipped with a partial order $\prec$ such that $\{(x,y) : x \prec y\}$ is a measurable subset of $\mathcal{S} \times \mathcal{S}$ (i.e., belongs to the product σ-field $\mathcal{F} \times \mathcal{F}$).

A *(poset) kernel* on an ordered probability space $(\mathcal{S},\mathcal{F},\mu,\prec)$ is a measurable function $W : \mathcal{S} \times \mathcal{S} \to [0,1]$ such that, for $x,y,z \in \mathcal{S}$,

$$W(x,y) > 0 \implies x \prec y, \tag{1.1}$$

$$W(x,y) > 0 \text{ and } W(y,z) > 0 \implies W(x,z) = 1. \tag{1.2}$$

A *strict kernel* is a kernel such that $W(x,y) > 0 \iff x \prec y$.

When convenient, we may omit parts of the notation that are clear from the context and say, e.g., that $\mathcal{S}$ or $(\mathcal{S},\mu)$ is a probability space or an ordered probability space. In this paper, 'kernel' always means poset kernel in the above sense.

Remark 1.2. We may when convenient suppose that the kernel is strict, by replacing the order $\prec$ on $\mathcal{S}$ by $\prec'$ defined by $x \prec' y$ if $W(x,y) > 0$. Note further that by (1.2), a strict kernel $W(x,y)$ is typically determined to be 0 or 1 for many (x, y); it is only when (x, y) forms a gap in the order $\prec'$ that we have the freedom to choose $W(x, y) \in (0, 1)$.

For $n \in \mathbb{N} := \{1, 2, \dots\}$, let $[n] := \{1, \dots, n\}$, and let $[\infty] := \mathbb{N}$. Thus $[n]$ is a set of cardinality n for all $n \in \mathbb{N} \cup \{\infty\}$.

Definition 1.3. Given a kernel W on an ordered probability space $(\mathcal{S}, \mathcal{F}, \mu, \prec)$, we define for every $n \in \mathbb{N} \cup \{\infty\}$ a random poset $P(n, W)$ of cardinality n by taking a sequence $(X_i)_{i=1}^{\infty}$ of i.i.d. points in $\mathcal{S}$ with distribution μ, and independent uniformly distributed random variables $\xi_{ij} \sim U(0, 1)$, $i, j \in \mathbb{N}$, and then defining $P(n, W)$ to be $[n]$ with the partial order $\prec^* = \prec_{P(n,W)}$ defined by: $i \prec^* j$ if and only if $\xi_{ij} < W(X_i, X_j)$. In other words, given (X_i), we define the partial order randomly such that $i \prec^* j$ with probability $W(X_i, X_j)$, with (conditionally) independent choices for different pairs (i, j).

Note that $\prec^*$ really is a partial order because of (1.1), which implies irreflexivity and asymmetry, and (1.2), which implies transitivity. (This is the reason why we have to insist that $W(x, z) = 1$ in (1.2).)

Remark 1.4. We insist in Definition 1.1 that (1.1)–(1.2) hold for *all* x, y, z, and not just a.e.; this will require some technical arguments in proofs in Section 5 to replace a candidate kernel by a kernel that is a.e. equal to it. Note that we can define $P(n, W)$ as above also if W only satisfies (1.1)–(1.2) a.e.; $P(n, W)$ then will be a poset a.s. (We will use this in the proof of Theorem 1.9 below.) Hence it would be possible to use such functions W too, but we find our definition more convenient.

Example 1.5. For any ordered probability space, $W(x, y) = \mathbf{1}[x \prec y]$ is a strict kernel. (We use $\mathbf{1}[\mathcal{E}]$ to denote the indicator function of the event $\mathcal{E}$, which is 1 if $\mathcal{E}$ occurs and 0 otherwise.) In this case $i \prec_{P(n,W)} j \iff X_i \prec X_j$ and we do not need the auxiliary random variables ξ_{ij}. In other words, $P(n, W)$ then is (apart from the labelling) just the subset $\{X_1, \dots, X_n\}$ of $\mathcal{S}$ with the induced order, provided $X_1, \dots, X_n$ are distinct (or, in general, if we regard $\{X_1, \dots, X_n\}$ as a multiset). In this case we use also the notation $P(n, \mathcal{S})$.

Note that every strict kernel with values in $\{0, 1\}$ is of this type by definition. Further, (1.2) implies that every strict kernel on an ordered probability space with a continuous order is of this type.

Example 1.6. Let $\mathcal{S} = \{0, 1\}$ with $\mu\{0\} = \mu\{1\} = 1/2$ and $0 \prec 1$; let further $W(0, 1) = p$ for some given $p \in [0, 1]$, and, as required by (1.1), $W(0, 0) = W(1, 0) = W(1, 1) = 0$. Then $P(n, W)$ consists of a random 'lower' set of roughly half the vertices and a complementary 'upper' set, and $u \prec_{P(n,W)} v$ with probability p for all lower u and upper v (and never otherwise), independently for all pairs (u, v).

Further examples are given below and in Section 9.

One of our main results is the following representation theorem, parallel to the result for graph limits by Lovász and Szegedy [19]. The proofs of this and other theorems in the introduction are given in later sections.

Theorem 1.7. *Every kernel W on an ordered probability space $(\mathcal{S}, \mathcal{F}, \mu, \prec)$ defines a poset limit $\Pi_W \in \mathcal{P}_\infty$ such that the following holds.*

(i) *$P(n, W) \xrightarrow{\text{a.s.}} \Pi_W$ as $n \to \infty$.*

(ii) *For every poset $Q \in \mathcal{P}$,*

$$t(Q, \Pi_W) = t(Q, W) := \int_{\mathcal{S}^{|Q|}} \prod_{ij: i <_Q j} W(x_i, x_j) \, \mathrm{d}\mu(x_1) \dots \, \mathrm{d}\mu(x_{|Q|}). \tag{1.3}$$

Moreover, every poset limit $\Pi \in \mathcal{P}_\infty$ can be represented in this way, i.e., $\Pi = \Pi_W$ for some kernel W on an ordered probability space $(\mathcal{S}, \mathcal{F}, \mu, \prec)$.

Unfortunately, the ordered probability space and the kernel W in Theorem 1.7 are not unique (just as in the corresponding representation of graph limits). We discuss the question of when two kernels represent the same poset limit in Section 7. We note, however, the following important fact.

Theorem 1.8. *If W and W' are kernels on ordered probability spaces both representing the same poset limit $\Pi \in \mathcal{P}_\infty$, then the random posets $P(n, W)$ and $P(n, W')$ have the same distribution.*

We may consequently define the random poset $P(n, \Pi)$, for any poset limit $\Pi \in \mathcal{P}_\infty$, as $P(n, W)$ for any kernel W such that $\Pi_W = \Pi$.

Recall the definition of $t(Q, W)$ in (1.3). It is easy to see that if Q is a finite poset, $n \geq |Q|$, and W is a kernel, then

$$\mathbb{E}\, t_{\mathrm{inj}}(Q, P(n, W)) = t(Q, W); \tag{1.4}$$

indeed, both sides are equal to the probability that any given injective map from Q to $P(n, W)$ is a poset homomorphism, noting that this does not depend on the map or on n. Hence, for every poset limit Π, finite poset Q and finite $n \geq |Q|$,

$$\mathbb{E}\, t_{\mathrm{inj}}(Q, P(n, \Pi)) = t(Q, \Pi). \tag{1.5}$$

If Q is a finite labelled poset with ground set $\subset \mathbb{N}$ we similarly find

$$\mathbb{P}\big(Q \subset P(\infty, \Pi)\big) = t(Q, \Pi). \tag{1.6}$$

This is easily seen to be equivalent to (see (3.8)–(3.9) and (5.16)–(5.17))

$$\mathbb{P}\big(P(\infty, \Pi)|_{[n]} = Q\big) = \mathbb{P}\big(P(n, \Pi) = Q\big) = t_{\mathrm{ind}}(Q, \Pi), \tag{1.7}$$

for every (labelled) poset Q on $[n]$, which describes the distribution of $P(\infty, \Pi)$.

We can use the non-uniqueness of the representation to our advantage by imposing further conditions (normalizations) that may be useful in various situations.

Theorem 1.9. *We may in Theorem 1.7 choose one of the following further conditions and impose it on the representing kernel W:*

(i) *W is a strict kernel.*
(ii) *$(\mathcal{S}, \mathcal{F}, \mu) = [0,1]$ with Lebesgue measure; i.e., W is a kernel on $([0,1], \mathcal{B}, \lambda, \prec)$, where $\mathcal{B}$ is the Borel σ-field, λ is Lebesgue measure and $\prec$ is some (measurable) partial order, not necessarily the standard order.*
(iii) *$(\mathcal{S}, \mathcal{F}, \mu) = [0,1]^2$ with Lebesgue measure, and $(x_1, y_1) \prec (x_2, y_2)$ if and only if $x_1 < x_2$ in the standard order.*

When μ is the Lebesgue measure λ (in one or several dimensions), we take $\mathcal{F}$ to be the Borel σ-field. (We could use the Lebesgue σ-field instead; this would not make any essential difference since a Lebesgue measurable function into $[0,1]$ is a.e. equal to a Borel measurable function.)

We have, however, not yet been able to see whether it always is possible to use $(\mathcal{S}, \mathcal{F}, \mu) = [0,1]$ with Lebesgue measure and the standard order $<$. (This would supersede both (ii) and (iii) in Theorem 1.9, and yield a simplified representation of poset limits.) We state this as an open problem:

Problem 1.10. Can every poset limit be represented by a kernel on $([0,1], \mathcal{B}, \lambda, <)$, with the standard order $<$?

Example 1.11 (Continuation of Example 1.6)**.** Let $\mathcal{S}$ and $W = W(p)$ be as in Example 1.6, and let $Q = \mathcal{S} = \{0,1\}$. Theorem 1.7(ii) then yields

$$t(Q, \Pi_W) = \int_{\mathcal{S}^2} W(x,y) \, \mathrm{d}\mu(x) \, \mathrm{d}\mu(y) = p/4.$$

This shows that different p yield different Π_W. Consequently, $\mathcal{P}_\infty$ is uncountable.

Example 1.12. Let $\mathcal{S}$ be an ordered probability space and let $W_{\mathcal{S}}(x,y) := \mathbf{1}[x <_{\mathcal{S}} y]$ be the strict $\{0,1\}$-valued kernel on $\mathcal{S}$ as in Example 1.5. Let $\Pi_{\mathcal{S}} := \Pi_{W_{\mathcal{S}}}$. For any $Q \in \mathcal{P}$, Theorem 1.7(ii) shows that $t(Q, \Pi_{\mathcal{S}}) = \mathbb{P}(X_i <_{\mathcal{S}} X_j \text{ when } i <_Q j)$ for i.i.d. random points X_i in $\mathcal{S}$. Similarly, by (1.7), if Q is labelled, $t_{\mathrm{ind}}(Q, \Pi_{\mathcal{S}})$ equals the probability that the induced order on $\{X_i\}_1^n$ equals the order on Q.

Every poset limit defined by a $\{0,1\}$-valued kernel is, by Remark 1.2, of the type $\Pi_{\mathcal{S}}$. Conversely, every kernel representing a poset limit $\Pi_{\mathcal{S}}$ is a.e. $\{0,1\}$-valued; this follows (by ignoring directions) from the corresponding result for graph limits, see [15, Section 10].

Example 1.13. Let P be a finite poset. Take $\mathcal{S} = P$ and let the probability measure μ be the uniform distribution on P: $\mu\{i\} = |P|^{-1}$ for every $i \in P$. Then P becomes an ordered probability space, and, as a special case of Example 1.12, $W_P(x,y) := \mathbf{1}[x <_P y]$ is a strict kernel on P, defining a poset limit $\Pi_P := \Pi_{W_P}$. Example 1.12 shows that, for any $Q \in \mathcal{P}$, $t(Q, \Pi_P) = \mathbb{P}(x_i <_P x_j \text{ when } i <_Q j)$ for i.i.d. random vertices x_i in $\mathcal{P}$, which is just the probability that the random mapping $i \mapsto x_i$ is a poset homomorphism. Thus,

$$t(Q, \Pi_P) = t(Q, P) \tag{1.8}$$

for all $Q \in \mathcal{P}$.

We have shown that for every finite poset P there is a poset limit $\Pi_P \in \mathcal{P}_\infty$ such that (1.8) holds for all $Q \in \mathcal{P}$. Note that this defines Π_P uniquely. However, as in the graph case, the map $P \mapsto \Pi_P$ is not injective; blowing up a poset P by replacing each element by k incomparable ones and extending the order in the natural way gives a poset P_k with $t(Q, P_k) = t(Q, P)$ for all Q. (Formally, P_k has ground set $P \times [k]$ and $(x, i) <_{P_k} (y, j)$ if $x <_P y$.) Note also that the mapping is not surjective, since $\mathcal{P}$ is countable and $\mathcal{P}_\infty$ is uncountable, see Example 1.11. Hence only some (exceptionally simple) poset limits can be represented as Π_P for a finite poset P.

We can now state a convergence criterion in terms of the cut metric $\delta_\square$ defined in Section 6. (See [6] for the graph version.)

Theorem 1.14. *Let (P_n) be a sequence of finite posets with $|P_n| \to \infty$ and let $\Pi \in \mathcal{P}_\infty$. Let W_{P_n} be the kernel defined by P_n as in Example 1.13, and let W be any kernel that represents Π. Then, as $n \to \infty$, $P_n \to \Pi \iff \delta_\square(W_{P_n}, W) \to 0$.*

There is also a similar result for a sequence of poset limits.

Theorem 1.15. *Let W and $W_1, W_2, \dots$ be kernels on ordered probability spaces $\mathcal{S}, \mathcal{S}_1, \mathcal{S}_2, \dots$. Then, as $n \to \infty$, $\Pi_{W_n} \to \Pi_W \iff \delta_\square(W_n, W) \to 0$.*

Our second representation of poset limits uses exchangeable random posets.

Definition 1.16. A random infinite poset (or digraph) on $\mathbb{N}$ is *exchangeable* if its distribution is invariant under every permutation of $\mathbb{N}$.

Similarly, an array $\{I_{ij}\}_{i,j=1}^\infty$, of random variables is *(jointly) exchangeable* if the array $\{I_{\sigma(i)\sigma(j)}\}_{i,j=1}^\infty$ has the same distribution as $\{I_{ij}\}_{i,j=1}^\infty$ for every permutation σ of $\mathbb{N}$.

Consequently, if R is a random poset on $\mathbb{N}$ and $I_{ij} := \mathbf{1}[i <_R j]$, then R is exchangeable if and only if the array $\{I_{ij}\}$ is.

The random poset $P(\infty, W)$ defined in Definition 1.3 is evidently exchangeable, and thus so is $P(\infty, \Pi)$ in Theorem 1.8. More generally, we can construct exchangeable random infinite posets by taking mixtures of such distributions, i.e., by taking $P(\infty, W)$ or $P(\infty, \Pi)$ with a random kernel W or a random poset limit Π (which of course is assumed to be independent of the other random variables X_i and ξ_{ij} in the construction); cf. the classical de Finetti's theorem for exchangeable sequences of random variables, see e.g. Kallenberg [18, Theorem 1.1]. Another of our main results is that this yields all exchangeable random infinite posets, which can be seen as a de Finetti theorem for posets. (It is a special case of the general representation theorem for exchangeable arrays by Aldous and Hoover [1; 13; 18]. Cf. the graph case in [10].) Moreover, the poset limits correspond to exchangeable random infinite posets whose distribution is an extreme point in the set of all such distributions, and this yields a unique representation of poset limits as follows.

Theorem 1.17. (i) *There is a one-to-one correspondence between distributions of random elements $\Pi \in \mathcal{P}_\infty$ and distributions of exchangeable random infinite posets R on $\mathbb{N}$ given by $R \overset{\mathrm{d}}{=} P(\infty, \Pi)$; this relation between Π and R is equivalent to either of*

$$\mathbb{E}\, t(Q, \Pi) = \mathbb{P}(R \supset Q) \tag{1.9}$$

or

$$\mathbb{E}\, t_{\mathrm{ind}}(Q, \Pi) = \mathbb{P}(R|_A = Q) \tag{1.10}$$

for every finite labelled poset Q with a ground set $A \subset \mathbb{N}$. Furthermore, then $R|_{[n]} \xrightarrow{\mathrm{d}} \Pi$ in $\overline{\mathcal{P}}$ as $n \to \infty$.

(ii) *There is a one-to-one correspondence between poset limits $\Pi \in \mathcal{P}_\infty$ and extreme points of the set of distributions of exchangeable random infinite posets R. This correspondence is given by $R \overset{\mathrm{d}}{=} P(\infty, \Pi)$, or, equivalently, either of*

$$t(Q, \Pi) = \mathbb{P}(R \supset Q) \tag{1.11}$$

or

$$t_{\mathrm{ind}}(Q, \Pi) = \mathbb{P}(R|_A = Q) \tag{1.12}$$

for every finite labelled poset Q with a ground set $A \subset \mathbb{N}$. Furthermore, then $R|_{[n]} \xrightarrow{\mathrm{a.s.}} \Pi$ in $\overline{\mathcal{P}}$ as $n \to \infty$.

We can characterize these extreme point distributions of exchangeable random infinite posets as follows. Let $\mathcal{P}^L_{\mathbb{N}}$ be the space of all (labelled) infinite posets on $\mathbb{N}$. This can be seen as a subset of the product space $\{0,1\}^{\mathbb{N}\times\mathbb{N}}$, using indicators I_{ij} as above. We equip this product space with the product topology, which is compact and metric; then $\mathcal{P}^L_{\mathbb{N}}$ is a closed subset and thus itself a compact metric space.

Theorem 1.18. *Let R be an exchangeable random infinite poset. Then the following are equivalent.*

(i) *The distribution of R is an extreme point in the set of exchangeable distributions in the space $\mathcal{P}^L_{\mathbb{N}}$ of all labelled infinite posets on $\mathbb{N}$.*

(ii) *If Q_1 and Q_2 are two finite posets with disjoint ground sets contained in $\mathbb{N}$, then*

$$\mathbb{P}(R \supset Q_1 \cup Q_2) = \mathbb{P}(R \supset Q_1)\, \mathbb{P}(R \supset Q_2).$$

(Here $Q_1 \cup Q_2$ denotes the poset with ground set $Q_1^\circ \cup Q_2^\circ$ and $x <_Q y \iff x <_{Q_1} y$ or $x <_{Q_2} y$; in particular $x \not<_Q y$ if $x \in Q_1^\circ$ and $y \in Q_2^\circ$ or conversely.)

(iii) *The restrictions $R|_{[k]}$ and $R|_{[k+1,\infty)}$ are independent for every k.*

(iv) *Let $\mathcal{F}_n$ be the σ-field generated by $R|_{[n,\infty)}$. Then the tail σ-field $\bigcap_{n=1}^\infty \mathcal{F}_n$ is trivial, i.e., contains only events with probability 0 or 1.*

There is also a more direct relation between poset limits and exchangeable random infinite posets, without going through kernels and $P(\infty,\Pi)$. Poset limits are limits of unlabelled finite posets. For labelled finite posets there is another, more elementary, notion of a limit as an infinite poset. More precisely, if P_n is a labelled poset on the ground set $[N_n]$ for some finite N_n with $N_n \to \infty$ as $n \to \infty$, and R is a poset on $\mathbb{N}$, we say that $P_n \to R$ if, for every pair $(i,j) \in \mathbb{N}^2$, $\mathbf{1}[i <_{P_n} j] \to \mathbf{1}[i <_R j]$, i.e., if $i <_R j$ then $i <_{P_n} j$ for all large n and if $i \not<_R j$ then $i \not<_{P_n} j$ for all large n. Equivalently, we may regard each P_n as an element of the space $\mathcal{P}^L_{\mathbb{N}}$ of posets on $\mathbb{N}$ by adding an infinite number of points incomparable to everything else (in fact, any extension to $\mathbb{N}$ would do, but it seems natural to choose the trivial one); then $P_n \to R$ in this sense just means convergence in $\mathcal{P}^L_{\mathbb{N}}$ with the topology just introduced. For unlabelled posets, we can always choose a labelling. Of course, the choice of labelling may affect the result, so we choose a random labelling. Thus, if P is a finite unlabelled poset, we let $\widehat{P}$ be the labelled poset obtained by randomly labelling P by $1,\dots,|P|$, with the same probability $1/|P|!$ for each possible labelling. As above, we can also extend $\widehat{P}$ to a random poset on $\mathbb{N}$, which we by abuse of notation still denote by $\widehat{P}$. The appropriate limit for a sequence (P_n) of finite unlabelled posets then is the limit in distribution of $(\widehat{P}_n)$ as random elements of the compact metric space $\mathcal{P}^L_{\mathbb{N}}$. This turns out to be equivalent to convergence of the unlabelled posets P_n as defined above (and in Definition 3.2 below), i.e., in $\overline{\mathcal{P}}$.

Theorem 1.19. *Let (P_n) be a sequence of finite unlabelled posets and assume that $|P_n| \to \infty$. Then the following are equivalent.*

(i) $P_n \to \Pi$ *in* $\overline{\mathcal{P}}$ *for some* $\Pi \in \overline{\mathcal{P}}$.
(ii) $\widehat{P}_n \xrightarrow{\mathrm{d}} R$ *in* $\mathcal{P}^L_{\mathbb{N}}$ *for some random* $R \in \mathcal{P}^L_{\mathbb{N}}$.

If these hold, then $\Pi \in \mathcal{P}_\infty$, R is exchangeable, and $R \overset{\mathrm{d}}{=} P(\infty,\Pi)$; consequently, (1.11) *and* (1.12) *hold for every finite labelled poset Q with a ground set $A \subset \mathbb{N}$.*

Finally, we note that by regarding posets as digraphs, we obtain an embedding $\mathcal{P} \subset \mathcal{D}$ which extends to an embedding $\overline{\mathcal{P}} \subset \overline{\mathcal{D}}$ with $\mathcal{P}_\infty \subset \mathcal{D}_\infty$. The poset limits can thus be seen as special digraph limits. We characterize the digraph limits that are poset limits in several ways in Theorem 10.1.

Sections 2–3 contain definitions and some basic properties of poset limits. The theorems above are proven in Sections 4–5 and 8. The cut metric is defined and studied in Section 6; in particular we show that it makes the set of all kernels into a compact metric space, which by Theorem 1.15 is homeomorphic to the space of poset limits $\mathcal{P}_\infty$. The (lack of) uniqueness of the representation by kernels is discussed in Section 7, where conditions for equivalence are given. Further examples are given in Section 9. The relation between poset limits and digraph limits is discussed further in Section 10.

2. Preliminaries

We consider both labelled and unlabelled posets and digraphs. We use for convenience $[n]$ as our standard ground set for labelled posets and vertex set for labelled digraphs, i.e., we use the labels $1, 2, \dots$.

A *digraph* (directed graph) G consists of a vertex set $V(G)$ and an edge set $E(G) \subseteq V(G) \times V(G)$; the edge indicators thus form an arbitrary zero–one matrix $\{X_{ij}\}$, $i, j \in V(G)$. We let $|G|$ denote the number of vertices. Unless we state otherwise explicitly, we assume that $1 \leq |G| < \infty$, but we will also sometimes consider infinite digraphs.

For $n \in \mathbb{N}$, let $\mathcal{D}_n^L$ be the set of the 2^{n^2} labelled digraphs with vertex set $[n]$ and let $\mathcal{D}_n$ be the set of unlabelled digraphs with n vertices; $\mathcal{D}_n$ can formally be defined as the quotient set $\mathcal{D}_n^L / \cong$ modulo isomorphisms. Further, let $\mathcal{D}^L := \bigcup_{n \geq 1} \mathcal{D}_n^L$ and $\mathcal{D} := \bigcup_{n \geq 1} \mathcal{D}_n$; thus $\mathcal{D}$ is the set of finite unlabelled digraphs.

Similarly, let $\mathcal{P}_n^L$ be the set of all posets with ground set $[n]$ and let $\mathcal{P}_n$ be the quotient set $\mathcal{P}_n^L / \cong$ of unlabelled posets with n vertices, and let $\mathcal{P}^L := \bigcup_{n \geq 1} \mathcal{P}_n^L$ and $\mathcal{P} := \bigcup_{n \geq 1} \mathcal{P}_n$, the set of finite unlabelled posets.

As said above, we can regard every poset as a digraph. This works for both labelled and unlabelled posets and yields the inclusions $\mathcal{P}_n^L \subset \mathcal{D}_n^L$, $\mathcal{P}_n \subset \mathcal{D}_n$, $\mathcal{P}^L \subset \mathcal{D}^L$, $\mathcal{P} \subset \mathcal{D}$. Further, every labelled poset or digraph can be regarded as an unlabelled one by ignoring the labels. Hence it often does not matter whether the posets and digraphs are labelled or not, but we shall be explicit the times it does matter.

We can characterize the digraphs that are posets using a few special digraphs. For $n \geq 1$, let C_n be the directed cycle with n vertices and n edges, and let P_n be the directed path with $n+1$ vertices and n edges. (Thus C_1 is a loop and C_2 a double edge.) We regard these as unlabelled digraphs. Note that, except for P_1, the digraphs P_n and C_n are *not* posets. Moreover, it is easy to check the following characterization.

Lemma 2.1. *A (finite or infinite) digraph G is a poset if and only if it does not have any induced subgraph C_1, C_2, C_3, or P_2.* □

3. Digraph and poset limits

We repeat some of the notation and results for digraphs in [10] and give corresponding results for posets.

If G is an (unlabelled) digraph and $v_1, \dots, v_k$ is a sequence of vertices in G, then $G(v_1, \dots, v_k)$ denotes the labelled digraph with vertex set $[k]$ where we put an edge $i \to j$ if $v_i \to v_j$ in G. We allow the possibility that $v_i = v_j$ for some i and j.

We let $G[k]$, for $k \geq 1$, be the random digraph $G(v_1, \dots, v_k)$ obtained by sampling $v_1, \dots, v_k$ uniformly at random among the vertices of G, with replacement. In other words, $v_1, \dots, v_k$ are independent uniformly distributed random vertices of G.

For $k \leq |G|$, we further let $G[k]'$ be the random digraph $G(v_1', \dots, v_k')$ where we sample $v_1', \dots, v_k'$ uniformly at random without replacement; the sequence $v_1', \dots, v_k'$ is thus a uniformly distributed random sequence of k distinct vertices. Hence, $G[k]'$ is the induced subgraph on a random set of k vertices, with the vertices relabelled $1, \dots, k$.

For a finite poset P, we similarly define $P(v_1, \dots, v_k)$, $P[k]$ and $P[k]'$ (the latter if $k \leq |P|$); these are posets with ground set $[k]$, and $P[k]$ and $P[k]'$ are random. Note that these definitions are consistent with our identification of posets and (certain) digraphs: for example, $P[k]$ is the same for the poset P as for P regarded as a digraph.

The graph limit theory in [19] and subsequent papers is based on the study of the functional $t(F, G)$ which is defined for two graphs F and G as the proportion of all mappings $V(F) \to V(G)$ that are graph homomorphisms $F \to G$. In probabilistic terms, $t(F, G)$ is the probability that a uniform random mapping $V(F) \to V(G)$ is a graph homomorphism. For the digraph version, see [10], $\varphi : V(F) \to V(G)$ is a homomorphism if $i \to j$ in F implies $\varphi(i) \to \varphi(j)$ in G. Thus, using the notation just introduced and assuming that F is labelled and $k = |F|$, we can write the definition as

$$t(F, G) := \mathbb{P}\big(F \subseteq G[k]\big). \tag{3.1}$$

Note that both F and $G[k]$ are digraphs on the same vertex set $[k]$, so the relation $F \subseteq G[k]$ simply means $E(F) \subseteq E(G[k])$. Again following [19] (and the notation of [6] and [10]), with $k = |F|$ as in (3.1), let

$$t_{\mathrm{inj}}(F, G) := \mathbb{P}\big(F \subseteq G[k]'\big), \tag{3.2}$$

be the proportion of injective maps $V(F) \to V(G)$ that are graph homomorphisms, and

$$t_{\mathrm{ind}}(F, G) := \mathbb{P}\big(F = G[k]'\big), \tag{3.3}$$

provided F and G are digraphs with $|F| \leq |G|$. If $|F| > |G|$ we set $t_{\mathrm{inj}}(F, G) := t_{\mathrm{ind}}(F, G) := 0$. Note that although the relations $F \subseteq G[k]$, $F \subseteq G[k]'$ and $F = G[k]'$ may depend on the labelling of F, the probabilities in (3.1)–(3.3) do not, by symmetry, so $t(F, G)$, $t_{\mathrm{inj}}(F, G)$ and $t_{\mathrm{ind}}(F, G)$ are well defined for unlabelled F and G (by choosing any labellings).

The definitions (3.1)–(3.3) can be used for finite posets too. Thus, if P and Q are finite (unlabelled) posets, then $t(P, Q)$, $t_{\mathrm{inj}}(P, Q)$ and $t_{\mathrm{ind}}(P, Q)$ are defined as numbers in $[0, 1]$. Note that these numbers are the same as if we regard P and Q as digraphs; we will therefore use the same notation for the poset case as for the digraph case.

The basic definition of Lovász and Szegedy [19] and Borgs, Chayes, Lovász, Sós and Vesztergombi [6] is that a sequence (G_n) of graphs converges if $t(F, G_n)$ converges for every graph F. As in [10], we modify this by requiring also that $|G_n|$ converges to some finite or infinite limit. We let, as in [10], $\mathcal{D}^+$ be the union of $\mathcal{D}$ and some one-point set $\{*\}$ and define the mappings

$\tau, \tau_{\text{inj}}, \tau_{\text{ind}} : \mathcal{D} \to [0,1]^{\mathcal{D}}$ and $\tau^+ : \mathcal{D} \to [0,1]^{\mathcal{D}^+} = [0,1]^{\mathcal{D}} \times [0,1]$ by

$$\tau(G) := (t(F,G))_{F\in\mathcal{D}} \in [0,1]^{\mathcal{D}}, \tag{3.4}$$

$$\tau_{\text{inj}}(G) := (t_{\text{inj}}(F,G))_{F\in\mathcal{D}} \in [0,1]^{\mathcal{D}}, \tag{3.5}$$

$$\tau_{\text{ind}}(G) := (t_{\text{ind}}(F,G))_{F\in\mathcal{D}} \in [0,1]^{\mathcal{D}}, \tag{3.6}$$

$$\tau^+(G) := \bigl(\tau(G),\, |G|^{-1}\bigr) \in [0,1]^{\mathcal{D}^+}. \tag{3.7}$$

For posets we similarly define $\mathcal{P}^+ := \mathcal{P} \cup \{*\}$ and the mappings $\hat\tau, \hat\tau_{\text{inj}}, \hat\tau_{\text{ind}} : \mathcal{P} \to [0,1]^{\mathcal{P}}$ and $\hat\tau^+ : \mathcal{P} \to [0,1]^{\mathcal{P}^+} = [0,1]^{\mathcal{P}} \times [0,1]$ by considering F in $\mathcal{P}$ only; these mappings can thus be obtained from $\tau, \tau_{\text{inj}}, \tau_{\text{ind}}, \tau^+$ by a projection selecting some coordinates only.

The mappings τ and $\hat\tau$ are not injective on $\mathcal{P}$. Indeed, if P is a poset and P_k is obtained by blowing-up P as in Example 1.13, then P and P_k have the same images under τ and $\hat\tau$ for all $k \in \mathbb{N}$. (This is the poset version of an example in [19] and [6].) However, it is easy to see that $\tau^+, \tau_{\text{inj}}$ and τ_{ind} are injective on $\mathcal{D}$, cf. [10], and, similarly, that $\hat\tau^+, \hat\tau_{\text{inj}}$ and $\hat\tau_{\text{ind}}$ are injective on $\mathcal{P}$, (This uses the special definitions of $\tau_{\text{inj}}(F,G)$ and $\tau_{\text{ind}}(F,G)$ when $|F| > |G|$.)

Although the mappings $\hat\tau^+, \hat\tau_{\text{inj}}, \hat\tau_{\text{ind}}$ contain only part of the information in $\tau^+, \tau_{\text{inj}}, \tau_{\text{ind}}$, the injectivity of them shows that they in fact contain all possible information. This is also seen in the following stronger result concerning limits.

Theorem 3.1. *Suppose that P_n is a sequence of finite posets. Then the following conditions are equivalent.*

(i) *$\hat\tau^+(P_n)$ converges in $[0,1]^{\mathcal{P}^+}$, i.e. $t(Q,P_n)$ converges for every poset $Q \in \mathcal{P}$ and $|P_n|$ converges to some limit in $\mathbb{N} \cup \{\infty\}$.*

(ii) *$\hat\tau_{\text{inj}}(P_n)$ converges in $[0,1]^{\mathcal{P}}$, i.e. $t_{\text{inj}}(Q,P_n)$ converges for every poset $Q \in \mathcal{P}$.*

(iii) *$\hat\tau_{\text{ind}}(P_n)$ converges in $[0,1]^{\mathcal{P}}$, i.e. $t_{\text{ind}}(Q,P_n)$ converges for every poset $Q \in \mathcal{P}$.*

(iv) *$\tau^+(P_n)$ converges in $[0,1]^{\mathcal{D}^+}$, i.e. $t(F,P_n)$ converges for every digraph $F \in \mathcal{D}$ and $|P_n|$ converges to some limit in $\mathbb{N} \cup \{\infty\}$.*

(v) *$\tau_{\text{inj}}(P_n)$ converges in $[0,1]^{\mathcal{D}}$, i.e. $t_{\text{inj}}(F,P_n)$ converges for every digraph $F \in \mathcal{D}$.*

(vi) *$\tau_{\text{ind}}(P_n)$ converges in $[0,1]^{\mathcal{D}}$, i.e. $t_{\text{ind}}(F,P_n)$ converges for every digraph $F \in \mathcal{D}$.*

Proof. It is easily seen that each of the conditions implies that $|P_n|$ converges to a limit in $\mathbb{N} \cup \{\infty\}$. Further, if $|P_n|$ converges to a finite limit, each of the statements implies that $P_n = P$ for all sufficiently large n and some (unlabelled) poset $P \in \mathcal{P}$.

It thus suffices to consider the case $|P_n| \to \infty$. In this case, for every $F \in \mathcal{D}$, $t(F,P_n) - t_{\text{inj}}(F,P_n) = O(|F|^2/|P_n|) \to 0$, see [19] and [10], and thus (i) $\iff$ (ii) and (iv) $\iff$ (v).

Further, see [6], [19] or [10] for the easy details, one can go between the two families $\{t_{\mathrm{inj}}(F,\cdot)\}_{F\in\mathcal{D}}$ and $\{t_{\mathrm{ind}}(F,\cdot)\}_{F\in\mathcal{D}}$ of functionals on $\mathcal{D}$ by summation and inclusion-exclusion, and for posets a similar argument holds for the families $\{t_{\mathrm{inj}}(F,\cdot)\}_{F\in\mathcal{P}}$ and $\{t_{\mathrm{ind}}(F,\cdot)\}_{F\in\mathcal{P}}$; hence it follows that (ii) $\iff$ (iii) and (v) $\iff$ (vi).

Finally, (iii) $\iff$ (vi) because $t_{\mathrm{ind}}(F,P_n)=0$ for every digraph F that is not a poset. $\square$

Definition 3.2. A sequence (P_n) of finite posets *converges* if one, and thus all, of the conditions in Theorem 3.1 holds.

Remark 3.3. As seen in the proof of Theorem 3.1, the case when $|P_n|\not\to\infty$ is not very interesting since then (P_n) converges if and only if the sequence is eventually constant. The interesting case is thus $|P_n|\to\infty$, and then convergence of (P_n) is also equivalent to convergence of $\hat\tau(P_n)$ in $[0,1]^{\mathcal{P}}$ or $\tau(P_n)$ in $[0,1]^{\mathcal{D}}$.

Since $\hat\tau^+$ is injective, we can identify $\mathcal{P}$ with its image $\hat\tau^+(\mathcal{P})\subseteq[0,1]^{\mathcal{P}^+}$ and define $\overline{\mathcal{P}}\subseteq[0,1]^{\mathcal{P}^+}$ as its closure. Alternatively, we can consider $\hat\tau_{\mathrm{inj}}$ or $\hat\tau_{\mathrm{ind}}$; we can again identify $\mathcal{P}$ with its image and consider its closure $\overline{\mathcal{P}}$ in $[0,1]^{\mathcal{P}}$. It follows from Theorem 3.1 that the three compactifications $\overline{\hat\tau^+(\mathcal{P})}$, $\overline{\hat\tau_{\mathrm{inj}}(\mathcal{P})}$, $\overline{\hat\tau_{\mathrm{ind}}(\mathcal{P})}$ are homeomorphic and we can use any of them for $\overline{\mathcal{P}}$. Moreover, we can also, again by Theorem 3.1, use τ^+, τ_{inj} or τ_{ind} and embed $\mathcal{P}$ in $[0,1]^{\mathcal{D}^+}$ or $[0,1]^{\mathcal{D}}$ and obtain $\overline{\mathcal{P}}$ as a compact subset of $[0,1]^{\mathcal{D}^+}$ or $[0,1]^{\mathcal{D}}$. This is equivalent to regarding posets as digraphs and using the embeddings $\mathcal{P}\subset\mathcal{D}\subset\overline{\mathcal{D}}$ and defining $\overline{\mathcal{P}}$ as the closure of $\mathcal{P}$ in $\overline{\mathcal{D}}$. (Thus $\overline{\mathcal{P}}$ can be regarded as a subset of $\overline{\mathcal{D}}$.) Since all these constructions yield homeomorphic results it does not matter which one we use. Note that $\overline{\mathcal{P}}$ is a compact metric space. Different, equivalent, metrics are given by the embeddings above into $[0,1]^{\mathcal{P}^+}$, $[0,1]^{\mathcal{P}}$, $[0,1]^{\mathcal{D}^+}$, $[0,1]^{\mathcal{D}}$.

We let $\mathcal{P}_\infty:=\overline{\mathcal{P}}\setminus\mathcal{P}$; this is the set of all limit objects of sequences (P_n) in $\mathcal{P}$ with $|P_n|\to\infty$; i.e., $\mathcal{P}_\infty$ is the set of all poset limits.

For every fixed digraph F, the functions $t(F,\cdot)$, $t_{\mathrm{inj}}(F,\cdot)$ and $t_{\mathrm{ind}}(F,\cdot)$ have unique continuous extensions to $\overline{\mathcal{D}}$, for which we use the same notation. (Indeed, this is built into our constructions; for example, if we use τ^+ to define $\overline{\mathcal{D}}$ then $t(F,\cdot)$ is simply one of the coordinates, and so continuous by definition of the product topology.) In particular, $t(Q,\Pi)$ is defined for every finite poset Q and poset limit Π. We similarly extend $|\cdot|^{-1}$ continuously to $\overline{\mathcal{D}}$ by defining $|\Gamma|=\infty$ and thus $|\Gamma|^{-1}=0$ for $\Gamma\in\mathcal{D}_\infty:=\overline{\mathcal{D}}\setminus\mathcal{D}$. It is easily seen that

$$t_{\mathrm{inj}}(F,\Gamma)=t(F,\Gamma) \tag{3.8}$$

for every $F\in\mathcal{D}$ and $\Gamma\in\mathcal{D}_\infty$ [10]; in particular for $F\in\mathcal{P}$ and $\Gamma\in\mathcal{P}_\infty$. Moreover, for any $Q,P\in\mathcal{P}$, $t_{\mathrm{inj}}(Q,P)=\sum_{Q'\supseteq Q}t_{\mathrm{ind}}(Q',P)$, where we sum over all posets $Q' \supseteq Q$ with the same ground set Q°, and thus by continuity

$$t_{\mathrm{inj}}(Q,\Pi) = \sum_{Q' \supseteq Q} t_{\mathrm{ind}}(Q',\Pi) \tag{3.9}$$

for every $Q \in P$ and $\Pi \in \overline{\mathcal{P}}$.

Note that $\mathcal{P}_\infty = \overline{\mathcal{P}} \cap \mathcal{D}_\infty = \{\Pi \in \overline{\mathcal{P}} : |\Pi|^{-1} = 0\}$, which shows that $\mathcal{P}_\infty$ is a closed and thus compact subset of $\overline{\mathcal{P}}$. Conversely, $\mathcal{P}$ is an open subset of $\overline{\mathcal{P}}$; by Remark 3.3, it has the discrete topology. Note further that $\mathcal{P}$ is countable while $\overline{\mathcal{P}}$ and $\mathcal{P}_\infty$ are uncountable, e.g. by Example 1.11.

We summarize the results above on convergence.

Theorem 3.4. *A sequence (P_n) of finite posets converges in the sense of Definition 3.2 if and only if it converges in the compact metric space $\overline{\mathcal{P}}$.* □

The construction of $\mathcal{P}$ further immediately implies the following related characterization of convergence in $\mathcal{P}_\infty$.

Theorem 3.5. *A sequence Π_n of poset limits (i.e., elements of $\mathcal{P}_\infty$) converges [to a poset limit Π] if and only if $t(Q,\Pi_n)$ converges [to $t(Q,\Pi)$] for every finite poset Q.*

We can here replace t by t_{inj} or t_{ind}; further, we may let Q range over all finite digraphs instead of posets. □

4. Exchangeable random infinite posets

It is straightforward to verify that Sections 3–5 of Diaconis and Janson [10] hold with only notational changes for the poset case as well as for the graph case treated there. Rather than repeating the details, we therefore omit them and refer to [10], giving only a few comments. We first obtain the following basic result on convergence in distribution of *random* unlabelled posets, corresponding to [10, Theorem 3.1].

Theorem 4.1. *Let P_n, $n \ge 1$, be random unlabelled posets and assume that $|P_n| \xrightarrow{\mathrm{p}} \infty$. The following are equivalent, as $n \to \infty$.*

(i) *$P_n \xrightarrow{\mathrm{d}} \Pi$ for some random $\Pi \in \overline{\mathcal{P}}$.*
(ii) *For every finite family $Q_1, \dots, Q_m$ of (non-random) finite posets, the random variables $t(Q_1,P_n), \dots, t(Q_m,P_n)$ converge jointly in distribution.*
(iii) *For every (non-random) $Q \in \mathcal{P}$, the random variables $t(Q,P_n)$ converge in distribution.*
(iv) *For every (non-random) $Q \in \mathcal{P}$, the expectations $\mathbb{E}\, t(Q,P_n)$ converge.*

If these properties hold, then the limits in (ii), (iii) *and* (iv) *are $\bigl(t(Q_i,\Pi)\bigr)_{i=1}^m$, $t(Q,\Pi)$ and $\mathbb{E}\, t(Q,\Pi)$, respectively; conversely, if* (ii), (iii) *or* (iv) *holds with these limits for some random $\Pi \in \overline{\mathcal{P}}$, then* (i) *holds with the same Π. Furthermore, $\Pi \in \mathcal{P}_\infty$ a.s.*

The same results hold if t is replaced by t_{inj} or t_{ind}. □

Using this we then obtain Theorem 1.19, which corresponds to [10, Theorems 4.1 and 5.2]; note that (1.12) is the poset version of a formula in [10, Theorem 4.1], which follows because, if n is so large that $A \subseteq [n]$, $\mathbb{P}(\widehat{P_n}|_A = Q) = t_{\mathrm{ind}}(Q, P_n) \to t_{\mathrm{ind}}(Q, \Pi)$, and that (1.11) easily follows from (1.12) by summing over $Q' \supseteq Q$ on the same ground set. We really cannot prove the equality $R \overset{\mathrm{d}}{=} P(\infty, \Pi)$ yet, since we have defined $P(\infty, \Pi)$ using kernels and Theorem 1.7, which is not yet proven. Instead, we note only that $R \overset{\mathrm{d}}{=} P(\infty, \Pi)$ will follow by (1.11) and (1.6) or (1.12) and (1.7) once we have proven Theorem 1.7 and thus verified (1.6) and (1.7) in Section 5. (Alternatively, we could have used (1.6) and (1.7) as a definition of $P(\infty, \Pi)$.)

Remark 4.2. Actually, [10, Theorems 4.1] is stated more generally for sequences of random graphs, and similarly Theorem 1.19 extends to the case of random finite posets P_n with $|P_n| \xrightarrow{\mathrm{p}} \infty$; then the limit $\Pi \in \overline{\mathcal{P}}$ is in general random too, and (i) becomes $P_n \xrightarrow{\mathrm{d}} \Pi$ while (1.11) and (1.12) have to be replaced by (1.9) and (1.10).

We then obtain Theorem 1.17, which corresponds to [10, Theorems 5.3 and Corollary 5.4], and Theorem 1.18, which corresponds to [10, Theorems 5.5]. The a.s. convergence of $R|_{[n]} = P(\infty, \Pi)|_{[n]}$ in Theorem 1.17(ii) follows, as in [10, Remark 5.1], because $t_{\mathrm{inj}}(Q, R|_{[n]})$, $n \geq |Q|$, is a reverse martingale for every $Q \in \mathcal{P}$.

Remark 4.3. Brightwell and Georgiou [8] define the *continuum limit* of a sequence of random finite posets P_n to be an (atomless) ordered probability space $\mathcal{S}$ such that, in our notation, $\mathbb{E}\, t_{\mathrm{ind}}(Q, P_n) \to t_{\mathrm{ind}}(Q, \Pi_{\mathcal{S}})$ for every $Q \in \mathcal{P}$, with $\Pi_{\mathcal{S}}$ as in Example 1.12. By Theorem 4.1, this is equivalent to $P_n \xrightarrow{\mathrm{p}} \Pi_{\mathcal{S}}$ (with a non-random limit).

Brightwell and Georgiou [8] further define a sequence of random finite posets to be *compatible* if Theorem 4.1(iv) holds (with t_{ind}, which is equivalent), and ask whether every compatible sequence has a limit. Theorem 4.1 shows that this holds if we allow the limit to be a random poset limit. (But not always with a non-random limit, and not always with a limit of the form $\Pi_{\mathcal{S}}$, since given any random Π, we can construct $P_n = P(n, \Pi)$ satisfying Theorem 4.1(i).)

5. Kernels

Proof of Theorem 1.7. First, let W be a kernel on an ordered probability space. Then $R = P(\infty, W)$ defined in Definition 1.3 is an exchangeable random infinite poset, which satisfies the independence condition Theorem 1.18(ii); hence, by Theorem 1.18(i) its distribution is an extreme point in the set of exchangeable distributions, and by Theorem 1.17(ii) there exists a poset limit Π (which we call Π_W) such that (1.11) and (1.12) hold, and $P(n, W) = R|_{[n]} \xrightarrow{\mathrm{a.s.}} \Pi = \Pi_W$ in $\overline{\mathcal{P}}$. This proves (i). Further, it follows directly from the definition of $P(\infty, W)$ that if Q is a finite labelled poset, then

$$\mathbb{P}\big(P(\infty,W) \supset Q\big) = \int_{\mathcal{S}^{|Q|}} \prod_{ij:i<_Q j} W(x_i,x_j)\,\mathrm{d}\mu(x_1)\dots\,\mathrm{d}\mu(x_{|Q|}),$$

and thus (ii) follows by (1.11).

For the converse, suppose that $\Pi \in \mathcal{P}_\infty$, and consider the corresponding exchangeable random infinite poset R given by Theorem 1.17. (I.e., $P(\infty,\Pi)$, although we have not yet shown this, so we have to use only (1.11) and (1.12) until Theorem 1.7 is proven.) Let $I_{ij} := \mathbf{1}[i <_R j]$, $i, j \in \mathbb{N}$. Then (I_{ij}) is a jointly exchangeable random arrays of zero–one variables, with the diagonal entries $I_{ii} = 0$. For such exchangeable random arrays, the Aldous–Hoover representation theorem takes the form, see Kallenberg [18, Theorem 7.22],

$$I_{ij} = f(\xi_\emptyset, \xi_i, \xi_j, \xi_{ij}), \qquad i \neq j, \tag{5.1}$$

where $f : [0,1]^4 \to \{0,1\}$ is a (Borel) measurable function, $\xi_{ji} = \xi_{ij}$, and $\xi_\emptyset$, ξ_i $(1 \le i)$ and ξ_{ij} $(1 \le i < j)$ are independent random variables uniformly distributed on $[0,1]$. By Theorem 1.17(ii), the distribution of the array (I_{ij}) is an extreme point in the set of exchangeable distributions, and thus by Theorem 1.18 and [18, Lemma 7.35], there exists such a representation where f does not depend on $\xi_\emptyset$, so (5.1) becomes $I_{ij} = f(\xi_i, \xi_j, \xi_{ij})$, $i \neq j$. We then further define

$$W_0(x,y) := \mathbb{P}\big(f(x,y,\xi) = 1\big) = \mathbb{E}\, f(x,y,\xi), \tag{5.2}$$

where $\xi \sim U(0,1)$. (In general, the variable $\xi_\emptyset$ can be interpreted as making W random; this is needed if we consider a random Π as in Theorem 1.17(i), but not in the present case.)

As our ordered probability space we take $[0,1]$ with Lebesgue measure, with an order to be defined later. The function W_0 is almost the sought kernel, but not quite. The problem is that the function f, and thus W_0, can be arbitrarily changed on a null set without affecting the distribution of (I_{ij}); consequently we can only show properties such as (1.2) a.e. for W_0. We thus have to make a suitable choice of W among all functions that are a.e. equal to W_0.

Recall that a point (x,y) is a *Lebesgue point* of an integrable function F on $\mathbb{R}^2$ if

$$(2\varepsilon)^{-2} \iint\limits_{|x'-x|<\varepsilon,\,|y'-y|<\varepsilon} |F(x',y') - F(x,y)|\,\mathrm{d}x'\,\mathrm{d}y' \to 0 \qquad \text{as } \varepsilon \to 0, \tag{5.3}$$

and that a.e. point is a Lebesgue point of F, see e.g. Stein [23, §1.8]. This applies trivially to functions defined on $(0,1)^2$ too, by extending the functions to $\mathbb{R}^2$ by defining them as 0 outside $(0,1)^2$. We modify the function

W_0 in two steps. We first define

$$W_1(x,y) := \liminf_{\varepsilon\to 0}(2\varepsilon)^{-2} \iint\limits_{|x'-x|<\varepsilon,\,|y'-y|<\varepsilon} W_0(x',y')\,\mathrm{d}x'\,\mathrm{d}y', \tag{5.4}$$

and note that $W_1 = W_0$ at every Lebesgue point of W_0 and thus a.e. Next, we let E be the set of all Lebesgue points of W_1 in $(0,1)^2$ and define $W(x,y) := W_1(x,y)\mathbf{1}[(x,y)\in E]$. Then $0 \le W(x,y) \le 1$ and $W = W_1 = W_0$ a.e. Moreover, if $W(x,y) > 0$, then $W_1(x,y) = W(x,y)$, $(x,y)\in(0,1)^2$ and (x,y) is a Lebesgue point of W_1; hence, using $W_1(x,y) = W(x,y)$ and $W_1 = W$ a.e., (x,y) is a Lebesgue point of W. Finally, if $(x,y)\in(0,1)^2$ and

$$(2\varepsilon)^{-2} \iint\limits_{|x'-x|<\varepsilon,\,|y'-y|<\varepsilon} W(x',y')\,\mathrm{d}x'\,\mathrm{d}y' \to 1 \tag{5.5}$$

as $\varepsilon\to 0$, then $W_1(x,y) = 1$ by (5.4); thus, using $W_1 \le 1$, (5.5) implies that (x,y) is a Lebesgue point of W_1, and hence $(x,y)\in E$ and $W(x,y) = W_1(x,y) = 1$.

After these preliminaries, note that $I_{12} = I_{23} = 1$ implies $I_{13} = 1$ since R is a poset. Hence, using (5.1) and (5.2), and the independence of $\{\xi_i,\xi_{jk}\}$,

$$\begin{aligned}
0 &= \mathbb{P}(I_{12} = I_{23} = 1,\ I_{13} = 0)\\
&= \mathbb{E}\big(f(\xi_1,\xi_2,\xi_{12})f(\xi_2,\xi_3,\xi_{23})(1 - f(\xi_1,\xi_3,\xi_{13}))\big)\\
&= \mathbb{E}\big(W_0(\xi_1,\xi_2)W_0(\xi_2,\xi_3)(1 - W_0(\xi_1,\xi_3))\big)\\
&= \mathbb{E}\big(W(\xi_1,\xi_2)W(\xi_2,\xi_3)(1 - W(\xi_1,\xi_3))\big);
\end{aligned}$$

thus

$$W(x_1,x_2)W(x_2,x_3)(1 - W(x_1,x_3)) = 0 \quad \text{a.e.} \tag{5.6}$$

Similarly, since R does not contain a directed cycle $1 <_R 2 <_R 3 <_R 1$, $P(I_{12} = I_{23} = I_{31} = 1) = 0$ and

$$W(x_1,x_2)W(x_2,x_3)W(x_3,x_1) = 0 \quad \text{a.e.} \tag{5.7}$$

Now assume that x, y and z are such that $W(x,y) > 0$ and $W(y,z) > 0$. Let $\varepsilon > 0$ and let X_x^ε be a random, uniformly distributed, point in $(x-\varepsilon, x+\varepsilon)$ and let similarly X_y^ε and X_z^ε be random points in $(y-\varepsilon, y+\varepsilon)$ and $(z-\varepsilon, z+\varepsilon)$; these three variables being independent. Since $W(x,y) > 0$, (x,y) is a Lebesgue point of W, and thus (5.3) shows that $\mathbb{E}\,|W(X_x^\varepsilon, X_y^\varepsilon) - W(x,y)| \to 0$ as $\varepsilon\to 0$. In particular, using Markov's inequality, $\mathbb{P}(W(X_x^\varepsilon, X_y^\varepsilon) = 0) \to 0$ as $\varepsilon \to 0$. Similarly, $\mathbb{P}(W(X_y^\varepsilon, X_z^\varepsilon) = 0) \to 0$ as $\varepsilon\to 0$. On the other hand, (5.6) implies $W(X_x^\varepsilon, X_y^\varepsilon)W(X_y^\varepsilon, X_z^\varepsilon)(1 - W(X_x^\varepsilon, X_z^\varepsilon)) = 0$ a.s., and thus

$$\mathbb{P}\big(W(X_x^\varepsilon, X_z^\varepsilon) < 1\big) \le \mathbb{P}(W(X_x^\varepsilon, X_y^\varepsilon) = 0) + \mathbb{P}(W(X_y^\varepsilon, X_z^\varepsilon) = 0) \to 0,$$

as $\varepsilon\to 0$. It follows that (5.5) holds at (x,z), and thus, by the remarks above, $W(x,z) = 1$. Consequently,

$$W(x,y) > 0 \text{ and } W(y,z) > 0 \implies W(x,z) = 1, \tag{5.8}$$

which is (1.2).

Similarly, still assuming $W(x,y) > 0$ and $W(y,z) > 0$, (5.7) implies $W(X^\varepsilon_x, X^\varepsilon_y)W(X^\varepsilon_y, X^\varepsilon_z)W(X^\varepsilon_z, X^\varepsilon_x) = 0$ a.s., and thus

$$\mathbb{P}\bigl(W(X^\varepsilon_z, X^\varepsilon_x) > 0\bigr) \le \mathbb{P}(W(X^\varepsilon_x, X^\varepsilon_y) = 0) + \mathbb{P}(W(X^\varepsilon_y, X^\varepsilon_z) = 0) \to 0, \tag{5.9}$$

as $\varepsilon \to 0$. If further $W(z,x) > 0$, then (z,x) is a Lebesgue point of W and $\mathbb{P}(W(X^\varepsilon_z, X^\varepsilon_x) = 0) \to 0$ as $\varepsilon \to 0$, which contradicts (5.9). Consequently,

$$W(x,y) > 0 \text{ and } W(y,z) > 0 \implies W(z,x) = 0. \tag{5.10}$$

Now suppose that $W(x,x) > 0$ for some x. Taking $y = z = x$ in (5.10) we find $W(x,x) = 0$, a contradiction. Hence, $W(x,x) = 0$ for every x. Further, if both $W(x,y) > 0$ and $W(y,x) > 0$ for some x and y, then (5.8) yields $W(x,x) = 1$, which was just shown to be impossible. Hence $W(x,y)W(y,x) = 0$ for all x and y. These properties and (5.8) show that we may define a partial order $\prec$ on $\mathcal{S} = [0,1]$ by $x \prec y$ if $W(x,y) > 0$, and then W is a (strict) kernel on the ordered probability space $([0,1], \mathcal{B}, \lambda, \prec)$, where λ is the Lebesgue measure. (We took f Borel measurable, and then W_0, W_1, E and W are Borel measurable too.)

Finally, it follows from (5.1) and (5.2) that for every finite poset Q with ground set $A \subset \mathbb{N}$,

$$\begin{aligned}
\mathbb{P}(Q \subset R) &= \mathbb{P}\Bigl(\prod_{ij:i<_Q j} I_{ij} = 1\Bigr) = \mathbb{E}\prod_{ij:i<_Q j} I_{ij} = \mathbb{E}\prod_{ij:i<_Q j} f(\xi_i, \xi_j, \xi_{ij}) \\
&= \mathbb{E}\prod_{ij:i<_Q j} W_0(\xi_i, \xi_j) = \mathbb{E}\prod_{ij:i<_Q j} W(\xi_i, \xi_j) \\
&= \int_{\mathcal{S}^{|Q|}} \prod_{ij:i<_Q j} W(x_i, x_j)\,\mathrm{d}\mu(x_1)\dots\,\mathrm{d}\mu(x_{|Q|}).
\end{aligned} \tag{5.11}$$

This equals $t(Q,W)$ defined in (1.3). Hence, by (1.11) and (1.3), $t(Q,\Pi) = \mathbb{P}(Q \subset R) = t(Q, \Pi_W)$ for all such posets Q, and thus $\Pi_W = \Pi$. □

Remark 5.1. Remember that we have $\xi_{ij} = \xi_{ji}$, which in principle may give a dependence between ij and ji terms. This is an important complication in other situations, for example for digraphs [10], but is of no concern for posets, where at most one of $W(\xi_i, \xi_j)$ and $W(\xi_j, \xi_i)$ is non-zero, and similarly, in (5.11), at most one of $i <_Q j$ and $j <_Q i$ holds.

As remarked above, it now follows that $R \overset{\mathrm{d}}{=} P(\infty, \Pi)$ in Theorems 1.17 and 1.19, for example by (1.7) and (1.12), or directly by (5.11).

Remark 5.2. Alternatively, we can regard R as an exchangeable random infinite digraph, and use the representation by a quintuple of functions $\mathbf{W} = (W_{00}, W_{01}, W_{10}, W_{11}, w)$ as in Diaconis and Janson [10, Theorem 9.1], see Section 10; here $W_{\alpha\beta} : [0,1]^2 \to [0,1]$ and $w : [0,1] \to [0,1]$. The function w generates loops and W_{11} generates doubly directed edges (i.e., cycles C_2); hence $w = 0$ and $W_{11} = 0$ in the poset case. Further, $W_{01}(x,y) = W_{10}(y,x)$

and $\sum_{\alpha,\beta=0}^{1} W_{\alpha\beta}(x,y)=1$, so the quintuple $\mathbf{W}$ is determined by W_{10}. We then can replace (5.2) by $W_0 := W_{10}$, and complete the proof by adjusting W_0 on a null set as above.

Lemma 5.3. *Let $(\mathcal{S},\mathcal{F},\mu,\prec)$ be an ordered probability space, and let $g(x) := \mu\{z\in\mathcal{S} : z\prec x\}$. Then, the set $\{(x,y)\in\mathcal{S}^2 : x\prec y$ and $g(x)\geq g(y)\}$ is a null set in $\mathcal{S}^2$.*

Proof. For $x\in\mathcal{S}$, let $D_x := \{z : z\prec x\}$ and $E_x := \{z : z\prec x$ and $g(z)\geq g(x)\}$. If $z\in E_x$, then $z\prec x$ and thus $D_z\subseteq D_x$, and $\mu(D_z)=g(z)\geq g(x)=\mu(D_x)$; hence $g(z)=g(x)$ and $\mu(D_x\setminus D_z)=0$. In particular, then $\mu(E_x\setminus D_z)=0$, because $E_x\subseteq D_x$.

For two points $y,z\in E_x$, at least one of $y\notin D_z$ and $z\notin D_y$ holds, and thus by symmetry

$$\mu(E_x)^2\leq 2\int_{E_x}\int_{E_x}\mathbf{1}[y\notin D_z]\,\mathrm{d}\mu(y)\,\mathrm{d}\mu(z)=2\int_{E_x}\mu(E_x\setminus D_z)\,\mathrm{d}\mu(z)=0.$$

Hence, $\mu(E_x)=0$ for every x, and thus

$$\mu\times\mu\{(z,y)\in\mathcal{S}^2 : z\prec y \text{ and } g(z)\geq g(y)\}=\int_{\mathcal{S}}\mu(E_y)\,\mathrm{d}\mu(y)=0. \qquad \square$$

Proof of Theorem 1.9. The proof of Theorem 1.7 above gives a kernel satisfying (i) and (ii).

For (iii), we start with a kernel W_1 on an ordered probability space $([0,1],\mathcal{B},\lambda,\prec)$ as in (ii); thus $\prec$ is some partial order on $[0,1]$, in general different from the standard order $<$. Define $g(x) := \lambda\{z\in[0,1] : z\prec x\}$. Then $x\preceq y \implies g(x)\leq g(y)$. Moreover, by Lemma 5.3, for a.e. (x,y), $x\prec y\implies g(x)<g(y)$.

Let $W_2(x,y) := W_1(x,y)\mathbf{1}[g(x)<g(y)]$; this too is a kernel on $([0,1],\mathcal{B},\lambda,\prec)$. Since $W_1(x,y)>0 \implies x\prec y \implies g(x)<g(y)$ for a.e. (x,y) by Lemma 5.3, we have $W_2=W_1$ a.e., and thus $\Pi_{W_2}=\Pi_{W_1}=\Pi$. Moreover,

$$W_2(x,y)>0 \implies g(x)<g(y). \tag{5.12}$$

Let $U_1,U_2\sim U(0,1)$ be independent uniform random variables. Let $\xi := g(U_1)$, and let $h:[0,1]\to[0,1]$ be the right-continuous inverse of its distribution function $s\mapsto\mathbb{P}(g(U_1)\leq s)$; then h is a non-decreasing function such that $h(U_1)\overset{\mathrm{d}}{=}g(U_1)=\xi$.

By the transfer theorem [17, Theorem 6.10] with $\xi := g(U_1)$, $\eta := U_1$, $\tilde\xi := h(U_1)\overset{\mathrm{d}}{=}\xi$, there exists a measurable function $f:[0,1]^2\to[0,1]$ such that if $\tilde\eta := f(\tilde\xi,U_2)$, then $(\tilde\xi,\tilde\eta)\overset{\mathrm{d}}{=}(\xi,\eta)=(g(U_1),U_1)$. This implies $\tilde\xi-g(\tilde\eta)\overset{\mathrm{d}}{=}\xi-g(\eta)=0$ and thus $\tilde\xi=g(\tilde\eta)$ a.s., i.e.

$$h(U_1)=g(\tilde\eta)=g(f(\tilde\xi,U_2))=g(f(h(U_1),U_2)) \qquad \text{a.s.;}$$

hence,

$$h(x_1)=g(f(h(x_1),x_2)) \qquad \text{a.e. on } [0,1]^2. \tag{5.13}$$

Let $\mathcal{S} := [0,1]^2$ with Lebesgue measure, and define the functions $W_3, W_4 : \mathcal{S}^2 \to [0,1]$ by

$$W_3\big((x_1,x_2),(y_1,y_2)\big) := W_2\big(f(h(x_1),x_2), f(h(y_1),y_2)\big) \tag{5.14}$$

and

$$W_4\big((x_1,x_2),(y_1,y_2)\big) := W_3\big((x_1,x_2),(y_1,y_2)\big)\mathbf{1}[x_1 < y_1]. \tag{5.15}$$

Then W_4 is a kernel on $(\mathcal{S},\mathcal{F},\mu,\prec)$. Further, if $W_3\big((x_1,x_2),(y_1,y_2)\big) > 0$, then (5.14) and (5.12) yield $g(f(h(x_1),x_2)) < g(f(h(y_1),y_2))$, which by (5.13) implies, except on a null set in $\mathcal{S}^2 = [0,1]^4$, that $h(x_1) < h(y_1)$ and thus, since h is non-decreasing, $x_1 < y_1$. Consequently, $W_3 = W_4$ a.e. on $\mathcal{S}^2 = [0,1]^4$.

Since $f(h(U_1),U_2) = f(\tilde\xi,U_2) = \tilde\eta \overset{\mathrm{d}}{=} \eta = U_1$ is uniformly distributed on [0,1], it follows from the construction of $P(n,W)$ that $P(n,W_4) \overset{\mathrm{d}}{=} P(n,W_3) \overset{\mathrm{d}}{=} P(n,W_2)$ for every $n \le \infty$. Thus by Theorem 1.7(i) or 1.17(ii), $\Pi_{W_4} = \Pi_{W_2} = \Pi$ and W_4 is a kernel on $(\mathcal{S},\mathcal{F},\mu,\prec)$ that represents Π. □

Proof of Theorem 1.8. Let W be a kernel with $\Pi_W = \Pi$. By (1.4), (1.3) and (3.8), for every finite poset Q and $n \ge |Q|$,

$$\mathbb{E}\, t_{\mathrm{inj}}(Q, P(n,W))) = t(Q,\Pi_W) = t(Q,\Pi) = t_{\mathrm{inj}}(Q,\Pi). \tag{5.16}$$

By (3.9), it follows that for any finite n and labelled poset Q on $[n]$,

$$\mathbb{P}(Q = P(n,W)) = \mathbb{E}\, t_{\mathrm{ind}}(Q, P(n,W))) = t_{\mathrm{ind}}(Q,\Pi). \tag{5.17}$$

Hence, the distribution of $P(n,W)$ is determined by Π for finite n, and does not depend on the choice of W. Further, the distribution of $P(\infty,W)$ is determined by the distribution of $P(n,W) = P(\infty,W)|_{[n]}$, $1 \le n < \infty$, so this distribution too is determined by Π. □

Moreover, (1.5)–(1.7) follow from (5.16)–(5.17).

6. Cut norm and metric

In this section it will be convenient to (usually) ignore orders and study general probability spaces.

Let $(\mathcal{S},\mu)$ be a probability space. We define the *cut norm* $\|W\|_\square$ of $W \in L^1(\mathcal{S}^2)$ by, see [12; 6; 3],

$$\|W\|_{\square,1} := \sup_{S,T}\Big|\int_{S\times T} W(x,y)\,\mathrm{d}\mu(x)\,\mathrm{d}\mu(y)\Big|, \tag{6.1}$$

where the supremum is taken over all pairs of measurable subsets of $\mathcal{S}$. Alternatively, one can take

$$\|W\|_{\square,2} := \sup_{\|f\|_\infty,\|g\|_\infty\le 1}\Big|\int_{\mathcal{S}^2} f(x)W(x,y)g(y)\,\mathrm{d}\mu(x)\,\mathrm{d}\mu(y)\Big|. \tag{6.2}$$

It is easily seen that $\|W\|_{\square,1} \le \|W\|_{\square,2} \le 4\|W\|_{\square,1}$; thus the two norms $\|\cdot\|_{\square,1}$ and $\|\cdot\|_{\square,2}$ are equivalent. It will for our purposes not be important which one we use, and we shall write $\|\cdot\|_\square$ for either norm. (There are further, equivalent versions of the cut norm; see [6].) Note that for either definition of the cut norm we have $\left|\int W\right| \le \|W\|_\square \le \|W\|_{L^1}$.

If W is a function defined on $\mathcal{S}^2$ for some space $\mathcal{S}$, and $\varphi : \mathcal{S}' \to \mathcal{S}$ is a function, we define the function W^φ on $\mathcal{S}'^2$ by

$$W^\varphi(x, y) = W\bigl(\varphi(x), \varphi(y)\bigr). \tag{6.3}$$

Given two integrable functions $W_j : \mathcal{S}_j^2 \to \mathbb{R}$, $j = 1, 2$, where $(\mathcal{S}_1, \mu_1)$ and $(\mathcal{S}_2, \mu_2)$ are two, in general different, probability spaces, we define the *cut metric* [6] by

$$\delta_\square(W_1, W_2) = \inf_{\varphi_1, \varphi_2} \|W_1^{\varphi_1} - W_2^{\varphi_2}\|_\square, \tag{6.4}$$

taking the infimum over all pairs (φ_1, φ_2) of measure preserving maps $\varphi_1 : \mathcal{S} \to \mathcal{S}_1$ and $\varphi_2 : \mathcal{S} \to \mathcal{S}_2$ defined on a common probability space $(\mathcal{S}, \mu)$; in other words, we take the infimum over all couplings of the distributions μ_1 and μ_2. (See further [6], [4] and [15] where some equivalent versions are given and discussed, and Lemma 6.2 below.) Note that

$$\delta_\square(W, W^\varphi) = 0 \tag{6.5}$$

for every W and every measure preserving φ; this is the point of using $\delta_\square$. Clearly, $0 \le \delta_\square(W_1, W_2) < \infty$ and $\delta_\square(W_1, W_2) = \delta_\square(W_2, W_1)$. The triangle inequality holds too, so $\delta_\square$ is a semimetric; this is not completely obvious but we refer to [14] or [15] for a detailed proof.

We let our kernels, and in this section more general functions, be defined on arbitrary probability spaces. Sometimes it is convenient to use the special space $([0, 1], \mathcal{B}, \lambda)$. (For simplicity we write often $[0, 1]$ instead of $([0, 1], \mathcal{B}, \lambda)$. Thus, $[0, 1]$ is assumed to be equipped with Lebesgue measure unless we state otherwise.) The next lemma shows that this can be done without loss of generality. Again, we refer to [14] or [15] (in the symmetric case, which is not simpler) for a detailed proof.

Lemma 6.1. *If $W \in L^1(\mathcal{S}^2)$ for some probability space $\mathcal{S}$, then there exists a function $W' \in L^1([0, 1]^2)$ with $\delta_\square(W, W') = 0$.* □

For such functions, the cut norm has a simpler descriptions.

Lemma 6.2. *If $W_1, W_2 \in L^1([0, 1]^2)$, then*

$$\delta_\square(W_1, W_2) = \inf_\varphi \|W_1 - W_2^\varphi\|_\square,$$

taking the infimum over all measure preserving bimeasurable bijections $\varphi : [0, 1] \to [0, 1]$. □

For the proof we refer to [6] (for the symmetric case; the proof is the same), [15] (the symmetric case) or [14].

The triangle inequality implies that the relation $W \cong W'$ if $\delta_\square(W, W') = 0$ defines an equivalence relation between functions W, possibly defined for different probability spaces. We let, for a probability space $\mathcal{S}$, $\mathcal{W}(\mathcal{S})$ be the set of all measurable $W : \mathcal{S}^2 \to [0, 1]$, and let $\overline{\mathcal{W}}$ be the quotient space of $\bigcup_{\mathcal{S}} \mathcal{W}(\mathcal{S})$ modulo $\cong$. (The careful reader might correctly object that the collection of all probability spaces is not a set, so $\bigcup_{\mathcal{S}}$ is not defined. However, Lemma 6.1 implies that it actually suffices to consider $\mathcal{W}([0,1])$ modulo $\cong$, or the union for $\mathcal{S}$ in any set of probability spaces containing $[0,1]$.)

By construction, $(\overline{\mathcal{W}}, \delta_\square)$ is a metric space. The following important result is a minor variation of the symmetric version in Lovász and Szegedy [20]; the proof is essentially the same as in the symmetric case (see [14] for details).

Theorem 6.3. *$(\overline{\mathcal{W}}, \delta_\square)$ is a compact metric space.* □

We use the construction in Definition 1.3 for an arbitrary $W \in \mathcal{W}(\mathcal{S})$; in general, $\prec^*$ will not be a partial order so $P(n, W)$ will not be a poset, but we can always regard $P(n, W)$ as a random digraph (with $i \prec^* j$ interpreted as a directed edge ij). We further extend the definition in (1.3) and define

$$t(F, W) := \int_{\mathcal{S}^{|F|}} \prod_{ij \in F} W(x_i, x_j)\, \mathrm{d}\mu(x_1) \dots \,\mathrm{d}\mu(x_{|F|}) \tag{6.6}$$

for every $W \in \mathcal{W}(\mathcal{S})$ and every finite digraph F. Equivalently,

$$t(F, W) = \mathbb{P}(i \prec^* j \text{ for every edge } ij \text{ in } F), \tag{6.7}$$

where $\prec^*$ is the relation in $P(\infty, W)$.

We say that a digraph is *simple* if it can be obtained by orienting a simple graph; in other words, a digraph is simple if it has no loops or double edges (i.e., no induced C_1 or C_2). In particular, a poset is a simple digraph.

Lemma 6.4. *Let $W_1 \in \mathcal{W}(\mathcal{S}_1)$ and $W_2 \in \mathcal{W}(\mathcal{S}_2)$ where $\mathcal{S}_1$ and $\mathcal{S}_2$ are probability spaces. Then, for every simple finite digraph F, if m is the number of edges in F, then*

$$\bigl|t(F, W_1) - t(F, W_2)\bigr| \le m\delta_\square(W_1, W_2).$$

In particular, for every finite poset Q (with m the number of pairs (i, j) with $i <_Q j$),

$$\bigl|t(Q, \Pi_{W_1}) - t(Q, \Pi_{W_2})\bigr| \le m\delta_\square(W_1, W_2).$$

Proof. This is identical to the proof in the symmetric case (when F is a finite undirected graph) given in [6] (with an unimportant extra factor in the constant); see also [4, Lemma 2.2] for a nice formulation (with the constant given above). □

Note that we exclude digraphs F with a loop or a double edge (an induced C_1 or C_2) since we do not want factors of the type $W(x_i, x_i)$ or $W(x_i, x_j)W(x_j, x_i)$ in the integrals. (In fact, Lemma 6.4 fails for $F = \mathsf{C}_1$ or C_2.)

We now focus on functions $W \in \mathcal{W}(\mathcal{S})$ that are poset kernels (recall Definition 1.1. We define three special digraphs $\mathsf{D}_1, \mathsf{D}_2, \mathsf{D}_3$ with vertex sets $\{1, 2, 3\}$ and edge sets $E(\mathsf{D}_1) = \{12, 23\}$, $E(\mathsf{D}_2) = \{12, 23, 13\}$ and $E(\mathsf{D}_3) = \{12, 23, 31\}$. (Thus D_2 is a poset, but not D_1 and D_3, and $\mathsf{D}_3 = \mathsf{C}_3$.)

Lemma 6.5. *Let $W \in \mathcal{W}([0,1])$. Then the following are equivalent.*

(i) *For every finite n, $P(n, W)$ is a.s. a poset.*
(ii) *$P(\infty, W)$ is a.s. a poset.*
(iii) *There exists a partial order $\prec$ on $[0,1]$ and a poset kernel W' on $([0,1], \mathcal{B}, \lambda, \prec)$ such that $W = W'$ a.e.*
(iv) *$t(\mathsf{D}_1, W) = t(\mathsf{D}_2, W)$ and $t(\mathsf{D}_3, W) = 0$.*

Proof. (i) $\iff$ (ii) is obvious because $P(n, W) = P(\infty, W)|_{[n]}$.

(iii) $\implies$ (i),(ii) is clear since $P(n, W) = P(n, W')$ a.s.

(ii) $\implies$ (iii). If (ii) holds, then $R := P(\infty, W)$ is an exchangeable random infinite poset. We follow the proof of Theorem 1.7 in Section 5, noting that by Definition 1.3, $I_{ij} := \mathbf{1}[\xi_{ij} < W(X_i, X_j)]$ so we already have the representation (5.1) (with $\xi_i = X_i$), and (5.2) yields $W_0(x, y) := \mathbb{P}\big(\xi < W(x, y)\big) = W(x, y)$. The remainder of the proof of Theorem 1.7 shows that we may modify W_0 on a null set such that the result (denoted W there and W' here) is a poset kernel on $([0,1], \mathcal{B}, \lambda, \prec)$ for some partial order $\prec$ on $[0,1]$.

(iv) $\implies$ (iii). We have

$$0 = t(\mathsf{D}_1, W) - t(\mathsf{D}_2, W) = \int_{[0,1]^3} W(x_1, x_2) W(x_2, x_3) \big(1 - W(x_1, x_3)\big) \, \mathrm{d}x_1 \, \mathrm{d}x_2 \, \mathrm{d}x_3$$

and

$$0 = t(\mathsf{D}_3, W) = \int_{[0,1]^3} W(x_1, x_2) W(x_2, x_3) W(x_3, x_1) \, \mathrm{d}x_1 \, \mathrm{d}x_2 \, \mathrm{d}x_3.$$

Thus, (5.6) and (5.7) in the proof of Theorem 1.7 hold. The proof of Theorem 1.7 actually used the assumption that $R = P(\infty, W)$ is a poset only to show (5.6) and (5.7); hence we may argue exactly as for (ii) $\implies$ (iii).

(ii) $\implies$ (iv). By (6.7),

$$t(\mathsf{D}_3, W) = \mathbb{P}(1 \prec^* 2,\ 2 \prec^* 3,\ 3 \prec^* 1)$$

and

$$t(\mathsf{D}_1, W) - t(\mathsf{D}_2, W) = \mathbb{P}(1 \prec^* 2,\ 2 \prec^* 3,\ 1 \not\prec^* 3),$$

and both are 0 if $P(\infty, W)$ a.s. is a poset. $\square$

Remark 6.6. The implications (i) $\iff$ (ii) $\implies$ (iv) hold for $W \in \mathcal{W}(\mathcal{S})$ for any probability space $\mathcal{S}$. We do not know whether that is true for the other implications, or whether there might be measure theoretic complications.

We prove a poset kernel version of Lemma 6.1.

Lemma 6.7. *If W is a poset kernel on an ordered probability space $(\mathcal{S}, \mathcal{F}, \mu, \prec)$, then there exists a poset kernel W' on $([0,1], \mathcal{B}, \lambda, \prec)$, for some partial order $\prec$ on $[0,1]$, such that $\delta_\square(W, W') = 0$.*

Proof. By Lemma 6.1, there exists $W_1 \in \mathcal{W}([0,1])$ such that $\delta_\square(W, W_1) = 0$. If F is any simple finite digraph, then Lemma 6.4 implies $t(F, W) = t(F, W_1)$. Since $P(\infty, W)$ is a random infinite poset, Lemma 6.5 and Remark 6.6

show that $t(\mathsf{D}_1, W_1) = t(\mathsf{D}_1, W) = t(\mathsf{D}_2, W) = t(\mathsf{D}_2, W_1)$ and $t(\mathsf{D}_3, W_1) = t(\mathsf{D}_3, W_1) = 0$, and thus Lemma 6.5 shows the existence of a poset kernel W' with $W' = W_1$ a.e. and thus $\delta_\square(W, W') = \delta_\square(W, W_1) = 0$. ☐

We define $\overline{\mathcal{W}}_{\mathsf{P}}$ as

$$\{W : W \text{ is a poset kernel on some ordered probability space } \mathcal{S}\}, \tag{6.8}$$

or

$$\{W : W \text{ is a poset kernel on } ([0,1], \mathcal{B}, \lambda, \prec) \text{ for some } \prec\}, \tag{6.9}$$

modulo the equivalence relation $\cong$; note that (6.8) and (6.9) are equivalent by Lemma 6.7. Thus $\overline{\mathcal{W}}_{\mathsf{P}}$ is a subset of the metric space $\overline{\mathcal{W}}$, and we equip $\overline{\mathcal{W}}_{\mathsf{P}}$ with the inherited metric $\delta_\square$.

By Lemma 6.4, the functionals $t(F, \cdot)$ are well-defined and continuous on the quotient space $\overline{\mathcal{W}}$.

Lemma 6.8. $\overline{\mathcal{W}}_{\mathsf{P}} = \{\overline{W} \in \overline{\mathcal{W}} : t(\mathsf{D}_1, \overline{W}) = t(\mathsf{D}_2, \overline{W})$ *and* $t(\mathsf{D}_3, \overline{W}) = 0\}$.

Proof. If $\overline{W} \in \overline{\mathcal{W}}$, we may by Lemma 6.1 choose a representative in $\mathcal{W}([0,1])$, and the result then follows by Lemma 6.5. ☐

Theorem 6.9. *The metric space* $(\overline{\mathcal{W}}_{\mathsf{P}}, \delta_\square)$ *is compact.*

Proof. $\overline{\mathcal{W}}_{\mathsf{P}}$ is a closed subset of $\overline{\mathcal{W}}$ by Lemma 6.8 and the fact that the functionals $t(\mathsf{D}_\ell, \cdot)$ are continuous on $\overline{\mathcal{W}}$. Hence the result follows from Theorem 6.3. ☐

7. Equivalence of kernels

Suppose that $(\mathcal{S}_1, \mu_1)$ and $(\mathcal{S}_2, \mu_2)$ are two probability spaces and that $\varphi : \mathcal{S}_1 \to \mathcal{S}_2$ is a measure preserving map. If $\mathcal{S}_2$ is an ordered probability space with order $\prec_2$ and W is a kernel on $\mathcal{S}_2$, then we can define a partial order $\prec_1$ on $\mathcal{S}_1$ by $x \prec_1 y \iff W^\varphi(x,y) > 0$; then $\mathcal{S}_1$ is an ordered probability space, W^φ is a (strict) kernel on $\mathcal{S}_1$, and $\varphi : \mathcal{S}_1 \to \mathcal{S}_2$ is order preserving. Furthermore, in this case, if $(X_i)_{i=1}^\infty$ are i.i.d. points in $\mathcal{S}_1$, then $(\varphi(X_i))_{i=1}^\infty$ are i.i.d. points in $\mathcal{S}_2$, and it follows from Definition 1.3 that

$$P(n, W^\varphi) \overset{\mathrm{d}}{=} P(n, W) \qquad \text{for every } n \le \infty; \tag{7.1}$$

hence Theorem 1.7 implies that the kernels W^φ and W define the same poset limit Π_W. As in the case of graph limits, see [6; 4; 10; 5], this is not quite the only source of non-uniqueness of the representing kernel W, but it is 'almost' so, in a sense made precise below.

A *Borel space* is a measurable space $(\mathcal{S}, \mathcal{F})$ that is isomorphic to a Borel subset of $[0,1]$, see e.g. [17, Appendix A1] and [21]. In fact, a Borel space is either isomorphic to $([0,1], \mathcal{B})$ or it is countable infinite or finite. Moreover, every Borel subset of a Polish topological space (with the Borel σ-field) is a Borel space. A *Borel probability space* is a probability space $(\mathcal{S}, \mathcal{F}, \mu)$ such that $(\mathcal{S}, \mathcal{F})$ is a Borel space.

We state a general equivalence theorem, which is the poset version of [10, Theorem 7.1] for graph limits. (This theorem in [10] is for simplicity stated only for functions defined on $([0,1],\mathcal{B},\lambda,<)$, but it extends to arbitrary Borel probability spaces without problems, see [15, Theorem 8.6].) The parts (viii) and (ix) are modelled after similar results for graph limits in [5]. (For graph limits, [5] also gives an equivalent condition with $W_1 = V^{\varphi_1}$ and $W_2 = V^{\varphi_2}$ for some φ_1, φ_2 and V. We conjecture that a similar result is true for poset limits too, but we have not yet investigated this.)

If W is a function $\mathcal{S}^2 \to [0,1]$, where $\mathcal{S}$ is a probability space, we say following [5] that $x_1, x_2 \in \mathcal{S}$ are *twins* if $W(x_1,y) = W(x_2,y)$ and $W(y,x_1) = W(y,x_2)$ for a.e. $y \in \mathcal{S}$. We say that W is *almost twinfree* if there exists a null set $N \subset \mathcal{S}$ such that there are no twins $x_1, x_2 \in \mathcal{S} \setminus N$ with $x_1 \neq x_2$.

Theorem 7.1. *Suppose that $W_1 : \mathcal{S}_1^2 \to [0,1]$ and $W_2 : \mathcal{S}_2^2 \to [0,1]$ are two kernels defined on two ordered probability spaces $(\mathcal{S}_1,\mathcal{F}_1,\mu_1,\prec_1)$ and $(\mathcal{S}_2,\mathcal{F}_2,\mu_2,\prec_2)$ such that $(\mathcal{S}_1,\mu_1)$ and $(\mathcal{S}_2,\mu_2)$ are Borel spaces, and let $\Pi_1 = \Pi_{W_1}$ and $\Pi_2 = \Pi_{W_2}$ be the corresponding poset limits in $\mathcal{P}_\infty$. Then the following are equivalent.*

(i) $\Pi_1 = \Pi_2$ *in* $\mathcal{P}_\infty$.

(ii) $t(Q,\Pi_1) = t(Q,\Pi_2)$ *for every poset* Q.

(iii) *The exchangeable random infinite posets $P(\infty,W_1)$ and $P(\infty,W_2)$ have the same distribution.*

(iv) *The random posets $P(n,W_1)$ and $P(n,W_2)$ have the same distribution for every finite n.*

(v) *There exist measure preserving maps $\varphi_j : [0,1] \to \mathcal{S}_j$, $j = 1,2$, such that $W_1^{\varphi_1} = W_2^{\varphi_2}$ a.s., i.e. $W_1\big(\varphi_1(x),\varphi_1(y)\big) = W_2\big(\varphi_2(x),\varphi_2(y)\big)$ a.e. on $[0,1]^2$.*

(vi) *There exists a measurable mapping $\psi : \mathcal{S}_1 \times [0,1] \to \mathcal{S}_2$ that maps $\mu_1 \times \lambda$ to μ_2 such that $W_1(x,y) = W_2\big(\psi(x,t_1),\psi(y,t_2)\big)$ for a.e. $x,y \in \mathcal{S}_1$ and $t_1,t_2 \in [0,1]$. (This is equivalent to $W_1^\pi = W_2^\psi$ a.s. on $\mathcal{S}^2$, where $\mathcal{S} := \mathcal{S}_1 \times [0,1]$ and $\pi : \mathcal{S} \to \mathcal{S}_1$ is projection onto the first coordinate.)*

(vii) $\delta_\square(W_1,W_2) = 0$.

If further W_2 is almost twinfree, then these are also equvalent to:

(viii) *There exists a measure preserving map $\varphi : \mathcal{S}_1 \to \mathcal{S}_2$ such that $W_1 = W_2^\varphi$ a.s., i.e. $W_1(x,y) = W_2\big(\varphi(x),\varphi(y)\big)$ a.e. on $\mathcal{S}_1^2$.*

If both W_1 and W_2 are almost twinfree, then these are also equvalent to:

(ix) *There exists a measure preserving map $\varphi : \mathcal{S}_1 \to \mathcal{S}_2$ such that φ is a bimeasurable bijection of $\mathcal{S}_1 \setminus N_1$ onto $\mathcal{S}_2 \setminus N_2$ for some null sets $N_1 \subset \mathcal{S}_1$ and $N_2 \subset \mathcal{S}_2$, and $W_1 = W_2^\varphi$ a.s., i.e. $W_1(x,y) = W_2\big(\varphi(x),\varphi(y)\big)$ a.e. on $\mathcal{S}_1^2$. (If further $(\mathcal{S}_2,\mu_2)$ has no atoms, for example if $\mathcal{S}_2 = [0,1]$, then we may take $N_1 = N_2 = \emptyset$.)* □

The proof is very similar to the proofs of [10, Theorem 7.1] and [15, Theorem 8.6]; see [14] for details.

8. Proofs of Theorems 1.14–1.15

Theorem 1.15 says that the mapping $W \mapsto \Pi_W$ is a homeomorphism of this space $(\overline{\mathcal{W}}_{\mathsf{P}}, \delta_\square)$ onto $\mathcal{P}_\infty$.

Proof of Theorem 1.15. The mapping $W \mapsto \Pi_W \in \mathcal{P}_\infty$ is well-defined and continuous on $\overline{\mathcal{W}}_{\mathsf{P}}$ by Lemma 6.4 and the construction of $\mathcal{P}_\infty$ (see Theorem 3.5); further, the mapping is surjective by Theorem 1.7 and it is injective by Theorem 7.1 ((i) $\implies$ (vii)), using the definition (6.9). Since $\overline{\mathcal{W}}_{\mathsf{P}}$ is compact by Theorem 6.9, the mapping is thus a homeomorphism. $\square$

Proof of Theorem 1.14. Let $W_n = W_{P_n}$. Thus $\Pi_{P_n} := \Pi_{W_n}$ and $t(Q, P_n) = t(Q, \Pi_{P_n}) = t(Q, \Pi_{W_n})$ for every $Q \in \mathcal{P}$ by Example 1.13. It follows from Theorems 3.4 and 3.5 that $P_n \to \Pi \iff \Pi_{W_n} \to \Pi$, and the result follows from Theorem 1.15. $\square$

9. Further examples

Example 9.1. For each finite n, all totally ordered sets with n elements are isomorphic, and there is thus a unique unlabelled totally ordered poset in $\mathcal{P}_n$ which we denote by T_n. Let $(\mathcal{S}, \mathcal{F}, \mu, <)$ be a totally ordered set with a continuous probability measure μ (i.e., a probability measure such that $\mu\{x\} = 0$ for every $x \in \mathcal{S}$), and let $W(x, y) = \mathbf{1}[x < y]$ as in Example 1.5. Since μ is continuous, the random points X_i in Definition 1.3 are (a.s.) distinct, and thus, see Example 1.5, $P(n, W) = P(n, \mathcal{S})$ is isomorphic to a subset of $\mathcal{S}$ and thus totally ordered. In other words, $P(n, W) = T_n$ as unlabelled posets. (As labelled posets, $P(n, W) \overset{\mathrm{d}}{=} \widehat{T}_n$, which is obtained by applying a random permutation to $[n]$ with the usual order.) By Theorem 1.7(i), thus $T_n \to \Pi_W$, which shows that Π_W does not depend on the choices of $\mathcal{S}$ and μ. We write Π_T for this poset limit and have thus shown that there exists a (unique) poset limit $\Pi_T \in \mathcal{P}_\infty$ such that $T_n \to \Pi_T$ and $P(n, \Pi_T) = T_n$ for all finite n. We may call Π_T the *total poset limit.*

It is convenient to choose $\mathcal{S}$ as $[0, 1]$ with Lebesgue measure; we then see that $P(\infty, \Pi_T)$ is the random infinite total order defined by a sequence of i.i.d. random points in $[0, 1]$ with the standard order.

Note that μ has to be continuous in this example; otherwise (i.e., if μ has an atom), there will (a.s.) be repetitions in $X_1, X_2, \dots$ and thus incomparable points in $P(\infty, W)$ (and with positive probability in $P(n, W)$ for finite $n \geq 2$); hence $P(\infty, \Pi_W) = P(\infty, W) \overset{\mathrm{d}}{\neq} P(\infty, \Pi_T)$ and $\Pi_W \neq \Pi_T$ by Theorem 1.17(ii). In particular, Π_P defined in Example 1.13 for a finite totally ordered set $P = T_m$ does *not* equal Π_T. (Although, as a consequence of (1.8), $\Pi_{T_m} \to \Pi_T$ in $\mathcal{P}_\infty$ as $m \to \infty$.)

Example 9.2. The other extreme is the poset where $x \not< y$ for all x, y; we call these posets trivial, and let E_n denote the (unique) unlabelled trivial poset with $|E_n| = n$. Then, trivially, $t(Q, E_n) = 0$ for every finite poset Q

that is not itself trivial, while $t(E_m, E_n) = 1$ for all m and n. Consequently the sequence (E_n) converges, and the limit is a poset limit $\Pi_0 \in \mathcal{P}_\infty$ with

$$t(Q, \Pi_0) = \begin{cases} 1, & Q = E_m \text{ for some } m, \\ 0, & \text{otherwise.} \end{cases} \tag{9.1}$$

If $\mathcal{S}$ is any ordered probability space and $W = 0$, which always is a poset kernel, then $P(n, W)$ is trivial for all $n < \infty$, and by Theorem 1.7(i) or (ii), $\Pi_W = \Pi_0$. (This explains our notation Π_0.)

Taking $P = E_n$ in Example 1.13, we see further by (1.8) that $\Pi_{E_n} = \Pi_0$ for every n. Trivially, $\widehat{E}_n = E_n$, and by (1.6) $P(\infty, \Pi_0)$ is the trivial infinite poset on $\mathbb{N}$.

Example 9.3. Let $\mathcal{S} = [0, 1]^2$ with Lebesgue measure and the product order $(x_1, x_2) < (y_1, y_2)$ if $x_1 < y_1$ and $x_2 < y_2$. Again, let $W(x, y) = \mathbf{1}[x < y]$ as in Example 1.5. Then $P(n, W)$ is the poset defined by n random points in $[0, 1]^2$, which also can be described as the intersection of two independent random total orders on $[n]$.

Example 9.4. Let $G(n, p)$ denote the random graph with n vertices $\{1, \dots, n\}$ where each possible edge ij appears with probability p, independently of all other edges. We make $G(n, p)$ into a (random) poset by directing each edge from the smaller endpoint to the larger, and then taking the transitive closure, see [22] and the references therein. In other words, $i \prec j$ in $G(n, p)$ if and only if there is an increasing path $i = i_1, i_2, \dots, i_n = j$ in $G(n, p)$. We use $G(n, p)$ to denote this random poset too.

Brightwell and Georgiou [8] showed that if $n \to \infty$ and $p \to 0$ such that $pn/\log n \to a \in [0, \infty]$, then

$$\mathbb{E}\, t(Q, G(n, p)) \to t(Q, W_a) = t(Q, \Pi_{W_a}),$$

where W_a is the kernel on $([0, 1], \mathcal{B}, \lambda, <)$ given by $W_a(x, y) := \mathbf{1}[y - x > a^{-1}]$. By Theorem 4.1, thus $G(n, p) \xrightarrow{\mathrm{d}} \Pi_{W_a}$; since Π_{W_a} is non-random, this means $G(n, p) \xrightarrow{\mathrm{p}} \Pi_{W_a}$.

Note that if $a \le 1$, then $W_a = 0$ and the limit $\Pi_{W_a} = \Pi_0$; see Example 9.2. The other extreme is $a = \infty$; then $W_a(x, y) = \mathbf{1}[y > x]$ on the totally ordered set $[0, 1]$, so the limit is Π_T, see Example 9.1.

The limits Π_{W_a} are such that $P(n, \Pi_{W_a})$ a.s. is a semiorder; they are thus examples of *semiorder limits* (although $G(n, p)$ in general is not a semiorder), see [16] for definitions, details and further results.

Example 9.5. A (finite) poset has an *interval order* if it is isomorphic to a set of intervals in $\mathbb{R}$ with $I \prec J$ if and only if $x < y$ for all $x \in I$, $y \in J$ (i.e., I lies entirely to the left of J). We define an *interval order limit* as a poset limit that is a limit of a sequence of finite posets with interval orders.

Let $\mathcal{S} = \{(x, y) : 0 \le x \le y \le 1\}$ with the partial order $(x_1, y_1) \prec (x_2, y_2)$ if $y_1 < x_2$. We can interpret $\mathcal{S}$ as the set of closed intervals in [0,1], with the order $I \prec J$ just defined. Any probability measure μ on $\mathcal{S}$

thus defines a distribution of random intervals, and the kernel $W(\mathbf{x}_1, \mathbf{x}_2) := \mathbf{1}[\mathbf{x}_1 \prec \mathbf{x}_2]$ as in Example 1.5 yields random posets $P(n, W)$ with interval orders, and an interval order limit Π. Conversely, the complement of the comparability graph of an interval order is an interval graph, and it follows by the arguments used for interval graph limits in [9] that every interval order limit can be represented by the kernel W on the ordered probability space $(\mathcal{S}, \mu, \prec)$ for some probability measure μ on $\mathcal{S}$. Since W is fixed, we thus represent interval order limits by probability measures μ on $\mathcal{S}$. The representation is non-unique, and we may further, for example, require that the left marginal distribution of μ (i.e., the distribution of the left endpoint of the random interval described by μ) is uniform on [0,1]. See [16] for details and further results, including a unique representation.

We note that although it is natural to represent an interval order limit Π by the kernel W on $(\mathcal{S}, \mu)$, Π can also be represented by a kernel on $([0,1], \mathcal{B}, \lambda, <)$. (Thus Problem 1.10 has a positive answer in this case.) To see this, we construct a measure preserving map $\varphi : ([0,1], \lambda) \to (\mathcal{S}, \mu)$ such that $\varphi(s) \prec \varphi(t) \implies s < t$; then W^φ is a kernel on $[0,1]$ that represents Π. We may construct φ by first partitioning $\mathcal{S}$ into $\mathcal{S}_0 := \{(x,y) : x \le y < 1/2\}$, $\mathcal{S}_{01} := \{(x,y) : x < 1/2 \le y\}$ and $\mathcal{S}_1 := \{(x,y) : 1/2 \le x \le y\}$, and a corresponding partitioning of $[0,1]$ into $I_0 := [0, \mu(\mathcal{S}_0))$, $I_{01} := [\mu(\mathcal{S}_0), 1 - \mu(\mathcal{S}_1))$ and $I_1 := [1 - \mu(\mathcal{S}_1), 1]$. Noting that all elements of $\mathcal{S}_{01}$ are incomparable, we define φ on I_{01} as any measure preserving map $I_{01} \to \mathcal{S}_{01}$. We then continue recursively and define $\varphi : I_0 \to \mathcal{S}_0$ and $I_1 \to \mathcal{S}_1$ by partitioning $\mathcal{S}_0$ and $\mathcal{S}_1$ into three parts each, and so on. (In the kth stage, the partitioning is according to the kth binary digit of x and y.) Let $\Delta := \{(x,x)\}$ be the diagonal in $\mathcal{S}$. If $\mu(\Delta) = 0$, then the recursive procedure just described defines φ at least a.e. on $[0,1]$. If $\mu(\Delta) > 0$, there will remain a Cantor like subset of $[0,1]$ of measure $\mu(\Delta)$; the construction then is completed by mapping this set to Δ by an increasing measure preserving map.

Example 9.6. Given any poset limit Π and $p \in [0,1]$, modify the exchangeable random infinite poset $P(\infty, \Pi)$ by randomly selecting elements with probability $1-p$ each, independently, and making them uncomparable to everything. This gives a new exchangeable random infinite poset, which satisfies Theorem 1.18(iii) and thus by Theorem 1.17(ii) equals $P(\infty, \Pi_{(p)})$ for some poset limit $\Pi_{(p)}$. We can regard $\Pi_{(p)}$ as a thinning of Π. It is easily seen that for every finite poset Q,

$$t(Q, \Pi_{(p)}) = p^{|Q|_+} t(Q, \Pi), \tag{9.2}$$

where $|Q|_+ := |\{x \in Q : x < y \text{ or } y < x \text{ for some } y\}|$ is the number of elements of Q that are comparable to at least one other element.

Clearly, $\Pi_{(1)} = \Pi$, while, by (9.2) and (9.1), $\Pi_{(0)} = \Pi_0$ defined in Example 9.2, for every $\Pi \in \mathcal{P}_\infty$. Moreover, (9.2) shows that the map $(\Pi, s) \mapsto \Pi_{(s)}$ is a continuous map $\mathcal{P}_\infty \times [0,1] \to \mathcal{P}_\infty$.

Consequently, this map defines a homotopy between the identity map $\mathcal{P}_\infty \to \mathcal{P}_\infty$ and a constant map, which shows that $\mathcal{P}_\infty$ is a contractible topological space. In particular, $\mathcal{P}_\infty$ is connected and simply connected.

10. Poset limits as digraph limits

As said repeatedly, we can regard posets as digraphs, which yields an inclusion mapping $\mathcal{P} \to \mathcal{D}$. We saw in Section 3 that this mapping extends to a (unique) continuous inclusion mapping $\overline{\mathcal{P}} \to \overline{\mathcal{D}}$; we may thus regard $\overline{\mathcal{P}}$ as a compact subset of $\overline{\mathcal{D}}$, with $\mathcal{P}_\infty$ a compact subset of $\mathcal{D}_\infty$. We can now characterize the subset $\mathcal{P}_\infty$ of $\mathcal{D}_\infty$ in several ways.

We first recall that, as shown in Diaconis and Janson [10], the digraph limits in $\mathcal{D}_\infty$ can be represented by quintuples $\mathbf{W} = (W_{00}, W_{01}, W_{10}, W_{11}, w)$ where $W_{\alpha\beta} : [0,1]^2 \to [0,1]$ and $w : [0,1] \to \{0,1\}$ are measurable functions such that $\sum_{\alpha,\beta=0}^{1} W_{\alpha\beta}(x,y) = 1$ and $W_{\alpha\beta}(x,y) = W_{\beta\alpha}(y,x)$ for $\alpha,\beta \in \{0,1\}$ and $x,y \in [0,1]$. Let $\mathcal{W}_5$ be the set of all such quintuples. For $\mathbf{W} \in \mathcal{W}_5$, we define a random infinite digraph $G(\infty,\mathbf{W})$ by specifying its edge indicators I_{ij} as follows (cf. Definition 1.3): we first choose a sequence $X_1, X_2, \dots$ of i.i.d. random variables uniformly distributed on $[0,1]$, and then, given this sequence, let $I_{ii} = w(X_i)$ and for each pair (i,j) with $i < j$ choose I_{ij} and I_{ji} at random such that

$$\mathbb{P}(I_{ij} = \alpha \text{ and } I_{ji} = \beta) = W_{\alpha\beta}(X_i, X_j), \qquad \alpha,\beta \in \{0,1\}; \tag{10.1}$$

this is done independently for all pairs (i,j) with $i < j$ (conditionally given $\{X_k\}$). The infinite random digraph $G(\infty,\mathbf{W})$ is exchangeable, and it is shown in [10], by digraph analogues of Theorems 1.18 and 1.17 above, that its distribution is an extreme point in the set of exchangeable distributions and that it corresponds to a digraph limit $\Gamma_{\mathbf{W}}$; for example, $G(n,\mathbf{W}) := G(\infty,\mathbf{W})|_{[n]} \to \Gamma_{\mathbf{W}}$ in $\overline{\mathcal{D}}$ a.s.

Theorem 10.1. *Let $\Gamma \in \mathcal{D}_\infty$ be a digraph limit. Then the following are equivalent.*

(i) $\Gamma \in \mathcal{P}_\infty$*, i.e.,* Γ *is a poset limit.*
(ii) $t_{\mathrm{ind}}(F,\Gamma) = 0$ *for every finite digraph F that is not a poset.*
(iii) $t_{\mathrm{ind}}(\mathsf{C}_1,\Gamma) = t_{\mathrm{ind}}(\mathsf{C}_2,\Gamma) = t_{\mathrm{ind}}(\mathsf{C}_3,\Gamma) = t_{\mathrm{ind}}(\mathsf{P}_2,\Gamma) = 0$.
(iv) *If $\mathbf{W} = (W_{00}, W_{01}, W_{10}, W_{11}, w)$ is some (any) quintuplet representing Γ, then $w = 0$ a.e., $W_{11} = 0$ a.e., and $\{(x,y,z) : W_{10}(x,y) > 0$ and $W_{10}(y,z) > 0$ and $W_{10}(x,z) < 1\}$ is a null set in* $[0,1]^3$.
(v) *There exists a quintuplet $\mathbf{W} = (W_{00}, W_{01}, W_{10}, W_{11}, w)$ representing Γ with $w = 0$, $W_{11}(x,y) = 0$, $W_{10}(x,x) = 0$, and W_{10} satisfying* (1.2). □

We omit the proof, which is rather straightforward; see [14] for details.

Acknowledgement. This work was stimulated by helpful discussions with Graham Brightwell and Malwina Luczak during the programme "Combinatorics and Statistical Mechanics" at the Isaac Newton Institute, Cambridge, 2008, where SJ was supported by a Microsoft fellowship. Parts of this work was done at Institut Mittag-Leffler, Djursholm, 2009. I thank an anonymous referee for many helpful comments.

References

[1] D. Aldous, Representations for partially exchangeable arrays of random variables. *J. Multivar. Anal.* **11**, 581–598, 1981.

[2] T. Austin. On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probab. Surv.* **5** (2008), 80–145.

[3] B. Bollobás, S. Janson & O. Riordan, The cut metric, random graphs, and branching processes. *Journal of Statistical Physics* **140**:2 (2010), 289-335

[4] B. Bollobás & O. Riordan, Metrics for sparse graphs. *Surveys in Combinatorics 2009*, LMS Lecture Notes Series 365, Cambidge Univ. Press, 2009, pp. 211–287.

[5] C. Borgs, J. T. Chayes & L. Lovász, Moments of two-variable functions and the uniqueness of graph limits. *Geom. Funct. Anal.* **19** (2010), no. 6, 1597–1619.

[6] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós & K. Vesztergombi, Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.* **219** (2008), no. 6, 1801–1851.

[7] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós & K. Vesztergombi, Convergent sequences of dense graphs II: Multiway cuts and statistical physics. Preprint, 2007. `http://research.microsoft.com/~borgs/`

[8] G. Brightwell & N. Georgiou, Continuum limits for classical sequential growth models. *Random Structures Algorithms* **36** (2010), no. 2, 218–250.

[9] P. Diaconis, S. Holmes & S. Janson, Interval graph limits. Preprint, 2011. `arXiv:1102.2841v1.`

[10] P. Diaconis & S. Janson, Graph limits and exchangeable random graphs. *Rendiconti di Matematica* **28** (2008), 33–61.

[11] G. Elek & B. Szegedy, Limits of hypergraphs, removal and regularity lemmas. A non-standard approach. `arXiv:0705.2179v1.`

[12] A. Frieze & R. Kannan, Quick approximation to matrices and applications. *Combinatorica* **19** (1999), 175–220.

[13] D. Hoover, *Relations on Probability Spaces and Arrays of Random Variables.* Preprint, Institute for Advanced Study, Princeton, NJ, 1979.

[14] S. Janson, Poset limits and exchangeable random posets. Preprint version of the present paper. `arXiv:0902.0306v1.`

[15] S. Janson, Graphons, cut norm and distance, couplings and rearrangements. Preprint, 2010–2011. `arXiv:1009.2376v3.`

[16] S. Janson, Limits of interval orders and semiorders. Preprint, 2011. `arXiv:1104.1264v1.`

[17] O. Kallenberg, *Foundations of Modern Probability*, 2nd ed., Springer, New York, 2002.
[18] O. Kallenberg, *Probabilistic Symmetries and Invariance Principles.* Springer, New York, 2005.
[19] L. Lovász & B. Szegedy, Limits of dense graph sequences. *J. Comb. Theory B* **96**, 933–957, 2006.
[20] L. Lovász & B. Szegedy, Szemerédi's lemma for the analyst. *Geom. Funct. Anal.* **17** (2007), no. 1, 252–270.
[21] K. R. Parthasarathy, *Probability Measures on Metric Spaces.* Academic Press, New York, 1967.
[22] B. Pittel & R. Tungol, A phase transition phenomenon in a random directed acyclic graph. *Random Struct. Alg.* **18** (2001), no. 2, 164–184.
[23] E. M. Stein, *Singular Integrals and Differentiability Properties of Functions.* Princeton University Press, Princeton, N.J., 1970.
[24] T. Tao, A correspondence principle between (hyper)graph theory and probability theory, and the (hyper)graph removal lemma. *J. Anal. Math.* **103** (2007), 1–45.

Department of Mathematics, Uppsala University, PO Box 480, SE-751 06 Uppsala, Sweden

E-mail address: svante.janson@math.uu.se
URL: http://www.math.uu.se/~svante/