

Exchangeable trait allocations

Trevor Campbell^{1,*}, Diana Cai^{2,**}, and Tamara Broderick^{1,†}

¹*Computer Science and Artificial Intelligence Laboratory (CSAIL)
Massachusetts Institute of Technology
Cambridge, MA, USA 02139*

e-mail: *tdjc@mit.edu; †tbroderick@csail.mit.edu

²*Department of Computer Science
Princeton University
Princeton, NJ, USA 08544*

e-mail: **dcai@cs.princeton.edu

Abstract: Trait allocations are a class of combinatorial structures in which data may belong to multiple groups and may have different levels of belonging in each group. Often the data are also exchangeable, i.e., their joint distribution is invariant to reordering. In clustering—a special case of trait allocation—exchangeability implies the existence of both a de Finetti representation and an exchangeable partition probability function (EPPF), distributional representations useful for computational and theoretical purposes. In this work, we develop the analogous de Finetti representation and exchangeable trait probability function (ETPF) for trait allocations, along with a characterization of all trait allocations with an ETPF. Unlike previous feature allocation characterizations, our proofs fully capture single-occurrence “dust” groups. We further introduce a novel constrained version of the ETPF that we use to establish an intuitive connection between the probability functions for clustering, feature allocations, and trait allocations. As an application of our general theory, we characterize the distribution of all edge-exchangeable graphs, a class of recently-developed models that captures realistic sparse graph sequences.

Keywords and phrases: Trait allocation, exchangeability, paintbox, probability function, partition, feature allocation, graph, vertex allocation, edge exchangeability.

Received September 2016.

1. Introduction

Representation theorems for exchangeable random variables are a ubiquitous and powerful tool in Bayesian modeling and inference. In many data analysis problems, we impose an order, or indexing, on our data points. This indexing can arise naturally—if we are truly observing data in a sequence—or can be artificially created to allow their storage in a database. In this context, exchangeability expresses the assumption that this order is arbitrary and should not affect our analysis. For instance, we often assume a sequence of data points is an *infinite exchangeable sequence*, i.e., that the distribution of any finite subsequence is invariant to reordering. Though this assumption may seem weak, de Finetti’s theorem (de Finetti, 1931; Hewitt and Savage, 1955) tells us that in this case, we can assume that a latent parameter exists, that our data are

independent and identically distributed (i.i.d.) conditional on this parameter, and that the parameter itself has a distribution. Thus, de Finetti's theorem may be seen as a justification for a Bayesian model and prior—and, in fact, for the infinite-dimensional priors provided by Bayesian nonparametrics (Jordan, 2010).

De Finetti-style representation theorems have provided many other useful insights for modeling and inference within Bayesian analysis. For example, consider *clustering* problems, where the inferential goal is to assign data points to mutually exclusive and exhaustive groups. It is typical to assume that the distribution of the clustering—i.e., the assignment of data points to clusters—is invariant to the ordering of the data points. In this case, two different representation theorems have proved particularly useful in practice. First, Kingman (1978) showed that exchangeability in clustering implies the existence of a latent set of probabilities (known as the “Kingman paintbox”) from which cluster assignments are chosen i.i.d.. It is straightforward to show from the Kingman paintbox representation that exchangeable clustering models enforce linear growth in cluster size as a function of the size of the total data. By contrast, many real-world clustering problems, such as disambiguating census data or clustering academic papers by originating lab, exhibit sublinear growth in cluster size (e.g., Wallach et al., 2010; Broderick and Steorts, 2014; Miller et al., 2016). Thus, the Kingman paintbox representation allows us to see that exchangeable clustering models are misspecified for these examples. Similarly, Pitman (1995) showed that clustering exchangeability is equivalent to the existence of an *exchangeable partition probability function* (EPPF). The EPPF and similar developments have led to algorithms that allow practical inference specifically with the Dirichlet process mixture (Escobar, 1994; Escobar and West, 1995) and more generally in other clustering models (Pitman and Yor, 1997; Ishwaran and James, 2001, 2003; Lee et al., 2013).

In this work, we develop and characterize a generalization of clustering models that we call *trait allocation models*. Trait allocations apply when data may belong to more than one group (a *trait*), and may exhibit nonnegative integer levels of belonging in each group. For example, a document might exhibit multiple words in a number of topics, a participant in a social network might send multiple messages to each of her friend groups, or a DNA sequence might exhibit different numbers of genes from different ancestral populations. Trait allocations generalize both clustering, where data must belong to exactly one group, and feature allocations (Griffiths and Ghahramani, 2005; Broderick, Pitman and Jordan, 2013), where data exhibit binary membership in multiple groups. Authors have recently proposed a number of models for trait allocations (e.g., Titsias, 2008; Zhou et al., 2012; Zhou, 2014; James, 2017; Broderick et al., 2015; Roychowdhury and Kulis, 2015). But as of yet, there is no characterization either of the class of exchangeable trait allocation models or of classes of exchangeable trait allocation models that are particularly amenable to inference. The consequences of the exchangeability assumption in this setting have not been explored. In this work, we provide characterizations of both the full class of exchangeable trait allocations and those with EPPF-like probability distributions.

This work not only unifies and generalizes past research on partitions and feature allocations, but provides a natural avenue for the study of other practical exchangeable combinatorial structures.

We begin by formally defining trait allocations, random sequences thereof, and exchangeability in Section 2. In Section 3, we introduce ordered trait allocations via the lexicographic ordering. We use these constructions to establish a de Finetti representation for exchangeable trait allocations in Section 4 that is analogous to the Kingman paintbox representation for clustering. Our new representation handles *dust*, the case where some traits may appear for just a single data point. This work therefore also extends previous work on the special case of exchangeable feature allocations to the fully general case, whereas previously it was restricted to the dustless case (Broderick, Pitman and Jordan, 2013). In Section 5, we develop an EPPF-like function to describe distributions over exchangeable trait allocations and characterize the class of trait allocations to which it applies. We call these *exchangeable trait probability functions* (ETPFs). Just as in the partition and feature allocation cases, the class of random trait allocations with probability functions represents a class of trait allocations that are particularly amenable to approximate posterior inference in practice—and therefore of particularly pressing interest to characterize. In Section 5, we introduce new concepts we call *constrained ETPFs*, which are the combinatorial analogue of earlier work on restricted nonparametric processes (Williamson, MacEachern and Xing, 2013; Doshi-Velez and Williamson, 2017). In Sections 5 and 6, we show how constrained ETPFs capture earlier probability functions for numerous exchangeable models within a single framework. In Section 6, we apply both our de Finetti representation and constrained ETPF to characterize *edge-exchangeable* graphs, a recently developed form of exchangeability for graph models that allows sparse projective sequences of graphs (Broderick and Cai, 2015; Crane and Dempsey, 2015; Cai, Campbell and Broderick, 2016; Crane and Dempsey, 2016a; Williamson, 2016). A similar representation generalizing partitions and edge-exchangeable (hyper)graphs has been studied in concurrent work (Crane and Dempsey, 2016b) on *relational exchangeability*, first introduced by Ackerman (2015); Crane and Towsner (2015)—but here we additionally explore the existence of a trait frequency model, the existence of a constrained trait frequency model and its connection to clustering and feature allocations, and the various connections between frequency models and probability functions.

1.1. Notation and conventions

Definitions are denoted by the symbol $:=$. The natural numbers are denoted $\mathbb{N} := \{1, 2, \dots\}$ and the nonnegative reals $\mathbb{R}_+ := [0, \infty)$. We let $[N] := \{1, 2, \dots, N\}$ for any $N \in \mathbb{N}$. Sequences are denoted with parentheses, with indices suppressed only if they are clear from context. For example, (x_k) is the sequence $(x_k)_{k \in \mathbb{N}}$ and (x_{kj}) is the sequence $(x_{kj})_{k, j \in \mathbb{N}}$, while $(x_{kj})_{j=1}^\infty$ is the sequence x_{k1}, x_{k2}, \dots with k fixed. The notation $A \subset B$ means A is a (not necessarily proper) subset of B . The indicator function is denoted $\mathbb{1}(\dots)$; for example, $\mathbb{1}(x \in A)$ is 1 if

$x \in A$, and 0 otherwise. For any multiset x of elements in a set \mathcal{X} , we denote $x(y)$ to be the multiplicity of y in x for each $y \in \mathcal{X}$. Two multisets x, x' of \mathcal{X} are said to be equal, denoted $x = x'$, if the multiplicity of all elements $y \in \mathcal{X}$ are equal in both x and x' , i.e. $\forall y \in \mathcal{X}, x(y) = x'(y)$. For any finite or infinite sequence, we use subscript k to denote the k^{th} element in the sequence. For sequences of (multi)sets, if k is beyond the end of the sequence, the subscript k operation returns the empty set. Equality in distribution and almost surely are denoted $\stackrel{d}{=}$ / $\stackrel{a.s.}{=}$, and convergence almost surely/in probability/in distribution is denoted $\xrightarrow{a.s.}$ / \xrightarrow{p} / \xrightarrow{d} . We often use cycle notation for permutations (see Dummit and Foote (2004, p. 29)): for example, $\pi = (12)(34)$ is the permutation π with $\pi(1) = 2, \pi(2) = 1, \pi(3) = 4, \pi(4) = 3$, and $\pi(k) = k$ for $k > 4$. We use the notation $X \sim (\theta_j)_{j=0}^{\infty}$ to denote sampling X from the categorical distribution on $\{0\} \cup \mathbb{N}$ with probabilities $\mathbb{P}(X = j) = \theta_j$ for $j \in \{0\} \cup \mathbb{N}$. The symbol $\times_N S_N$ for a sequence of sets (S_N) denotes $S_1 \times S_2 \times \dots$, their infinite product space.

2. Trait allocations

We begin by formalizing the concepts of a *trait* and *trait allocation*. We assume that our sequence of data points is indexed by \mathbb{N} . As a running example for intuition, consider the case where each data point is a document, and each trait is a topic. Each document may have multiple words that belong to each topic. The degree of membership of the document in the topic is the number of words in that topic. We wish to capture the assignment of data points to the traits they express but in a way that does not depend on the type of data at hand. Therefore, we focus on the indices to the data points. This leads to the definition of traits as multisets of the data indices, i.e., the natural numbers. E.g., $\tau = \{1, 1, 3\}$ is a trait in which the datum at index 1 has multiplicity 2, and the datum at index 3 has unit multiplicity. In our running example, this trait might represent the topic about sports; the first document has two sports words, and the third document has one sports word.

Definition 2.1. A *trait* is a finite, nonempty multiset of \mathbb{N} .

Let the set of all traits be denoted \mathbb{T} . A single trait is not sufficient to capture the combinatorial structure underlying the first $N \in \mathbb{N}$ data in the sequence: each datum may be a member of multiple traits (with varying degrees of membership). The traits have no inherent order just as the topics “sports”, “arts”, and “science” have no inherent order. And each document may contain words from multiple topics. Building from Definition 2.1 and motivated by these desiderata, we define a finite trait allocation as a finite multiset of traits. For example, $t_4 = \{\{1\}, \{3, 4\}, \{3, 3\}, \{3, 3\}, \{1, 1, 4\}\}$ represents a collection of traits expressed by the first 4 data points in a sequence. In this case, index 1 is a member of two traits, index 2 is a member of none, and so on. Throughout, we assume that each datum at index $n \in \mathbb{N}, n \leq N$ belongs to only finitely many latent traits. Further, for a data set of size N , any index $n > N$ should

not belong to any trait; the allocation t_N represents traits expressed by only the first N data. These statements are formalized in Definition 2.2.

Definition 2.2. A *trait allocation* of $[N]$ is a multiset t_N of traits, where

$$\forall n \in \mathbb{N} : n \leq N, \quad \sum_{\omega \in \mathbb{T}} t_N(\omega) \cdot \omega(n) < \infty \quad (2.1)$$

$$\forall n \in \mathbb{N} : n > N, \quad \sum_{\omega \in \mathbb{T}} t_N(\omega) \cdot \omega(n) = 0. \quad (2.2)$$

Let \mathcal{T}_N be the set of trait allocations of $[N]$, and define \mathcal{T} to be the set of all *finite trait allocations*, $\mathcal{T} := \bigcup_N \mathcal{T}_N$. Two notable special cases of finite trait allocations that have appeared in past work are *feature allocations* (Griffiths and Ghahramani, 2005; Broderick, Pitman and Jordan, 2013) and *partitions* (Kingman, 1978; Pitman, 1995). Feature allocations are the natural combinatorial structure underlying feature learning, where each datum expresses each trait with multiplicity at most 1. For example, $t_4 = \{\{1\}, \{3, 4\}, \{3, 1\}, \{3\}\}$ is a feature allocation of $[4]$. Note that each index may be a member of multiple traits. Partitions are the natural combinatorial structure underlying clustering, where the traits form a partition of the indices. For example, $t_4 = \{\{1, 3, 4\}, \{2\}\}$ is a partition of $[4]$, since its traits are disjoint and their union is $[4]$. The theory in the remainder of the paper will be applied to recover past results for these structures as corollaries.

Up until this point, we have dealt solely with finite sequences of N data. However, in many data analysis problems, it is more natural (or at least an acceptable simplifying approximation) to treat the observed sequence of N data as the beginning of an infinite sequence. As each datum arrives, it adds its own index to the traits it expresses, and in the process introduces any previously uninstantiated traits. For example, if after 3 observations we have $t_3 = \{\{1\}, \{1, 2\}\}$, then observing the next might yield $t_4 = \{\{1\}, \{1, 2, 4, 4\}, \{4, 4\}\}$. Note that when an index is introduced, none of the earlier indices' memberships to traits are modified; the sequence of finite trait allocations is *consistent*. To make this rigorous, we define the *restriction* of a trait (allocation), which allows us to relate two trait allocations $t_N, t_M \in \mathcal{T}$ of differing N and M . The restriction operator $|_M$ —provided by Definition 2.3 and acting on either traits or finite trait allocations—removes all indices greater than M from all traits, and does not modify the multiplicity of indices less than or equal to M . If any trait becomes empty in this process, it is removed from the allocation. For example, $\{\{1, 3, 4\}, \{1, 2\}, \{4\}\}_1 = \{\{1\}, \{1\}\}$. Two trait allocations are said to be *consistent*, per Definition 2.4, if one can be restricted to recover the other. Thus, $\{\{1, 3, 4\}, \{1, 2\}, \{4\}\}$ and $\{\{1\}, \{1\}\}$ are consistent finite trait allocations.

Definition 2.3. The *restriction* $|_M : \mathbb{T} \rightarrow \mathbb{T}$ of a trait τ to $M \in \mathbb{N}$ is defined as

$$\tau|_M(m) := \begin{cases} \tau(m) & m \leq M \\ 0 & m > M \end{cases}, \quad (2.3)$$

and is overloaded for finite trait allocations $|_M : \mathcal{T} \rightarrow \mathcal{T}_M$ as

$$t_N|_M(\tau) := \begin{cases} \sum_{\omega \in \mathbb{T}} \mathbb{1}(\omega|_M = \tau) \cdot t_N(\omega) & \tau \neq \emptyset \\ 0 & \tau = \emptyset \end{cases}. \quad (2.4)$$

Definition 2.4. A pair of trait allocations t_M of $[M]$ and t_N of $[N]$ with $M \leq N$ is said to be *consistent* if $t_N|_M = t_M$.

The consistency of two finite trait allocations allows us to define the notion of a consistent sequence of trait allocations. Such a sequence can be thought of as generated by the sequential process of data arriving; each data point adds its index to its assigned traits without modifying any previous index. For example, $(\{\{1\}, \{1\}\}, \{\{1, 2\}, \{1\}\}, \{\{1, 2\}, \{1, 3\}\}, \dots)$ is a valid beginning to an infinite sequence of trait allocations. The first datum expresses two traits with multiplicity 1, and the second and third each express a single one of those traits with multiplicity 1. As a counterexample, $(\{\{1, 1\}\}, \{\{1, 1\}\}, \{\{1, 3\}\}, \dots)$ is not a valid trait allocation sequence, as the third trait allocation is not consistent with either the first or second. This sequence does not correspond to building up the traits expressed by data in a sequence; when the third datum is observed, the traits expressed by the first are modified.

Definition 2.5. An *infinite trait allocation* $t_\infty = (t_N)$ is a sequence of trait allocations of $[N]$, $N = 1, 2, \dots$ for which

$$\forall N \in \mathbb{N}, \quad t_{N+1}|_N = t_N. \quad (2.5)$$

Note that since restriction is commutative ($\cdot|_K|_M = \cdot|_M|_K = \cdot|_K$ for $K \leq M$), Definition 2.5 implies that all pairs of elements of the sequence (t_N) are consistent. Restriction acts on infinite trait allocations in a straightforward way: given $t_\infty = (t_N)$, restriction to $M \in \mathbb{N}$ is equivalent to the corresponding projection, $t_\infty|_M := t_M$.

Denote the set of all infinite trait allocations $\mathcal{T}_\infty \subset \times_N \mathcal{T}_N$. Recall that the motivation for developing infinite trait allocations is to capture the latent combinatorial structure underlying a sequence of observed data. Since this sequence is random, its underlying structure may also be, and thus the next task is to develop a corresponding notion of a random infinite trait allocation. Given a sequence of probability spaces $(\mathcal{T}_N, 2^{\mathcal{T}_N}, \nu_N)$ for $N \in \mathbb{N}$ with consistent measures (ν_N) , i.e.

$$\forall N \in \mathbb{N}, \quad \nu_N(t_N) = \sum_{t_{N+1} \in \mathcal{T}_{N+1}} \mathbb{1}(t_{N+1}|_N = t_N) \cdot \nu_{N+1}(t_{N+1}), \quad (2.6)$$

the Kolmogorov extension theorem (Kallenberg, 1997, Theorem 5.16) guarantees the existence of a unique random infinite trait allocation T_∞ that satisfies $T_\infty \in \mathcal{T}_\infty$ a.s. and has finite marginal distributions equal to the ν_N induced by restriction, i.e.

$$\forall N \in \mathbb{N}, \quad T_\infty|_N \sim \nu_N. \quad (2.7)$$

The properties of the random infinite trait allocation T_∞ are intimately related to those of the observed sequence of data it represents. In many applications, the data sequence has the property that its distribution is invariant to finite permutation of its elements; in some sense, the order in which the data sequence is observed is immaterial. We expect the random infinite trait allocation T_∞ associated with such an *infinite exchangeable sequence*¹ to inherit a similar property. As a simple illustration of the extension of permutation to infinite trait allocations, suppose we observe the sequence of data (x_1, x_2, x_3, \dots) exhibiting trait allocation sequence $T_1 = \{\{1, 1\}\}$, $T_2 = \{\{1, 1, 2\}, \{2\}\}$, $T_3 = \{\{1, 1, 2\}, \{2\}, \{3, 3\}\}$, and so on. If we swap x_1 and x_2 in the data sequence—resulting in the new sequence (x_2, x_1, x_3, \dots) —the traits expressed by x_2 become those containing index 1, the traits for x_1 become those containing index 2, and the rest are unchanged. Therefore, the permuted infinite trait allocation is $T'_1 = \{\{1\}, \{1\}\}$, $T'_2 = \{\{2, 2, 1\}, \{1\}\}$, $T'_3 = \{\{2, 2, 1\}, \{1\}, \{3, 3\}\}$, and so on. Note that T'_1 (resp. T'_2) is equal to the restriction to 1 (resp. 2) of T_2 with permuted indices, while T'_N for $N \geq 3$ is T_N with its indices permuted. This demonstrates a crucial point—if the permutation affects only indices up to $M \in \mathbb{N}$ (there is always such an M for finite permutations), we can arrive at the sequence of trait allocations for the permuted data sequence in two steps. First, we permute the indices in T_M and then restrict to $1, 2, \dots, M$ to get the first M permuted finite trait allocations. Then we permute the indices in T_N for each $N > M$.

To make this observation precise, we let π be a *finite permutation* of the natural numbers, i.e.,

$$\pi : \mathbb{N} \rightarrow \mathbb{N}, \quad \pi \text{ is a bijection,} \quad \exists M \in \mathbb{N} : \forall m > M, \pi(m) = m, \quad (2.8)$$

and overload its notation to operate on traits and (in)finite trait allocations in Definition 2.6. Note that if π is a finite permutation, its inverse π^{-1} is also a finite permutation with the same value of $M \in \mathbb{N}$ for which $m > M$ implies $\pi(m) = m$. Intuitively, π operates on traits and finite trait allocations by permuting their indices. For example, if π has the cycle (123) and fixes all indices greater than 3, then $\pi\{1, 1, 2, 4\} = \{2, 2, 3, 4\}$.

Definition 2.6. Given a finite permutation of the natural numbers $\pi : \mathbb{N} \rightarrow \mathbb{N}$ that fixes all indices $m > M$, the permutation of a trait τ under π is defined as

$$\pi\tau(m) := \tau(\pi^{-1}(m)), \quad (2.9)$$

the permutation of a trait allocation t_N of $[N]$ under π is defined as

$$\pi t_N(\tau) := t_N(\pi^{-1}\tau), \quad (2.10)$$

and the permutation of an infinite trait allocation t_∞ under π is defined as

$$\pi t_\infty := \left((\pi t_{\max(M, N)}) \Big|_N \right)_{N=1}^\infty. \quad (2.11)$$

¹For an introduction to exchangeability and related theory, see Aldous (1985).




FIG 1. An example exchangeable trait allocation construction. For each $N \in \mathbb{N}$, the trait membership $\xi_N \in \mathbb{K}$ of index N is determined by sampling i.i.d. from the distribution $(\mu_\xi)_{\xi \in \mathbb{K}}$ (depicted by colored bars). The resulting (unordered) trait allocation for indices up to 4 is shown above. Here $\xi_1 = (1, 0, 2, 0, \dots)$, $\xi_2 = \xi_4 = (0, 0, 1, 0, \dots)$, and $\xi_3 = (1, 2, 0, 0, \dots)$.

As discussed above, the definition for infinite trait allocations ensures that the permuted infinite trait allocation is a consistent sequence that corresponds to rearranging the observed data sequence with the same permutation. Definition 2.6 provides the necessary framework for studying *infinite exchangeable trait allocations*, defined as random infinite trait allocations whose distributions are invariant to finite permutation.

Definition 2.7. An *infinite exchangeable trait allocation*, T_∞ , is a random infinite trait allocation such that for any finite permutation $\pi : \mathbb{N} \rightarrow \mathbb{N}$,

$$\pi T_\infty \stackrel{d}{=} T_\infty. \tag{2.12}$$

Note that if the random infinite trait allocation is a random infinite partition/feature allocation almost surely, the notion of exchangeability in Definition 2.7 reduces to earlier notions of exchangeability for random infinite partition/feature allocations (Kingman, 1978; Aldous, 1985; Broderick, Pitman and Jordan, 2013). Exchangeability also has an analogous definition for random *finite* trait allocations, though this is of less interest in the present work.

As a concrete example, consider the countable set \mathbb{K} of sequences of nonnegative integers $\xi \in (\{0\} \cup \mathbb{N})^\infty$ such that $\sum_k \xi_k < \infty$. For each data index, we will generate an element of \mathbb{K} and use it to represent a sequence of multiplicities in an ordered sequence of traits. In particular, we endow \mathbb{K} with probabilities μ_ξ for each $\xi \in \mathbb{K}$. We start from an empty ordered trait allocation. Then for each data index $N \in \mathbb{N}$, we sample a sequence $\xi_N \stackrel{i.i.d.}{\sim} (\mu_\xi)_{\xi \in \mathbb{K}}$; and for each $k \in \mathbb{N}$, we add index N to trait k with multiplicity ξ_{Nk} . The final trait allocation is the unordered collection of nonempty traits. Since each data index generates its membership in the traits i.i.d. conditioned on (μ_ξ) , the sequence of trait allocations is exchangeable. This process is depicted in Fig. 1. As we will show in Section 4, all infinite exchangeable trait allocations have a similar construction.

3. Ordered trait allocations and lexicographic ordering

We impose no inherent ordering on the traits in a finite trait allocation via the use of (multi)sets; the allocations $\{\{1\}, \{3, 3\}\}$ and $\{\{3, 3\}, \{1\}\}$ are identical. This correctly captures our lack of a preferred trait order in many data analysis problems. However, ordered trait allocations are nonetheless often useful from standpoints both practical—such as when we need to store a finite trait allocation in an array in physical memory—and theoretical—such as in developing the characterization of all infinite exchangeable trait allocations in Section 4.

A primary concern in the development of an ordering scheme is consistency. Intuitively, as we observe more data in the sequence, we want the sequence of finite ordered trait allocations to “grow” but not be “shuffled”; in other words, if two finite trait allocations are consistent, the traits in their ordered counterparts at the same index should each be consistent. For partitions, this task is straightforward: each trait receives as a label its lowest index (Aldous, 1985), and the labels are used to order the traits. This is known as the *order-of-appearance* labeling, as traits are labeled in the order in which they are instantiated by data in the sequence. For example, in the partition $t_4 = \{\{1, 3\}, \{2, 4\}\}$ of $[4]$, $\{1, 3\}$ would receive label 1 and $\{2, 4\}$ would receive label 2, so $\{1, 3\}$ would be before $\{2, 4\}$ in the order. Restricting these traits will never change their order—for instance, $\{1, 3\}|_2 = \{1\}$ and $\{2, 4\}|_2 = \{2\}$, which still each receive label 1 and 2, respectively. If a restriction leaves a trait empty, it is removed and does not interfere with any traits of a lower label. For finite feature allocations, this ordering is inapplicable, since multiple features may have a common lowest index. Instead, Griffiths and Ghahramani (2005) introduce a *left-ordered form* in which one feature precedes another if it contains an index n that the other does not, and all indices $0 < m < n$ have the same membership in both features. For example, $\{1, 2, 5\}$ precedes $\{1, 3, 5\}$ in this ordering, since the traits both have index 1, but only the first has index 2.² In this section, we show that the well-known *lexicographic ordering*—which generalizes these previous orderings for partitions and feature allocations—satisfies our desiderata for an ordering on traits. We begin by defining ordered trait allocations.

Definition 3.1. An *ordered trait allocation* ℓ_N of $[N]$ is a sequence $\ell_N = (\ell_{Nk})_{k=1}^K$, $K < \infty$, of traits $\ell_{Nk} \in \mathbb{T}$ such that no trait contains an index $n > N$.

Let \mathcal{L}_N be the set of ordered trait allocations of $[N]$, and let $\mathcal{L} = \bigcup_N \mathcal{L}_N$ be the set of all ordered finite trait allocations. As in the case of unordered trait allocations, the notion of consistency is intimately tied to that of restriction. We again require that restriction to $M \in \mathbb{N}$ removes all indices $m > M$, and removes all traits rendered empty by that process. However, we also require that the order of the remaining traits is preserved: for example, if $\ell_3 = (\{3\}, \{1, 2\}, \{2\}, \{1, 1, 2\})$, the restriction of ℓ_3 to 1 should yield $(\{1\}, \{1, 1\})$,

²Other past work (Broderick, Pitman and Jordan, 2013) uses auxiliary randomness to order features, but this technique does not guarantee that orderings of two consistent finite trait allocations t_N, t_M are themselves consistent.

not $(\{1, 1\}, \{1\})$. Definition 3.2 satisfies these desiderata, overloading the $|_M$ function again for notational brevity.

Definition 3.2. The restriction $|_M : \mathcal{L} \rightarrow \mathcal{L}_M$ of an ordered finite trait allocation ℓ_N to $M \in \mathbb{N}$ is defined as

$$\ell_N|_M := \text{filter} \left((\ell_{Nk}|_M)_{k=1}^K \right), \tag{3.1}$$

where the filter function removes any empty sets from a sequence while preserving the order of the nonempty sets.

In the example above, the basic restriction of ℓ_3 to 1 would yield $(\emptyset, \{1\}, \emptyset, \{1, 1\})$, which the filter function then processes to form $\ell_3|_1 = (\{1\}, \{1, 1\})$, as desired. Analogously to the unordered case, we say two ordered trait allocations ℓ_N, ℓ_M , of $[N], [M]$, with $M \leq N$, are consistent if $\ell_N|_M = \ell_M$, and define the set of infinite ordered trait allocations \mathcal{L}_∞ as the set of infinite sequences of ordered finite trait allocations with $\ell_{N+1}|_N = \ell_N \forall N \in \mathbb{N}$.

Given these definitions, we are now ready to make the earlier intuitive notion of a consistent trait ordering scheme precise. Definition 3.3 states that a function $[\cdot] : \mathcal{T} \rightarrow \mathcal{L}$ must satisfy two conditions to be a valid trait ordering. The first condition enforces that a trait ordering does not add, remove, or modify the traits in the finite trait allocation t_N ; this implies that trait orderings are injective. The second condition enforces that trait orderings commute with restriction; in other words, applying a trait ordering to a consistent sequence of finite trait allocations yields a consistent sequence of ordered finite trait allocations. For example, suppose $t_2 = \{\{2\}, \{1, 2\}\}$, $t_3 = \{\{2\}, \{1, 2\}, \{3\}\}$, and we are given a proposed trait ordering where $[t_2] = (\{2\}, \{1, 2\})$ and $[t_3] = (\{3\}, \{2\}, \{1, 2\})$. This would not violate either of the conditions and may be a valid trait ordering. If instead the ordering was $[t_3] = (\{3\}, \{1, 2\}, \{2\})$, the proposal would not be a valid trait ordering—the traits $\{2\}$ and $\{1, 2\}$ get “shuffled”, i.e., $[t_3|_2] = [t_2] = (\{2\}, \{1, 2\}) \neq (\{1, 2\}, \{2\}) = [t_3]_2$.

Definition 3.3. A trait ordering is a function $[\cdot] : \mathcal{T} \rightarrow \mathcal{L}$ such that:

1. The ordering is *exhaustive*: If $[t_N] = (\tau_k)_{k=1}^K$, then $t_N = \{\tau_1, \dots, \tau_K\}$.
2. The ordering is *consistent*: $[t_N|_M] = [t_N]_M$.

The trait ordering we use throughout is the *lexicographic ordering*: for two traits, we pick the lowest index with differing multiplicity, and order the one with higher multiplicity first. For example, $\{1, 1, 4\} < \{1, 2\}$ since 1 is the lowest index with differing multiplicity, and the multiplicity of 1 is greater in the first trait than in the second. Similarly, $\{2, 3\} < \{2, 4\}$ since 3 has greater multiplicity in the first trait than the second, and both 1 and 2 have the same multiplicity in both traits. Definition 3.4 makes this precise.

Definition 3.4. For two traits $\tau, \omega \in \mathbb{T}$, we say that $\tau < \omega$ if there exists $n \in \mathbb{N}$ such that $\tau(n) > \omega(n)$ and all $m \in [n - 1]$ satisfy $\omega(m) = \tau(m)$.

We define $[\cdot] : \mathcal{T} \rightarrow \mathcal{L}$ as the mapping from t_N to the ordered trait allocation ℓ_N induced by the lexicographic ordering. The mapping $[\cdot]$ is a trait ordering, as shown by Theorem 3.6. The proof of Lemma 3.5 is provided in Appendix A.

Lemma 3.5. *For any pair $\tau, \omega \in \mathbb{T}$, if $\tau \leq \omega$ then $\tau|_M \leq \omega|_M$ for all $M \in \mathbb{N}$.*

Theorem 3.6. *The mapping $[\cdot]$ is a trait ordering.*

Proof. $[\cdot]$ is trivially exhaustive: since the restriction operation $\cdot|_M$ acts identically to individual traits in both ordered and unordered finite trait allocations, and empty traits are removed, both $[t_N|_M]$ and $[t_N]|_M$ have the same multiset of traits (albeit in a potentially different order). The first trait τ of $[t_N]$ satisfies $\tau \leq \omega$ for any $\omega \in \mathbb{T}$ such that $t_N(\omega) > 0$, by definition of $[\cdot]$. By Lemma 3.5, this implies that $\tau|_M \leq \omega|_M$ for all $\omega \in t_N$. Therefore, the first trait in $[t_N]|_M$ is the same as the first trait in $[t_N|_M]$. Applying this logic recursively to t_N with τ removed, the result follows. \square

4. De Finetti representation of exchangeable trait allocations

We now derive a de Finetti-style representation theorem for infinite exchangeable trait allocations (Definition 2.7) that extends previous results for partitions and feature allocations (Kingman, 1978; Broderick, Pitman and Jordan, 2013). It turns out that *all* infinite exchangeable trait allocations have essentially the same form as in the example construction at the end of Section 2, with some additional nuance.

The high-level proof sketch is as follows. We first use the lexicographic ordering from Section 3 to associate an i.i.d. sequence of uniform random labels to the traits in the sequence, in the style of Aldous (1985). We collect the multiset of labels for each index into a sequence, called the *label multiset sequence*; the consistency of the ordering from Theorem 3.6 implies that this construction is well-defined. We show that the label multiset sequence itself is exchangeable in the traditional sequential sense in Lemma 4.3. And we use de Finetti's theorem (Kallenberg, 1997, Theorem 9.16) to uncover its construction from conditionally i.i.d. random quantities. Finally, we relate this construction back to the original set of infinite exchangeable trait allocations to arrive at its representation in Theorem 4.5. Throughout the remainder of the paper, $T_\infty := (T_N)$ is a random infinite trait allocation and $\phi_\infty := (\phi_k) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$.

As an example construction of the label multiset sequence, suppose we have $T_4 = \{\{1, 2, 2\}, \{2, 4\}\}$, and $\phi_\infty := (\phi_k) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$. The lexicographic ordering of T_4 is $[T_4] = (\{1, 2, 2\}, \{2, 4\})$. The first trait in the ordering $\{1, 2, 2\}$ receives the first label in the sequence, ϕ_1 , and the second trait $\{2, 4\}$ receives the second label, ϕ_2 . For each index $n \in [4]$, we now collect the multiset of labels to its assigned traits with the same multiplicity. Index 1 is a member of only the first trait with multiplicity 1, so its label multiset is $\{\phi_1\}$. Index 2 is a member of the first trait with multiplicity 2 and the second with multiplicity 1, so its label multiset is $\{\phi_1, \phi_1, \phi_2\}$. Similarly, for index 3 it is \emptyset , and for index 4 it is $\{\phi_2\}$. Putting the multisets in order (for index 1, then 2, 3, etc.), the label

multiset sequence is therefore $(\{\phi_1\}, \{\phi_1, \phi_1, \phi_2\}, \emptyset, \{\phi_2\}, \dots)$, where the ellipsis represents the continuation beyond T_4 to T_5, T_6 , and so on. While the ϕ_k may be seen as a mathematical convenience for the proof, an alternative interpretation is that they correspond to trait-specific parameters in a broader Bayesian hierarchical model. Indeed, our proof would hold for ϕ_k from any nonatomic distribution, not just the uniform. In the document modeling example, each ϕ_k could correspond to a distribution over English words; ϕ_k with high mass on “basketball”, “luge”, and “curling” could represent a “sports” topic. For this reason, we call the ϕ_k *labels*. Let the set of (possibly empty) finite multisets of $(0, 1)$ be denoted \mathbb{Y} .

Definition 4.1. The *label multiset sequence* $Y_\infty := (Y_N)$ of elements $Y_N \in \mathbb{Y}$ corresponding to T_∞ and ϕ_∞ is defined by

$$Y_N(\phi) := \sum_k \mathbb{1}(\phi = \phi_k) \cdot [T_N]_k(N). \tag{4.1}$$

In other words, Y_N is constructed by selecting the N^{th} component of T_∞ , ordering its traits τ_1, \dots, τ_K , and then adding $\tau_k(N)$ copies of ϕ_k to Y_N for each $k \in [K]$. Again, the ϕ_k can thus be thought of as labels for the traits, and Y_N is the multiset of labels representing the assignment of the N^{th} datum to its traits (hence the name *label multiset sequence*). This construction of Y_∞ ensures that the “same label applies to the same trait” as N increases: the a.s. consistency of the ordering $[\cdot]$ introduced in Section 3 immediately implies that

$$\forall N \leq M, \quad Y_N(\phi) \stackrel{\text{a.s.}}{=} \sum_k \mathbb{1}(\phi = \phi_k) \cdot [T_M]_k(N). \tag{4.2}$$

Definition 4.1 implicitly creates a mapping, which we denote $\varphi : \mathcal{T}_\infty \times (0, 1)^\infty \rightarrow \mathbb{Y}^\infty$. Since the ϕ_k are distinct a.s., we can partially invert φ to recover the infinite trait allocation T_∞ corresponding to Y_∞ a.s. via

$$T_N(\tau) \stackrel{\text{a.s.}}{=} \mathbb{1}(\forall n > N, \tau(n) = 0) \cdot |\{\phi \in (0, 1) : \forall n \leq N, \tau(n) = Y_n(\phi)\}|. \tag{4.3}$$

The first term in the product—the indicator function—ensures that $T_N(\tau)$ is nonzero only for traits $\tau \in \mathbb{T}$ that do not contain any index $n > N$. The second term counts the number of points $\phi \in (0, 1)$ for which the multiplicities in τ match those expressed by the label multiset sequence for $n \leq N$. Thus, there exists another mapping $\tilde{\varphi} : \mathbb{Y}^\infty \rightarrow \mathcal{T}_\infty$ such that

$$\tilde{\varphi}(\varphi(T_\infty, \phi_\infty)) \stackrel{\text{a.s.}}{=} T_\infty. \tag{4.4}$$

The existence of the partial inverse $\tilde{\varphi}$ is a crucial element in the characterization of all distributions on infinite exchangeable trait allocations in Theorem 4.5. In particular, it guarantees that the distributions over random infinite trait allocations are in bijection with the distributions on label multiset sequences \mathbb{Y}^∞ , allowing the characterization of those on \mathbb{Y}^∞ (a much simpler space) instead. As the primary focus of this work is infinite exchangeable trait allocations, we

therefore must deduce the particular family of distributions on \mathbb{Y}^∞ that are in bijection with the infinite exchangeable trait allocations on \mathcal{T}_∞ .

Lemma 4.3 shows that this family is, as one might suspect, the exchangeable (in the classical, sequential sense) label multiset sequences. The main result required for its proof is Lemma 4.2, which states that permutation of T_∞ essentially results in the same permutation of the components of Y_∞ , modulo reordering the labels in ϕ_∞ . In other words, permuting the data sequence represented by T_∞ leads to the same permutation of Y_∞ . As an example, consider a setting in which $T_4 = \{\{1, 3, 4\}, \{2\}, \{2\}\}$, $\phi_\infty = (0.5, 0.4, 0.8, \dots)$, and thus $Y_\infty = (\{0.5\}, \{0.4, 0.8\}, \{0.5\}, \{0.5\}, \dots)$. For a finite permutation π , we define $\pi Y_\infty := (Y_{\pi^{-1}(N)})$ and $\pi \phi_\infty := (\phi_{\pi^{-1}(k)})$, i.e., permutations act on sequences by reordering elements. If we permute the observed data sequence that T_4 represents by $\pi = (12)(34)$, this leads to the permutation of the indices in T_4 also by π , resulting in $\pi T_4 = \{\{2, 3, 4\}, \{1\}, \{1\}\}$. If we then reorder ϕ_∞ with a different permutation $\pi' = (213)$, so $\pi' \phi_\infty = (0.4, 0.8, 0.5, \dots)$, then the corresponding label multiset sequence is $Y'_\infty = (\{0.4, 0.8\}, \{0.5\}, \{0.5\}, \{0.5\}, \dots)$. This Y'_∞ is the reordering of Y_∞ by π , the same permutation that was used to reorder the observed data; the main result of Lemma 4.2 is that a π' *always exists* to reorder ϕ_∞ such that this is the case. The proof of Lemma 4.2 may be found in Appendix A.

Lemma 4.2. *For each finite permutation π and infinite trait allocation t_∞ , there exists a finite permutation π' such that*

$$\pi \varphi(t_\infty, \phi_\infty) \stackrel{a.s.}{=} \varphi(\pi t_\infty, \pi' \phi_\infty). \quad (4.5)$$

Lemma 4.3. *T_∞ is exchangeable iff $Y_\infty = \varphi(T_\infty, \phi_\infty)$ is exchangeable.*

Proof. Fix a finite permutation π . Then by Lemma 4.2 there exists a collection of finite permutations π_{T_∞} that depend on T_∞ such that

$$\pi Y_\infty \stackrel{a.s.}{=} \varphi(\pi T_\infty, \pi_{T_\infty} \phi_\infty). \quad (4.6)$$

If Y_∞ is exchangeable, then using Eq. (4.6) and the definition of $\tilde{\varphi}$ in Eq. (4.4),

$$T_\infty \stackrel{a.s.}{=} \tilde{\varphi}(Y_\infty) \stackrel{d}{=} \tilde{\varphi}(\pi Y_\infty) \stackrel{a.s.}{=} \pi T_\infty. \quad (4.7)$$

If T_∞ is exchangeable, then again using Eq. (4.6) and noting that ϕ_∞ is a sequence of i.i.d. random variables and hence also exchangeable,

$$\pi Y_\infty \stackrel{a.s.}{=} \varphi(\pi T_\infty, \pi_{T_\infty} \phi_\infty) \stackrel{d}{=} \varphi(T_\infty, \phi_\infty) = Y_\infty. \quad (4.8)$$

□

We are now ready to characterize all distributions on infinite exchangeable trait allocations in Theorem 4.5 using the *de Finetti representation* provided by Definition 4.4. At a high level, this is a constructive representation involving three steps. Recall that \mathbb{K} is the countable set of sequences of nonnegative

integers (ξ_k) such that $\sum_k \xi_k < \infty$. First, we generate a (possibly random) distribution over \mathbb{K}^2 , i.e., a sequence $(\mu_{\xi, \xi'})_{\xi, \xi' \in \mathbb{K}}$ of nonnegative reals such that

$$\sum_{\xi, \xi' \in \mathbb{K}} \mu_{\xi, \xi'} = 1 \quad \text{and} \quad \forall \xi, \xi' \in \mathbb{K}, \quad \mu_{\xi, \xi'} \geq 0. \tag{4.9}$$

Next, for each $N \in \mathbb{N}$, we sample i.i.d. from this distribution, resulting in two sequences ξ_N, ξ'_N . The sequence ξ_N determines the membership of index N in *regular* traits—which may be joined by other indices—and ξ'_N determines its membership in *dust* traits—which are unique to index N and will never be joined by any other index. In particular, for each $k \in \mathbb{N}$, index N joins trait k with multiplicity ξ_{Nk} ; and for each $j \in \mathbb{N}$, index N has ξ'_{Nj} additional unique traits of multiplicity j . For example, in a sequence of documents generated by latent topics, one author may write a single document with a number of words that are never again used by other authors (e.g. *Jabberwocky*, by Lewis Carroll); in the present context, these words would be said to arise from a dust topic. Meanwhile, common collections of words expressed by many documents will group together to form regular topics. Finally, we associate each trait with an i.i.d. $\text{Unif}(0, 1)$ label, construct the label multiset sequence Y_∞ , and use our mapping $\tilde{\varphi}$ to collect these results together to form an infinite trait allocation T_∞ . We say a random infinite trait allocation is *regular* if it has no dust traits with probability 1, and *irregular* otherwise.

Definition 4.4. A random infinite trait allocation T_∞ has a *de Finetti representation* if there exists a random distribution $(\mu_{\xi, \xi'})$ on \mathbb{K}^2 such that T_∞ has distribution induced by the following construction:

1. generate $(\phi_k), (\phi_{Nj\ell}) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$ and $(\xi_N, \xi'_N) \stackrel{\text{i.i.d.}}{\sim} (\mu_{\xi, \xi'})$,
2. for all $N \in \mathbb{N}$, define the multisets R_N, D_N, Y_N of $(0, 1)$ via

$$R_N(\phi) = \sum_{k, j} \mathbf{1}(\phi = \phi_k, \xi_{Nk} = j) \cdot j \quad (\text{regular traits}) \tag{4.10}$$

$$D_N(\phi) = \sum_{j, \ell} \mathbf{1}(\phi = \phi_{Nj\ell}, \ell \leq \xi'_{Nj}) \cdot j \quad (\text{dust traits}) \tag{4.11}$$

$$Y_N(\phi) = R_N(\phi) + D_N(\phi), \tag{4.12}$$

3. assemble the label multiset sequence $Y_\infty = (Y_N)$ and set $T_\infty = \tilde{\varphi}(Y_\infty)$.

Theorem 4.5 is the main result of this section, which shows that infinite exchangeable trait allocations—both regular and irregular—are precisely those which have a de Finetti representation per Definition 4.4. The proof of Theorem 4.5 approaches the problem by characterizing the distribution of the exchangeable label multiset sequence Y_∞ .

Theorem 4.5. T_∞ is exchangeable iff it has a de Finetti representation.

Proof. If T_∞ has a de Finetti representation, then it is exchangeable by the fact that the ξ_N, ξ'_N are i.i.d. random variables. In the other direction, if T_∞ is

exchangeable, then there is a random label multiset sequence $Y_\infty = \varphi(T_\infty, \phi_\infty)$ which is exchangeable by Lemma 4.3. Since we can recover T_∞ from Y_∞ via $T_\infty = \tilde{\varphi}(Y_\infty)$, it suffices to characterize Y_∞ and then reconstruct T_∞ .

We split Y_N into its regular R_N and dust D_N components—that represent, respectively, traits that are expressed by multiple data points and those that are expressed only by data point N —defined for $\phi \in (0, 1)$ by

$$D_N(\phi) = \begin{cases} 0 & \exists M \neq N : Y_M(\phi) > 0 \\ Y_N(\phi) & \text{otherwise} \end{cases} \quad (4.13)$$

$$R_N(\phi) = Y_N(\phi) - D_N(\phi). \quad (4.14)$$

Choose any ordering (ϕ_k) on the countable set $\{\phi \in (0, 1) : \sum_N R_N(\phi) > 0\}$. Next, we extract the multiplicities in R_N and D_N via the sequences $\xi_N, \xi'_N \in \mathbb{K}$,

$$\xi'_{Nj} := |\{\phi \in (0, 1) : D_N(\phi) = j\}| \quad \xi_{Nk} := R_N(\phi_k). \quad (4.15)$$

Note that we can recover the distribution of Y_∞ from that of $(\xi_N, \xi'_N)_{N=1}^\infty$ by generating sequences $(\phi'_k), (\phi'_{Nj\ell}) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$ and using steps 2 and 3 of Definition 4.4. Therefore it suffices to characterize the distribution of $(\xi_N, \xi'_N)_{N=1}^\infty$. Note that $(\xi_N, \xi'_N)_{N=1}^\infty$ is a function of Y_∞ such that permuting the elements of Y_∞ corresponds to permuting those of $(\xi_N, \xi'_N)_{N=1}^\infty$ in the same way. Thus since Y_∞ is exchangeable, so is $(\xi_N, \xi'_N)_{N=1}^\infty$. And since $(\xi_N, \xi'_N)_{N=1}^\infty$ is a sequence in a Borel space, de Finetti's theorem (Kallenberg, 1997, Theorem 9.16) states that there exists a directing random measure μ such that $(\xi_N, \xi'_N) \stackrel{\text{i.i.d.}}{\sim} \mu$. Since the set \mathbb{K}^2 is countable, we can represent μ with a probability $\mu_{\xi, \xi'}$ for each tuple $(\xi, \xi') \in \mathbb{K}^2$. \square

The representation in Theorem 4.5 generalizes de Finetti representations for both clustering (the Kingman paintbox) and feature allocation (the feature paintbox) (Kingman, 1978; Broderick, Pitman and Jordan, 2013), as shown by Corollaries 4.7 and 4.8. Further, Corollary 4.8 is the first de Finetti representation for feature allocations that accounts for the possibility of dust features; previous results were limited to regular feature allocations (Broderick, Pitman and Jordan, 2013). Theorem 4.5 also makes the distinction between regular and irregular trait allocations straightforward, as shown by Corollary 4.6.

Corollary 4.6. *An exchangeable trait allocation T_∞ is regular iff it has a de Finetti representation where $\mu_{\xi, \xi'} > 0$ implies $\sum_k \xi'_k = 0$.*

Corollary 4.7. *A partition T_∞ is exchangeable iff it has a de Finetti representation where $\mu_{\xi, \xi'} > 0$ implies either*

- $\sum_k \xi_k = 1$ and $\sum_k \xi'_k = 0$, or
- $\sum_k \xi_k = 0$, $\xi'_1 = 1$, and $\sum_k \xi'_k = 1$.

Corollary 4.8. *A feature allocation T_∞ is exchangeable iff it has a de Finetti representation where $\mu_{\xi, \xi'} > 0$ implies that*

- $\forall k \in \mathbb{N}, \xi_k \leq 1$, and $\forall j > 1, \xi'_j = 0$.

5. Frequency models and probability functions

The set of infinite exchangeable trait allocations encompasses a very expressive class of random infinite trait allocations: membership in different regular traits at varying multiplicities can be correlated, membership in dust traits can depend on membership in regular traits, etc. While interesting, this generality makes constructing models with efficient posterior inference procedures difficult. A simplifying assumption one can make is that given the directing measure μ , the membership of an index in a particular trait is independent of its membership in other traits. This assumption is often acceptable in practice, and limits the infinite exchangeable trait allocations to a subset—which we refer to as *frequency models*—for which efficient inference is often possible. Frequency models, as used in the present context, generalize the notion of a feature frequency model (Broderick, Pitman and Jordan, 2013) for feature allocations.

At a high level, this constructive representation consists of three steps. First, we generate random sequences of nonnegative reals (θ_{kj}) and (θ'_j) such that $\sum_{k,j} \theta_{kj} < \infty$, $\sum_j \theta'_j < \infty$, and $\forall k \in \mathbb{N}$, $\sum_j \theta_{kj} \leq 1$. The quantity θ_{kj} is the probability that an index joins regular trait k with multiplicity j , while θ'_j is the average number of dust traits of multiplicity j for each index. Next, each index $N \in \mathbb{N}$ independently samples its multiplicity ξ_{Nk} in regular trait k from the discrete distribution $(\theta_{kj})_{j=0}^{\infty}$, where $\theta_{k0} := 1 - \sum_j \theta_{kj}$ is the probability that the index is not a member of trait k . For each $j \in \mathbb{N}$, each index $N \in \mathbb{N}$ is a member of an additional $\xi'_{Nj} \stackrel{\text{indep}}{\sim} \text{Pois}(\theta'_j)$ dust traits of multiplicity j . Finally, we collect these results together to form an infinite trait allocation T_∞ . Note that the above essentially imposes a particular form for μ , as given by Definition 5.1.

Definition 5.1. A random infinite trait allocation T_∞ has a *frequency model* if there exist two random sequences (θ_{kj}) , (θ'_j) of nonnegative real numbers such that T_∞ has a de Finetti representation with

$$\mu_{\xi, \xi'} = \left(\prod_{k=1}^{\infty} \theta_{k\xi_k} \right) \cdot \left(\prod_{j=1}^{\infty} \frac{(\theta'_j)^{\xi'_j} e^{-\theta'_j}}{\xi'_j!} \right). \quad (5.1)$$

Although considerably simpler than general infinite exchangeable trait allocations, this representation still involves a potentially infinite sequence of parameters; a finitary representation would be more useful for computational purposes. In practice, the marginal distribution of T_N provides such a representation (Griffiths and Ghahramani, 2005; Thibaux and Jordan, 2007; James, 2017; Broderick, Wilson and Jordan, 2018). So rather than considering a simplified class of de Finetti representations, we can alternatively consider a simplified class of marginal distributions for T_N . In previous work on feature allocations (Broderick, Pitman and Jordan, 2013), the analog of frequency models was shown to correspond to those marginal distributions that depend only on the unordered feature sizes (the so-called *exchangeable feature probability functions (EFPPs)*). In the following, we develop the generalization of EFPPs for trait allocations

and show that the same correspondence result holds in this generalized framework.

We let $\kappa(t_N)$ be the number of unique orderings of a trait allocation t_N ,

$$\kappa(t_N) := \frac{(\sum_{\tau \in \mathbb{T}} t_N(\tau))!}{\prod_{\tau \in \mathbb{T}} t_N(\tau)!}, \quad (5.2)$$

and use the *multiplicity profile*³ of t_N , given by Definition 5.2, to capture the multiplicities of indices in its traits. The multiplicity profile of a trait is defined to be the multiset of multiplicities of its elements, while the multiplicity profile of a finite trait allocation is the multiset of multiplicity profiles of its traits. As an example, the multiplicity profile of a trait $\{1, 3, 4, 4, 2, 2, 2\}$ is $\{1, 1, 2, 4\}$, since there are two elements of multiplicity 1, one element of multiplicity 2, and one of multiplicity 4 in the trait. If we are given the finite trait allocation $\{\{1, 1, 2\}, \{2\}, \{3\}, \{3, 3, 3, 3, 1\}\}$, then its multiplicity profile is $\{\{1, 2\}, \{1\}, \{1\}, \{1, 4\}\}$. Note that a multiplicity profile is itself a trait allocation, though not always of the same indices. Here, the trait allocation is of $\{3\}$, and its multiplicity profile is a trait allocation of $\{4\}$.

Definition 5.2. The *multiplicity profile* $\bar{\tau} : \mathbb{T} \rightarrow \mathbb{T}$ of a trait $\tau \in \mathbb{T}$ is defined as

$$\bar{\tau}(n) := |\{m \in \mathbb{N} : \tau(m) = n\}|, \quad (5.3)$$

and is overloaded for finite trait allocations $\bar{\tau} : \mathcal{T} \rightarrow \mathcal{T}$ as

$$\bar{t}_N(\xi) := \sum_{\tau \in \mathbb{T}} \mathbf{1}(\bar{\tau} = \xi) \cdot t_N(\tau). \quad (5.4)$$

We also extend Definition 5.2 to ordered trait allocations ℓ_N , where the multiplicity profile is the ordered multiplicity profiles of its traits, i.e. $\bar{\ell}_N$ is defined such that $\forall k \in \mathbb{N}, \bar{\ell}_{Nk} := \bar{\ell}_{Nk}$.

The precise simplifying assumption on the marginal distribution of T_N that we employ in this work is provided in Definition 5.3, which generalizes past work on exchangeable probability functions (Pitman, 1995; Broderick, Pitman and Jordan, 2013).

Definition 5.3. A random infinite trait allocation T_∞ has an *exchangeable trait probability function (ETPF)* if there exists a function $p : \mathbb{N} \times \mathcal{T} \rightarrow \mathbb{R}_+$ such that for all $N \in \mathbb{N}$,

$$\mathbb{P}(T_N = t_N) = \kappa(t_N) \cdot p(N, \bar{t}_N). \quad (5.5)$$

One of the primary goals of this section is to relate infinite exchangeable trait allocations with frequency models to those with ETPFs. The main result of this section, Theorem 5.4, shows that these two assumptions are actually equivalent:

³A very similar quantity is known in the population genetics literature as the *site* (or *allele*) *frequency spectrum* (Bustamante et al., 2001), though it is typically defined there as an ordered sequence or vector rather than as a multiset.

any random infinite trait allocation T_∞ that has a frequency model (including those with random $(\theta_{kj}), (\theta'_j)$ of arbitrary distribution) has an ETPF, and any random infinite trait allocation with an ETPF has a frequency model. Therefore, we are able to use the simple construction of frequency models in practice via their associated ETPFs.

Theorem 5.4. T_∞ has a frequency model iff it has an ETPF.

The key to the proof of Theorem 5.4 is the *uniformly ordered infinite trait allocation*, defined below in Definition 5.6. Recall that \mathcal{L}_∞ is the space of consistent, ordered infinite trait allocations and that L_∞ denotes an ordering of T_∞ . Here, we develop the *uniform ordering* L_∞ : intuitively, for each $N \in \mathbb{N}$, L_{N+1} is constructed by inserting the new traits in T_{N+1} relative to T_N into uniformly random positions among the elements of L_N . This guarantees that L_N is marginally a uniform random permutation of $[T_N]$ for each $N \in \mathbb{N}$, and that L_∞ is a consistent sequence, i.e. $L_\infty \in \mathcal{L}_\infty$. There are two advantages to analyzing L_∞ rather than T_∞ itself. First, the ordering removes the combinatorial difficulties associated with analyzing T_∞ . Second, the traits are independent of their ordering, thereby avoiding the statistical coupling of the ordering based solely on $[\cdot]$.

The definition of the uniform ordering L_∞ in Definition 5.6 is based on associating traits with the uniformly distributed i.i.d. sequence ϕ_∞ , and ordering the traits based on the order of those values. To do so, we require a definition of the finite permutation π_n that rearranges the first n elements of ϕ_∞ to be in order and leaves the rest unchanged, known as the *n^{th} order mapping* π_n of ϕ_∞ . For example, if $\phi_\infty = (0.4, 0.1, 0.3, 0.2, 0.5, \dots)$, then π_3 is represented in cycle notation as (321) , and $\pi_3\phi_\infty = (0.1, 0.3, 0.4, 0.2, 0.5, \dots)$. The precise formulation of this notion is provided by Definition 5.5.

Definition 5.5. The *n^{th} order mapping* $\pi_n : \mathbb{N} \rightarrow \mathbb{N}$ of the sequence ϕ_∞ is the finite permutation defined by

$$\pi_n(k) := \begin{cases} |\{j \in \mathbb{N} : j \leq n, \phi_j \leq \phi_k\}| & k \leq n \\ k & k > n \end{cases} \quad (5.6)$$

Definition 5.6 shows how to use the n^{th} order mapping to uniformly order an infinite trait allocation: we rearrange the lexicographic ordering of T_N using the K_N^{th} order mapping π_{K_N} where K_N is the number of traits in T_N .

Definition 5.6. The *uniform ordering* $L_\infty := (L_N)$ of T_∞ is

$$L_{Nk} := [T_N]_{\rho_N^{-1}(k)}, \quad (5.7)$$

where $\rho_N := \pi_{K_N}$ and $K_N = \sum_{\tau \in \mathbb{T}} T_N(\tau)$ is the number of traits in T_N .

Note that we can also define the *uniformly ordered label multiset sequence* $Y_\infty = (Y_N) \in \mathbb{Y}^\infty$ from the uniform ordering L_∞ of T_∞ via

$$Y_N(\phi) := \sum_k L_{Nk}(N) \cdot \mathbf{1} \left(\phi = \phi_{\rho_N^{-1}(k)} \right), \quad (5.8)$$

and recover the original infinite random trait allocation $T_\infty \stackrel{a.s.}{=} \tilde{\varphi}(Y_\infty)$ from the mapping $\tilde{\varphi}$ in Eq. (4.4).

The proof of Theorem 5.4 relies on Lemma 5.7, a collection of two technical results associated with uniformly ordered infinite trait allocations L_∞ for which the associated unordered infinite trait allocation T_∞ has an ETPF. The first result states that L_N and $\overline{L_{N+k}}$ are conditionally independent given $\overline{L_N}$ for any $N, k \in \mathbb{N}$; essentially, if the distribution of L_N depends only on its multiplicity profile, knowing the multiplicity profiles of further uniformly ordered trait allocations in the sequence L_∞ does not provide any extra useful information about L_N . The second result states that the distribution of L_N conditioned on $\overline{L_N}$ is uniform. The proof of Lemma 5.7 may be found in Appendix A.

Lemma 5.7. *If T_∞ has an ETPF, and L_∞ is the uniform ordering of T_∞ , then for all $N \in \mathbb{N}$, $\ell_N \in \mathcal{L}_N$,*

$$\mathbb{P}(L_N = \ell_N \mid \overline{L_N}, \overline{L_{N+1}}, \overline{L_{N+2}}, \dots) = \mathbb{P}(L_N = \ell_N \mid \overline{L_N}) \quad a.s., \quad (5.9)$$

and $\mathbb{P}(L_N = \cdot \mid \overline{L_N})$ is a uniform distribution over the ordered trait allocations of $[N]$ consistent with $\overline{L_N}$.

Proof of Theorem 5.4. Let $L_\infty := (L_N)$ be the uniform ordering of $T_\infty := (T_N)$. For any $N \in \mathbb{N}$, $\ell_N \in \mathcal{L}_N$, and $t_N \in \mathcal{T}_N$ such that ℓ_N is an ordering of t_N ,

$$\mathbb{P}(L_N = \ell_N) = \sum_{t'_N \in \mathcal{T}_N} \mathbb{P}(L_N = \ell_N \mid T_N = t'_N) \mathbb{P}(T_N = t'_N) \quad (5.10)$$

$$= \mathbb{P}(L_N = \ell_N \mid T_N = t_N) \mathbb{P}(T_N = t_N) \quad (5.11)$$

$$= \kappa(t_N)^{-1} \mathbb{P}(T_N = t_N), \quad (5.12)$$

where the sum collapses to a single term since $t_N \in \mathcal{T}_N$ is the unique unordered version of ℓ_N , and $\mathbb{P}(L_N = \ell_N \mid T_N = t_N) = \kappa(t_N)^{-1}$ since L_N is uniformly distributed over the possible orderings of T_N . Thus

$$\mathbb{P}(T_N = t_N) = \kappa(t_N) \cdot \mathbb{P}(L_N = \ell_N). \quad (5.13)$$

Suppose T_∞ has a frequency model as in Definition 5.1. To show T_∞ has an ETPF, it remains to show that there exists a function p such that

$$\mathbb{P}(L_N = \ell_N) = p(N, \overline{t_N}). \quad (5.14)$$

The major difficulty in doing so is that there is ambiguity in how $L_N = \ell_N$ was generated from the frequency model; any trait ℓ_{Nk} for which $\overline{\ell_{Nk}}$ is a singleton (i.e., ℓ_{Nk} contains a single unique index) may correspond to *either* a dust or regular trait. Therefore, we must condition on both the frequency model parameters and the (random) dust/regular assignments of the K traits in ℓ_N . We let $A_j \subset [K]$, $j \in \mathbb{N}$ be the set of components of ℓ_N corresponding to dust traits of multiplicity j . We further let Q be the set of sequences (A_j) such that $k \in A_j \implies \overline{\ell_{Nk}} = \{j\}$ for all $k, j \in \mathbb{N}$, i.e., those that are possible dust/regular

assignments of the traits given ℓ_N . Note in particular that Q is a function of $\overline{\ell}_N$ but not ℓ_N . Then by the tower property,

$$\mathbb{P}(L_N = \ell_N) = \mathbb{E} [\mathbb{P}(L_N = \ell_N \mid (A_j), (\theta_{kj}), (\theta'_j))] . \tag{5.15}$$

Expanding the inner conditional probability, and defining $A = [K] \setminus \bigcup_j A_j$,

$$\mathbb{P}(L_N = \ell_N \mid \dots) = \frac{\mathbb{1}((A_j) \in Q)}{N^{\sum_j |A_j|}} \cdot \prod_{k=1}^{\infty} \theta_{k0}^N \cdot \sum_{\substack{\sigma: A \rightarrow \mathbb{N} \\ \sigma \text{ 1-to-1}}} \prod_{k \in A} \prod_{j=1}^{\infty} \left(\frac{\theta_{\sigma(k)j}}{\theta_{\sigma(k)0}} \right)^{\overline{\ell}_{Nk(j)}} . \tag{5.16}$$

The first term in the product relates to the dust. Given that we know the positions and multiplicities of dust in L_N , the only remaining randomness is in which index expresses each dust trait; and since L_N has a uniformly random order, the probability of any index expressing dust at an index is $1/N$. The indicator expresses the fact that the probability of observing $L_N = \ell_N$ is 0 if it is inconsistent with the dust assignments (A_j) . The second and third terms are the sum over the probabilities of all ways the (θ_{kj}) could have generated the observed regular traits.

Note that the expression in Eq. (5.16) is a function of only N and $\overline{\ell}_N$, and therefore so is $\mathbb{P}(L_N = \ell_N)$ in Eq. (5.15). But since L_N is a uniformly ordered trait allocation, $\mathbb{P}(L_N = \ell_N)$ is invariant to reordering ℓ_N , so it is invariant to reordering $\overline{\ell}_N$; and since ℓ_N is some ordering of the traits in t_N , $\mathbb{P}(L_N = \ell_N)$ is a function of only \overline{t}_N and N . Therefore, there exists some function p such that

$$\mathbb{P}(L_N = \ell_N) = p(N, \overline{t}_N), \tag{5.17}$$

and T_∞ has an ETPF as required.

Next, assume T_∞ has an ETPF. Consider the finite subsequence $(Y_m)_{m=1}^M$ and σ -algebra $\mathcal{G}_N := \sigma(\rho_N \phi_\infty, \overline{L}_N)$, where $M \leq N$, and recall that $\rho_N \phi_\infty$ is the N^{th} ordering of ϕ_∞ , L_N is the uniform ordering of T_N , and \overline{L}_N is its multiplicity profile. Note that

$$\mathbb{P}((Y_m)_{m=1}^M \mid \mathcal{G}_N) \tag{5.18}$$

$$= \sum_{\ell_N \in \mathcal{L}_N} \mathbb{P}((Y_m)_{m=1}^M \mid \rho_N \phi_\infty, L_N = \ell_N) \mathbb{P}(L_N = \ell_N \mid \overline{L}_N) \tag{5.19}$$

$$= \sum_{\ell_N \in \mathcal{L}_N} \mathbb{P}((Y_m)_{m=1}^M \mid \rho_N \phi_\infty, L_N = \ell_N) \mathbb{P}(L_N = \ell_N \mid (\overline{L}_K)_{K=N}^\infty) \tag{5.20}$$

$$= \mathbb{P}((Y_m)_{m=1}^M \mid (\rho_K \phi_\infty, \overline{L}_K)_{K=N}^\infty) \tag{5.21}$$

almost surely, where the steps follow from the law of total probability, the measurability of \overline{L}_N with respect to $\sigma(L_N)$, Lemma 5.7, and the measurability of $\rho_{N+K} \phi_\infty$ with respect to $\sigma(\rho_N \phi_\infty)$ for any $K \in \mathbb{N}$. Therefore $\mathbb{P}((Y_m)_{m=1}^M \mid \mathcal{G}_N)$ is a reverse martingale in N , since $\sigma(\rho_K \phi_\infty, \overline{L}_K)_{K=N}^\infty$ is a reverse filtration; so by the reverse martingale convergence theorem (Kallenberg, 1997, Theorem 6.23), there exists a σ -algebra \mathcal{G} such that

$$\mathbb{P}((Y_m)_{m=1}^M \mid \mathcal{G}_N) \xrightarrow{a.s.} \mathbb{P}((Y_m)_{m=1}^M \mid \mathcal{G}) \quad N \rightarrow \infty. \tag{5.22}$$

We now study the properties of the limiting distribution. Denoting $Y_{mk} := Y_m(\phi_{\rho_N^{-1}(k)})$ for brevity, note that the uniform distribution of L_N conditioned on \overline{L}_N implies that

$$\mathbb{P}(Y_{1k} = j | (Y_m)_{m=2}^M, \mathcal{G}_N) = \frac{\overline{L}_{Nk}(j) - \sum_{m=2}^M \mathbb{1}(Y_{mk} = j)}{N - M + 1}, \quad j \in \mathbb{N} \cup \{0\} \tag{5.23}$$

independently across the trait indices $k \in \mathbb{N}$. Since $\frac{\sum_{m=2}^M \mathbb{1}(Y_{mk} = j)}{N - M + 1} \xrightarrow{a.s.} 0$ as $N \rightarrow \infty$, we have that $Y_1 \perp\!\!\!\perp (Y_m)_{m=2}^M | \mathcal{G}$. By symmetry, $(Y_m)_{m=1}^M$ are conditionally independent given \mathcal{G} . Since this holds for all finite subsequences, the result extends to the infinite sequence: Y_∞ is an i.i.d. sequence conditioned on \mathcal{G} . It thus suffices to characterize the limit of $\mathbb{P}(Y_1 | \mathcal{G}_N)$.

Define \mathcal{D}_{Nj} to be the set of indices for “dust-like” traits of multiplicity j , and \mathcal{R}_N to be the remaining component indices corresponding to nonempty “regular-like” traits,

$$\mathcal{D}_{Nj} = \{k \in \mathbb{N} : \overline{L}_{Nk} = \{j\}\}, \quad j \in \mathbb{N} \tag{5.24}$$

$$\mathcal{R}_N = \{k \in \mathbb{N} : \overline{L}_{Nk} \neq \emptyset\} \setminus \cup_j \mathcal{D}_{Nj}. \tag{5.25}$$

Simulating from $\mathbb{P}(Y_1 | \mathcal{G}_N)$ can be performed in two steps. First, independently for every $k \in \mathcal{R}_N$, we set Y_{1k} to $j \in \mathbb{N}$ with probability $\overline{L}_{Nk}(j)/N$, and to 0 with probability $1 - \sum_j \overline{L}_{Nk}(j)/N$. Then for each $j \in \mathbb{N}$, we generate $S_j \sim \text{Binom}(|\mathcal{D}_{Nj}|, 1/N)$, select a subset of \mathcal{D}_{Nj} of size S_j uniformly at random, and set Y_{1k} for each k in the subset to j . Given the almost-sure convergence of $\mathbb{P}(Y_1 | \mathcal{G}_N)$ as $N \rightarrow \infty$, the first step implies the existence of a countable sequence (ϕ'_k) in $(0, 1)$ (a rearrangement of some subset of the sequence ϕ_∞) and sequences of nonnegative reals $(\theta_{kj})_{j=0}^\infty$ such that

$$\theta_{kj} = \lim_{N \rightarrow \infty} \frac{\overline{L}_{Nk}(j)}{N}, \quad \theta_{k0} = 1 - \sum_j \theta_{kj}, \quad \mathbb{P}(Y_1(\phi'_k) = j | \mathcal{G}) = \theta_{kj} \tag{5.26}$$

independently across $k \in \mathbb{N}$. Using the law of small numbers (Ross, 2011, Theorem 4.6) on the binomial distribution for S_j (with shrinking probabilities $1/N$ as $N \rightarrow \infty$), and the fact that $\phi_\infty \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$, the second step implies that there exists a sequence of positive reals (θ'_j) such that

$$\theta'_j = \lim_{N \rightarrow \infty} |\mathcal{D}_{Nj}|/N, \tag{5.27}$$

where Y_1 additionally has $\text{Pois}(\theta'_j)$ unique elements uniformly distributed on $(0, 1)$ with multiplicity j . Finally, $\sum_j \theta_{kj} \leq 1$ by the above construction, and both $\sum_{k,j} \theta_{kj} < \infty$ and $\sum_j \theta'_j < \infty$ almost surely, since otherwise the second Borel–Cantelli lemma combined with the i.i.d. nature of Y_∞ conditioned on \mathcal{G} would imply that each Y_n is not a finite multiset, which contradicts the assumption that any index is a member of only finitely many traits almost surely. Thus $T_\infty = \tilde{\varphi}(Y_\infty)$ has a frequency model. \square

By setting $\theta_{kj} = \theta'_j = 0$ for all $k, j \in \mathbb{N} : j > 1$, Theorem 5.4 can be used to recover the correspondence between random infinite feature allocations with an exchangeable feature probability function (EFPF) and those with a feature frequency model, both defined in earlier work by Broderick, Pitman and Jordan (2013). In the present context, an EFPF is an ETPF where $p(N, \overline{t_N}) > 0$ only for t_N that are feature allocations. These are exactly the t_N for which $\overline{t_N}$ only contains traits τ of the form $\{1, 1, 1, \dots, 1\}$, i.e., $\overline{t_N}(\tau) > 0$ only if $\forall n > 1, \tau(n) = 0$.

Corollary 5.8. *A random infinite feature allocation has a feature frequency model iff it has an EFPF.*

For infinite exchangeable partitions, the result is stronger: *all* exchangeable infinite partitions have an exchangeable partition probability function (EPPF) (Pitman, 1995), defined as a summable symmetric function of the partition sizes times $K!$, where K is the number of partition elements. Theorem 5.4 cannot be directly used to recover this result: no choice of $(\theta_{kj}), (\theta'_j)$ in Definition 5.1 or $p(N, \overline{t_N})$ in Definition 5.3 guarantees that the resulting T_∞ is a partition. The key issue is that in trait allocations with frequency models, the membership of each index in the traits is independent across the traits, while in partitions each index is a member of exactly one trait. In the EPPF, this manifests itself as an indicator function that tests whether the traits exhibit a partition structure, where no such test exists in the ETPF (or EFPF, by extension).

As trait allocations generalize not only partitions, but other combinatorial structures with restrictions on index membership as well (cf. Section 6), it is of interest to find a generalization of the correspondence between frequency models and ETPFs that applies to these constrained structures. We thus require a way of extracting the memberships of a single index in a trait allocation—referred to as its *membership profile*, as in Definition 5.9—so that we can check whether it satisfies constraints on the combinatorial structure. For example, if we have the trait allocation $t_4 = \{\{1, 1, 2\}, \{1, 2, 3\}, \{1\}\}$, then the membership profile of index 1 is $\{1, 1, 2\}$, since index 1 is a member of two traits with multiplicity 1, and one trait with multiplicity 2. The membership profile of an index may be empty; for example, here the membership profile of index 4 in t_4 is \emptyset . Finally, and crucially, the membership profile for an index does not change as more data are observed: for an infinite trait allocation $t_\infty \in \mathcal{T}_\infty$, if τ is the membership profile of index n in t_N for $n \leq N$, then for all $M \geq N$, τ is the membership profile of index n in t_M .

Definition 5.9. The *membership profile of index n* in a finite trait allocation t_N is the multiset $t_N^{(n)}$ of \mathbb{N} defined by

$$t_N^{(n)}(j) := \sum_{\tau \in \mathbb{T}} \mathbf{1}(\tau(n) = j) \cdot t_N(\tau). \tag{5.28}$$

Note that t_N is a partition of $[N]$ if and only if $\forall n \in [N] t_N^{(n)} = \{1\}$, and $\forall n > N t_N^{(n)} = \emptyset$. Likewise, t_N is a feature allocation of $[N]$ if and only if $\forall n \in [N]$ and $j \in \mathbb{N} : j > 1$, we have $t_N^{(n)}(j) = 0$, and $\forall n > N, t_N^{(n)} = \emptyset$.

Definitions 5.10 and 5.11 provide definitions of a frequency model and exchangeable probability function for combinatorial structures with constraints on the membership profiles that are analogous to the earlier unconstrained versions in Definitions 5.1 and 5.3. The intuitive connection to these earlier definitions is made through rejection sampling. First, we define an acceptable set of membership profiles, known as the *constraint set* $\mathfrak{C} \subset \mathbb{T} \cup \{\emptyset\}$. Then, for trait allocations with a *constrained exchangeable trait probability function (CETPF)* in Definition 5.11, we generate T_N from the associated unconstrained ETPF and check if all indices $n \in [N]$ have membership profiles falling in \mathfrak{C} . If this check fails, we repeat the process, and otherwise output T_N as a sample from the distribution. Likewise, for trait allocations with a *constrained frequency model*, we generate Y_n , $n = 1, 2, \dots, N$, progressively checking if all the indices in the associated T_n , $n = 1, 2, \dots, N$ have membership profiles in \mathfrak{C} . If any check fails, we repeat the generation of Y_n for that index $n \in \mathbb{N}$ until it passes. We continue this process until we reach $N \in \mathbb{N}$ and output T_N as a sample from the distribution. To sample T_∞ , we do the same thing but do not terminate the sequential construction at any finite $N \in \mathbb{N}$. Constrained frequency models and CETPFs are the combinatorial analogue of restricted nonparametric processes (Williamson, MacEachern and Xing, 2013; Doshi-Velez and Williamson, 2017).

Definition 5.10. A random infinite trait allocation T_∞ has a *constrained frequency model* with constraint set $\mathfrak{C} \subset \mathbb{T} \cup \{\emptyset\}$ if it has a frequency model with step (2) from Definition 4.4 replaced by

2. For $N = 1, 2, \dots$,

(a) generate $Y_N = R_N + D_N$ as in step (2) of Definition 4.4,

(b) let \overline{Y}_N be the multiset of \mathbb{N} defined by

$$\overline{Y}_N(n) := |\{\phi \in (0, 1) : Y_N(\phi) = n\}|, \quad (5.29)$$

(c) if $\overline{Y}_N \in \mathfrak{C}$, continue; otherwise, go to step 2a.

Note that in Definition 5.10, \overline{Y}_N is precisely the membership profile of index N . That is to say, if we were to construct T_∞ from $Y_\infty = (Y_1, \dots, Y_N, \emptyset, \emptyset, \dots)$, then $\overline{Y}_N = T_N^{(N)}$. Using \overline{Y}_N instead of this construction simplifies the definition considerably.

Definition 5.11. An infinite trait allocation T_∞ has a *constrained exchangeable trait probability function (CETPF)* with constraint set $\mathfrak{C} \subset \mathbb{T} \cup \{\emptyset\}$ if there exists a function $p : \mathbb{N} \times \mathcal{T} \rightarrow \mathbb{R}_+$ such that for all $N \in \mathbb{N}$,

$$\sum_{t_N \in \mathcal{T}_N} \kappa(t_N) \cdot p(N, \overline{t}_N) < \infty \quad (5.30)$$

and

$$\mathbb{P}(T_N = t_N) = \kappa(t_N) \cdot p(N, \overline{t}_N) \cdot \prod_{n=1}^N \mathbf{1}\left(t_N^{(n)} \in \mathfrak{C}\right). \quad (5.31)$$

The extension of Theorem 5.4—a correspondence between random infinite trait allocations T_∞ with constrained frequency models and CETPFs in Definitions 5.10 and 5.11—that applies to constrained combinatorial structures is given by Theorem 5.12.

Theorem 5.12. T_∞ has a constrained frequency model with constraint set \mathfrak{C} iff it has a CETPF with constraint set \mathfrak{C} .

Proof. Suppose T_∞ has a constrained frequency model with constraint set \mathfrak{C} . For finite $N \in \mathbb{N}$, generating T_N from the constrained frequency model is equivalent to generating it from the associated unconstrained frequency model (i.e., removing the rejection in step 2c of Definition 5.10), and then rejecting T_N if $\prod_{n=1}^N \mathbb{1}(T_N^{(n)} \in \mathfrak{C}) = 0$. Since generating T_N from an unconstrained frequency model implies it has an ETPF by Theorem 5.4—which inherently satisfies the summability condition in Definition 5.3 because it is itself a probability distribution—and the final rejection step is equivalent to multiplying the distribution of T_N by $\prod_{n=1}^N \mathbb{1}(T_N^{(n)} \in \mathfrak{C})$ and renormalizing, T_∞ has a CETPF with constraint set \mathfrak{C} .

Next, suppose T_∞ has a CETPF with constraint set \mathfrak{C} . We can reverse the above logic: since the associated ETPF is summable, we can generate T_N by simulating from the (normalized) ETPF and rejecting if $\prod_{n=1}^N \mathbb{1}(T_N^{(n)} \in \mathfrak{C}) = 0$. The ETPF has an associated frequency model by Theorem 5.4. Instead of rejecting T_N after generating all $Y_n, n = 1, 2, \dots, N$, we can reject after each index $n \in \mathbb{N}$ based on progressively constructing $T_n, n = 1, 2, \dots, N$. \square

We can, of course, recover Theorem 5.4 from Theorem 5.12 by setting $\mathfrak{C} = \mathbb{T} \cup \{\emptyset\}$. But Theorem 5.12 also allows us to recover earlier results—using a novel proof technique—about the correspondence of infinite exchangeable partitions and partitions with an EPPF in Corollary 5.13. The proof of Corollary 5.13 uses the fact that the EPPF is a constrained EFPF; it is noted that other connections between classes of probability functions for clustering and feature allocation have been previously established (Roy, 2014).

Corollary 5.13. An infinite partition T_∞ is exchangeable iff it has an EPPF.

Proof. Suppose T_∞ has an EPPF. The EPPF is a CETPF with $\mathfrak{C} = \{\{1\}\}$, and thus T_∞ is exchangeable by inspection of Definition 5.11; the probability is invariant to finite permutations of the indices. In the other direction, if T_∞ is an infinite exchangeable partition, then it has a de Finetti representation of the form specified in Corollary 4.7; for notational brevity define $w_k = \mu_{\xi, \xi'}$ when $\xi_k = 1$ and $w_0 = \mu_{\xi, \xi'}$ when $\xi'_1 = 1$. Note in particular that $\sum_{k=0}^\infty w_k = 1$, and each index $n \in \mathbb{N}$ selects its trait from the distribution $(w_k)_{k=0}^\infty$, where selecting 0 implies selecting a dust (or unique) trait. We seek a constrained frequency model equivalent to this de Finetti representation, so we set $\theta_{kj} = \theta'_j = 0$ for all $k, j \in \mathbb{N} : j > 1$ and seek (θ_{k1}) and θ'_1 such that

$$e^{-\theta'_1} \prod_k \theta_{k0} \propto w_0 \quad \text{and} \quad \forall k \in \mathbb{N}, e^{-\theta'_1} \theta_{k1} \prod_{\ell \neq k} \theta_{\ell 0} \propto w_k. \tag{5.32}$$

Dividing by $\prod_k \theta_{k0}$, this is equivalent to finding (θ_{k1}) and θ'_1 such that

$$\theta'_1 \propto w_0 \quad \text{and} \quad \forall k \in \mathbb{N}, \frac{\theta_{k1}}{\theta_{k0}} \propto w_k. \quad (5.33)$$

We have a degree of freedom in the proportionality constant, so set that equal to 1 and solve each equation by noting that $\theta_{k1} + \theta_{k0} = 1$, yielding

$$\theta_{k1} = \frac{w_k}{w_k + 1} \text{ for } k \in \mathbb{N}, \quad \theta'_1 = w_0. \quad (5.34)$$

The infinite exchangeable partition T_∞ has a constrained frequency model with constraint set $\mathfrak{C} = \{\{1\}\}$ based on (θ_{kj}) , (θ'_j) . By Theorem 5.12 it thus has a CETPF with the same constraint set \mathfrak{C} , which is an EPPF. \square

6. Application: vertex allocations and edge-exchangeable graphs

A natural assumption for random graph sequences with \mathbb{N} -labeled vertices—arising from online social networks, protein interaction networks, co-authorship networks, email communication networks, etc. (Goldenberg et al., 2010)—is that the distribution is projective and invariant to reordering the vertices, i.e., the graph is *vertex exchangeable*. Under this assumption, however, the Aldous–Hoover theorem (Aldous, 1981; Hoover, 1979) for exchangeable arrays guarantees that the resulting graph is either dense or empty almost surely, an inappropriate consequence when modeling the sparse networks that occur in most applications (Mitzenmacher, 2003; Newman, 2005; Clauset, Shalizi and Newman, 2009). Standard statistical models, which are traditionally vertex exchangeable (Lloyd et al., 2012), are therefore misspecified for modeling real-world networks. This model misspecification has motivated the development and study of a number of projective, exchangeable network models that do not preclude sparsity (Caron and Fox, 2017; Veitch and Roy, 2015; Borgs et al., 2018; Crane and Dempsey, 2016a; Cai, Campbell and Broderick, 2016; Herlau and Schmidt, 2016; Williamson, 2016). One class of such models assumes the network is generated by an exchangeable sequence of (multisets of) edges—the so-called *edge-exchangeable* models (Broderick and Cai, 2015; Crane and Dempsey, 2015; Cai, Campbell and Broderick, 2016; Crane and Dempsey, 2016a; Williamson, 2016). These models were studied in the generalized hypergraph setting in concurrent work by Crane and Dempsey (2016b). In this section we provide an alternate view of edge-exchangeable multigraphs as a subclass of infinite exchangeable trait allocations called *vertex allocations*, thus guaranteeing a de Finetti representation. We also show that the *vertex popularity model*, a standard example of an edge-exchangeable model, is a constrained frequency model per Definition 5.10, thus guaranteeing the existence of a CETPF which we call the *exchangeable vertex probability function (EVPF)*. We begin by considering multigraphs without loops, i.e., edges can occur with multiplicity and all edges contain exactly two vertices. We then discuss the generalization to multigraphs with edges that can contain one vertex (i.e., a loop) or two or (finitely many) more vertices (i.e., a hypergraph).




FIG 2. Top: the graph encoded by the vertex allocation $t_4 = \{\{1, 2, 4\}, \{2\}, \{1, 4\}, \{3\}, \{3\}\}$. The four steps show the sequential construction process of the graph. Edge labels correspond to indices, and each trait is a vertex. One or both of the vertices connected to edge 3 and the vertex connected only to edge 2 may be dust; the remaining two are guaranteed to be regular as they connect to multiple unique edge labels (i.e. both 1 and 4). Bottom: the same construction with the edges reordered by the permutation $\pi = (314)(2)$, resulting in the vertex allocation $\pi t_4 = \{\{4, 2, 3\}, \{2\}, \{4, 3\}, \{1\}, \{1\}\}$. If the vertex allocation is exchangeable, these sequences have equal probability.

In the graph setting, the traits correspond to vertices, and the data indices in each trait correspond to the edges of the graph. Each data index has multiplicity 1 in exactly two traits—encoding an edge between two separate vertices—specified in Definition 6.1. Fig. 2 shows an example encoding of a graph as a vertex allocation.

Definition 6.1. A vertex allocation of $[N]$ is a trait allocation of $[N]$ in which each index has membership profile equal to $\{1, 1\}$.

Definition 6.1 and Theorem 4.5 together immediately yield a de Finetti representation for edge-exchangeable graphs, provided by Corollary 6.2. There are three cases: an edge is either a member of two regular vertices, one dust vertex and one regular vertex, or two dust vertices. These three cases are listed in order in Corollary 6.2.

Corollary 6.2. An infinite vertex allocation T_∞ is exchangeable iff it has a de Finetti representation such that $\mu_{\xi, \xi'} > 0$ implies that either

1. $\exists k \neq j$ such that $\xi_k = \xi_j = 1$, $\sum_k \xi_k = 2$, and $\sum_k \xi'_k = 0$,
2. $\exists k$ such that $\xi_k = 1$, $\sum_k \xi_k = 1$, $\xi'_1 = 1$, and $\sum_k \xi'_k = 1$, or
3. $\sum_k \xi_k = 0$, $\xi'_1 = 2$, and $\sum_k \xi'_k = 2$.

Definition 6.1 and Corollary 6.2 can be modified in a number of ways to better suit the particular application at hand. For example, if loops are allowed—useful for capturing, for example, authors citing their own earlier work in a citation network—the membership profile of each index can be either $\{1, 1\}$ or $\{1\}$. This allows indices to be a member of a single trait with multiplicity 1, encoding a loop on a single vertex. If edges between more than two vertices are allowed—that is, we are concerned with hypergraphs—then we may repurpose the definition of a feature allocation, with associated de Finetti representation in Corollary 4.8, where we view the features as vertices. If \mathbb{N} -valued weights are allowed on the multigraph edges, they can be encoded using multiplicities

greater than 1. In this case, the index membership profiles must be of the form $\{j, j\}$ for $j \in \mathbb{N}$, which encodes an edge of weight j . Weighted loops may be similarly obtained by allowing membership profiles of the form $\{j\}$ for $j \in \mathbb{N}$. This might be used, for example, to capture an author citing the same work multiple times in a single document. Weighted hypergraphs are trait allocations without any restrictions.

Vertex popularity models (Caron and Fox, 2017; Cai, Campbell and Broderick, 2016; Crane and Dempsey, 2016a; Palla, Caron and Teh, 2016; Herlau and Schmidt, 2016; Williamson, 2016)⁴ are a simple yet powerful class of network models. There are a number of different versions, but all share the common feature that each vertex is associated with a nonnegative weight representing how likely it is to take part in an edge. Here we adopt a particular construction based on a sequence of edges: all (potentially infinitely many) vertices $k \in \mathbb{N}$ are associated with a weight $w_k \in (0, 1)$ such that $\sum_k w_k < \infty$, and we sample an edge between vertex k and ℓ with probability proportional to $w_k w_\ell$. For an edge-exchangeable vertex popularity model, assuming no loops, Theorem 5.4 enforces that this model has an associated *exchangeable vertex probability function (EVPF)*, given by Definition 6.3.

Definition 6.3. An *exchangeable vertex probability function (EVPF)* is a CETPF with constraint set $\mathfrak{C} = \{\{1, 1\}\}$.

Corollary 6.4. A regular infinite exchangeable vertex allocation has a vertex popularity model iff it has an EVPF.

Proof. We use a similar technique to the proof of Corollary 5.13—we seek a constrained frequency model (a sequence (θ_{kj}) and set \mathfrak{C}) that corresponds to the vertex popularity model with weights (w_i) , and then use Theorem 5.12 to obtain a correspondence with a CETPF (and in particular, an EVPF). We let $\theta'_j = 0$ for all $j \in \mathbb{N}$, let $\theta_{kj} = 0$ for all $j \in \mathbb{N} : j > 1$, and seek (θ_{k1}) such that

$$\forall k, \ell \in \mathbb{N} : k \neq \ell, \theta_{k1} \theta_{\ell 1} \prod_{m \neq k, \ell} \theta_{m0} \propto w_k w_\ell. \quad (6.1)$$

Dividing by $\prod_k \theta_{k0}$, and setting the proportionality constant to 1, Eq. (6.1) is equivalent to

$$\forall k, \ell \in \mathbb{N} : k \neq \ell, \frac{\theta_{k1} \theta_{\ell 1}}{\theta_{k0} \theta_{\ell 0}} = w_k w_\ell. \quad (6.2)$$

Eq. (6.2) may be solved, noting that $\forall k \in \mathbb{N}, \theta_{k0} + \theta_{k1} = 1$, by

$$\theta_{k1} = \frac{w_k}{1 + w_k} \text{ for } k \in \mathbb{N}. \quad (6.3)$$

Therefore the vertex popularity model with weights (w_i) is equivalent to a constrained frequency model with $\theta_{k1} = w_k / (1 + w_k)$ for $k \in \mathbb{N}$, $\theta_{kj} = 0$ for $j > 1$,

⁴These have appeared in previous work as “graph frequency models” (Cai, Campbell and Broderick, 2016) or left unnamed, and the weights w_k are occasionally referred to as “sociality parameters” (Caron and Fox, 2017; Palla, Caron and Teh, 2016).

$\theta'_j = 0$ for all $j \in \mathbb{N}$, and $\mathfrak{C} = \{\{1, 1\}\}$ as specified above. Theorem 5.12 guarantees that the vertex popularity model has a CETPF with constraint set \mathfrak{C} , and likewise that any CETPF with constraint set \mathfrak{C} yields a vertex popularity model by inverting the relation in Eq. (6.3). \square

7. Conclusions

In this work, we formalized the idea of trait allocations—the natural extension of well-known combinatorial structures such as partitions and feature allocations to data expressing latent factors with multiplicity greater than one. We then developed the framework of exchangeable random infinite trait allocations, which represent the latent memberships of an exchangeable sequence of data. The three major contributions in this framework are a de Finetti-style representation theorem for all exchangeable trait allocations, a correspondence theorem between random trait allocations with a frequency model and those with an ETPF, and finally the introduction and study of the constrained ETPF for capturing random trait allocations with constrained index memberships. These contributions apply directly to many other combinatorial structures, such as edge-exchangeable graphs and topic models.

Appendix A: Proofs of results in the main text

Proof of Lemma 3.5. If $\tau = \omega$, then $\tau|_M = \omega|_M$ (and hence $\tau|_M \leq \omega|_M$) for all $M \in \mathbb{N}$ trivially. Otherwise, $\tau < \omega$. Let $m \in \mathbb{N}$ be the minimum index in τ with $\tau(m) > \omega(m)$. If $M \geq m$, then $\tau|_M(m) > \omega|_M(m)$ and $\tau|_M(j) = \omega|_M(j)$ for $j < m$, so $\tau|_M < \omega|_M$ by Definition 3.4. If $M < m$, then $\tau|_M = \omega|_M$, since $\tau(n) = \omega(n)$ for any $n < m$. Therefore, $\tau|_M \leq \omega|_M$. \square

Proof of Lemma 4.2. We prove the result for nonrandom ϕ_∞ ; since the π' we develop does not depend on ϕ_∞ , the result holds for all distinct sequences of labels and thus almost surely for i.i.d. uniform ϕ_∞ as in the main text as well.

Suppose π fixes indices greater than N . Then using Definition 2.6, πt_N is t_N with indices permuted. Let $K_N = \sum_{\tau \in \mathbb{T}} t_N(\tau) < \infty$, the number of traits in t_N . Then let π' be the unique finite permutation that maps the index of each trait τ in $[t_N]$ to its corresponding trait $\pi\tau$ in $[\pi t_N]$, while preserving monotonicity for any traits of multiplicity greater than 1. Mathematically, π' fixes all $k > K_N$, sets $\pi \left([t_N]_{\pi'^{-1}(k)} \right) = [\pi t_N]_k$ for all $k \in [K_N]$, and satisfies $\pi'(k+1) = \pi'(k) + 1$ for all $k \in [K_N - 1]$ such that $[t_N]_k = [t_N]_{k+1}$. Clearly such a permutation exists because πt_N contains the same traits as t_N with indices permuted by Definition 2.6, and the permutation is unique because any ambiguity (where t_N contains traits with multiplicity greater than 1) is resolved by the monotonicity requirement. The monotonicity requirement also implies that π' satisfies the desired ordering condition for all $M \geq N$, i.e.

$$\forall k, M \in \mathbb{N} : M \geq N, \quad \pi \left([t_M]_{\pi'^{-1}(k)} \right) = [\pi t_M]_k, \quad (\text{A.1})$$

since if an index $M > N$ disambiguates two traits, the fact that π fixes all $M > N$ means that these two traits have the same relative order in $[t_M]$ and $[\pi t_M]$. Set $y'_\infty = \varphi(\pi t_\infty, \pi' \phi_\infty)$. By Definition 4.1 and Eq. (A.1), we have

$$\forall M > N, \quad y'_M(\phi_k) = [\pi t_M]_{\pi'(k)}(M) = \pi([t_M]_k)(M), \tag{A.2}$$

and since π fixes indices greater than N (in particular $\pi(M) = \pi^{-1}(M) = M$),

$$\pi([t_M]_k)(M) = [t_M]_k(M) = y_M(\phi_k) = y_{\pi^{-1}(M)}(\phi_k), \tag{A.3}$$

so $y'_\infty = \pi y_\infty$ at all indices greater than N . For the remaining indices, we use Definitions 2.6 and 4.1, the consistency of the trait ordering in Definition 3.4, and the definition of π' in sequence:

$$\forall M \leq N, \quad y'_M(\phi_k) = [(\pi t_N)|_M]_{\pi'(k)}(M) = [\pi t_N]_{\pi'(k)}(M) = \pi([t_N]_k)(M). \tag{A.4}$$

By the definition of permutations of traits in Eq. (2.9),

$$\pi([t_N]_k)(M) = [t_N]_k(\pi^{-1}(M)). \tag{A.5}$$

Finally, again using the consistency of the trait ordering and the fact that $\pi^{-1}(M) \leq N$, we recover the definition of an element in the original label multiset sequence,

$$[t_N]_k(\pi^{-1}(M)) = [t_{\pi^{-1}(M)}]_k(\pi^{-1}(M)) = y_{\pi^{-1}(M)}(\phi_k). \tag{A.6}$$

Thus $y'_\infty = \pi y_\infty$ at all indices less than or equal to N , and the result follows. \square

Proof of Lemma 5.7. For the first statement of the lemma, we need to show that L_N is independent of $\overline{L_{N+1}}, \dots, \overline{L_{N+M}}$ given $\overline{L_N}$ for any $M \in \mathbb{N}$. We abbreviate $\overline{L_{N+1}}, \dots, \overline{L_{N+M}}$ with $\overline{L_{N\uparrow M}}$, and abbreviate statements of probabilities by removing unnecessary equalities going forward, e.g. we replace $\mathbb{P}(L_N = \ell_N \dots)$ with $\mathbb{P}(L_N \dots)$. The fact that $\overline{L_N}$ is a function of L_N and Bayes' rule yields

$$\mathbb{P}(L_N | \overline{L_N}, \overline{L_{N\uparrow M}}) = \frac{\mathbb{P}(\overline{L_{N\uparrow M}} | L_N)}{\mathbb{P}(\overline{L_{N\uparrow M}} | \overline{L_N})} \mathbb{P}(L_N | \overline{L_N}), \tag{A.7}$$

so we require that

$$\mathbb{P}(\overline{L_{N\uparrow M}} | L_N) = \mathbb{P}(\overline{L_{N\uparrow M}} | \overline{L_N}). \tag{A.8}$$

Using the law of total probability,

$$\mathbb{P}(\overline{L_{N\uparrow M}} | L_N) = \sum_{L_{N+M}} \mathbb{P}(\overline{L_{N\uparrow M}} | L_{N+M}) \mathbb{P}(L_N | L_{N+M}) \frac{\mathbb{P}(L_{N+M})}{\mathbb{P}(L_N)}. \tag{A.9}$$

Since both L_N and $\overline{L_{N\uparrow M}}$ are functions of L_{N+M} , the first two probabilities on the right hand side are actually indicator functions. Moreover, knowing L_N

and $\overline{L_{N \uparrow M}}$ determines L_{N+M} uniquely, since the differences between $\overline{L_{N+m}}$ and $\overline{L_{N+m+1}}$ for $m = 0, 1, \dots, M$ allow one to build up to L_{N+M} sequentially from L_N . Thus there is a unique value L_{N+M}^* such that

$$\mathbb{P}(\overline{L_{N \uparrow M}} | L_N) = \frac{\mathbb{P}(L_{N+M}^*)}{\mathbb{P}(L_N)}, \quad (\text{A.10})$$

where L_{N+M}^* satisfies

$$L_{N+M}^*|_N = L_N \quad \text{and} \quad \forall m \in [M], \quad \overline{L_{N+M}^*|_{N+m}} = \overline{L_{N+m}}. \quad (\text{A.11})$$

If we replace L_N with any L'_N such that $\overline{L_N} = \overline{L'_N}$, we have that the corresponding L_{N+M}^* satisfies $\overline{L_{N+M}^*} = \overline{L_{N+M}^*}$. By the ETPF assumption, the marginal distributions of L_N and L_{N+M} depend on only their multiplicity profiles, so

$$\mathbb{P}(\overline{L_{N \uparrow M}} | L_N) = \frac{\mathbb{P}(L_{N+M}^*)}{\mathbb{P}(L_N)} = \frac{\mathbb{P}(L_{N+M}^*)}{\mathbb{P}(L'_N)} = \mathbb{P}(\overline{L_{N \uparrow M}} | L'_N), \quad (\text{A.12})$$

and so summing over all such L'_N ,

$$\mathbb{P}(\overline{L_{N \uparrow M}} | L_N) = \frac{1}{|\{L'_N : \overline{L'_N} = \overline{L_N}\}|} \sum_{L'_N : \overline{L'_N} = \overline{L_N}} \mathbb{P}(\overline{L_{N \uparrow M}} | L'_N). \quad (\text{A.13})$$

Therefore $\mathbb{P}(\overline{L_{N \uparrow M}} | L_N)$ is a function of only $\overline{L_N}$, as desired. To show that $L_N | \overline{L_N}$ has the uniform distribution over all ordered trait allocations L_N with the given multiplicity profile,

$$\mathbb{P}(L_N | \overline{L_N}) \propto \mathbb{P}(\overline{L_N} | L_N) \mathbb{P}(L_N), \quad (\text{A.14})$$

which by the ETPF assumption is a constant for any L_N with the given multiplicity profile, and 0 otherwise. \square

References

- ACKERMAN, N. (2015). Representations of $\text{aut}(M)$ -invariant measures: part 1. [arXiv:1509.06170](https://arxiv.org/abs/1509.06170).
- ALDOUS, D. (1981). Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis* **11** 581–598. [MR0637937](https://doi.org/10.1016/0167-7152(81)90031-3)
- ALDOUS, D. (1985). *Exchangeability and related topics*. *École d'été de probabilités de Saint-Flour, XIII*. Springer, Berlin. [MR0883646](https://doi.org/10.1007/BFb0075001)
- BORGS, C., CHAYES, J., COHN, H. and HOLDEN, N. (2018). Sparse exchangeable graphs and their limits via graphon processes. *Journal of Machine Learning Research* **18** 1–71.
- BRODERICK, T. and CAI, D. (2015). Edge-exchangeable graphs and sparsity. In *NIPS Workshop on Networks in the Social and Information Sciences*.

- BRODERICK, T., PITMAN, J. and JORDAN, M. (2013). Feature allocations, probability functions, and paintboxes. *Bayesian Analysis* **8** 801–836. [MR3150470](#)
- BRODERICK, T. and STEORTS, R. (2014). Variational Bayes for merging noisy databases. In *NIPS Workshop on Advances in Variational Inference*.
- BRODERICK, T., WILSON, A. and JORDAN, M. (2018). Posteriors, conjugacy, and exponential families for completely random measures. *Bernoulli* **24** 3181–3221. [MR3788171](#)
- BRODERICK, T., MACKEY, L., PAISLEY, J. and JORDAN, M. (2015). Combinatorial clustering and the beta negative binomial process. *IEEE transactions on pattern analysis and machine intelligence* **37** 290–306.
- BUSTAMANTE, C., WAKELEY, J., SAWYER, S. and HARTL, D. (2001). Directional selection and the site-frequency spectrum. *Genetics* **159** 1779–1788.
- CAI, D., CAMPBELL, T. and BRODERICK, T. (2016). Edge-exchangeable graphs and sparsity. In *Advances in Neural Information Processing Systems*.
- CARON, F. and FOX, E. (2017). Sparse graphs using exchangeable random measures. *Journal of the Royal Statistical Society Series B* **79** 1295–1366. [MR3731666](#)
- CLAUSET, A., SHALIZI, C. R. and NEWMAN, M. (2009). Power-law distributions in empirical data. *SIAM Review* **51** 661–703. [MR2563829](#)
- CRANE, H. and DEMPSEY, W. (2015). A framework for statistical network modeling. [arXiv:1509.08185](#).
- CRANE, H. and DEMPSEY, W. (2016a). Edge exchangeable models for network data. [arXiv:1603.04571v3](#).
- CRANE, H. and DEMPSEY, W. (2016b). Relational Exchangeability. [arXiv:1607.06762v1](#).
- CRANE, H. and TOWNSNER, H. (2015). Relatively exchangeable structures. [arXiv:1509.06733](#).
- DE FINETTI, B. (1931). Funzione caratteristica di un fenomeno aleatorio. *Atti della R. Accademia Nazionale dei Lincei, Serie 6.* **4** 251–299. In Italian.
- DOSHI-VELEZ, F. and WILLIAMSON, S. (2017). Restricted Indian buffet processes. *Statistical Computing* **27** 1205–1223. [MR3647093](#)
- DUMMIT, D. and FOOTE, R. (2004). *Abstract algebra*, 3rd ed. Wiley, Hoboken. [MR2286236](#)
- ESCOBAR, M. (1994). Estimating normal means with a Dirichlet process prior. *Journal of the American Statistical Association* **89** 268–277. [MR1266299](#)
- ESCOBAR, M. and WEST, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association* **90** 577–588. [MR1340510](#)
- GOLDENBERG, A., ZHENG, A., FIENBERG, S. and AIROLDI, E. (2010). A survey of statistical network models. *Foundations and Trends in Machine Learning* **2** 129–233.
- GRIFFITHS, T. and GHAHRAMANI, Z. (2005). Infinite latent feature models and the Indian buffet process. In *Advances in Neural Information Processing Systems*.
- HERLAU, T. and SCHMIDT, M. (2016). Completely random measures for mod-

- elling block-structured sparse networks. In *Advances in Neural Information Processing Systems*.
- HEWITT, E. and SAVAGE, L. (1955). Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society* **80** 470–501. [MR0076206](#)
- HOOVER, D. (1979). Relations on probability spaces and arrays of random variables. Preprint, Institute for Advanced Study, Princeton University.
- ISHWARAN, H. and JAMES, L. (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association* **96** 161–173. [MR1952729](#)
- ISHWARAN, H. and JAMES, L. (2003). Generalized weighted Chinese restaurant processes for species sampling mixture models. *Statistica Sinica* **13**. [MR2026070](#)
- JAMES, L. (2017). Bayesian Poisson calculus for latent feature modeling via generalized Indian buffet process priors. *Annals of Statistics* **45** 2016–2045. [MR3718160](#)
- JORDAN, M. (2010). Bayesian nonparametric learning: expressive priors for intelligent systems. In *Heuristics, probability and causality: a tribute to Judea Pearl* (R. Dechter, H. Geffner and J. Halpern, eds.) College Publications. [MR2841295](#)
- KALLENBERG, O. (1997). *Foundations of modern probability*, 1st ed. *Probability and its applications*. Springer, New York. [MR1464694](#)
- KINGMAN, J. F. C. (1978). The representation of partition structures. *Journal of the London Mathematical Society* **2** 374–380. [MR0509954](#)
- LEE, J., QUINTANA, F., MÜLLER, P. and TRIPPA, L. (2013). Defining predictive probability functions for species sampling models. *Statistical Science* **28** 209–222. [MR3112406](#)
- LLOYD, J. R., ORBANZ, P., GHAHRAMANI, Z. and ROY, D. (2012). Random function priors for exchangeable arrays with applications to graphs and relational data. In *Advances in Neural Information Processing Systems*.
- MILLER, J., BETANCOURT, B., ZAIDI, A., WALLACH, H. and STEORTS, R. (2016). Microclustering: when the cluster sizes grow sublinearly with the size of the data set. *Advances in Neural Information Processing Systems*.
- MITZENMACHER, M. (2003). A brief history of generative models for power law and lognormal distributions. *Internet Mathematics* **1** 226–251. [MR2077227](#)
- NEWMAN, M. (2005). Power laws, Pareto distributions and Zipf’s law. *Contemporary physics* **46** 323–351.
- PALLA, K., CARON, F. and TEH, Y. W. (2016). Bayesian nonparametrics for sparse dynamic networks. [arXiv:1607.01624](#).
- PITMAN, J. (1995). Exchangeable and partially exchangeable random partitions. *Probability Theory and Related Fields* **102** 145–158. [MR1337249](#)
- PITMAN, J. and YOR, M. (1997). The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability* 855–900. [MR1434129](#)
- ROSS, N. (2011). Fundamentals of Stein’s method. *Probability Surveys* **8** 210–293. [MR2861132](#)

- ROY, D. (2014). The continuum-of-urns scheme, generalized beta and Indian buffet processes, and hierarchies thereof. [arXiv:1501.00208](#).
- ROYCHOWDHURY, A. and KULIS, B. (2015). Gamma processes, stick-breaking, and variational inference. In *International Conference on Artificial Intelligence and Statistics*.
- THIBAUT, R. and JORDAN, M. (2007). Hierarchical beta processes and the Indian buffet process. In *International Conference on Artificial Intelligence and Statistics*.
- TITSIAS, M. (2008). The infinite gamma-Poisson feature model. In *Advances in Neural Information Processing Systems*.
- VEITCH, V. and ROY, D. (2015). The class of random graphs arising from exchangeable random measures. [arXiv:1512.03099](#). [MR3781748](#)
- WALLACH, H., JENSEN, S., DICKER, L. and HELLER, K. (2010). an alternative prior process for nonparametric Bayesian clustering. In *International Conference on Artificial Intelligence and Statistics*.
- WILLIAMSON, S. (2016). Nonparametric network models for link prediction. *Journal of Machine Learning Research* **17** 1–21. [MR3580355](#)
- WILLIAMSON, S., MACEACHERN, S. and XING, E. (2013). Restricting exchangeable nonparametric distributions. In *Advances in Neural Information Processing Systems*.
- ZHOU, M. (2014). Beta-negative binomial process and exchangeable random partitions for mixed-membership modeling. In *Advances in Neural Information Processing Systems*.
- ZHOU, M., HANNAH, L., DUNSON, D. and CARIN, L. (2012). Beta-negative binomial process and Poisson factor analysis. In *International Conference on Artificial Intelligence and Statistics*.