

A SIMPLE AND FLEXIBLE TEST OF SAMPLE EXCHANGEABILITY WITH APPLICATIONS TO STATISTICAL GENOMICS

BY ALAN J. AW^{1,a}, JEFFREY P. SPENCE^{2,b} AND YUN S. SONG^{3,c}

¹Department of Statistics, University of California, Berkeley, ^aalanaw1@berkeley.edu

²Department of Genetics, School of Medicine, Stanford University, ^bjspence@stanford.edu

³Department of Statistics and Computer Science Division, University of California, Berkeley, ^cyss@berkeley.edu

In scientific studies involving analyses of multivariate data, basic but important questions often arise for the researcher: Is the sample exchangeable, meaning that the joint distribution of the sample is invariant to the ordering of the units? Are the features independent of one another, or perhaps the features can be grouped so that the groups are mutually independent? In statistical genomics these considerations are fundamental to downstream tasks such as demographic inference and the construction of polygenic risk scores. We propose a nonparametric approach, which we call the V test, to address these two questions, namely, a test of sample exchangeability given dependency structure of features and a test of feature independence given sample exchangeability. Our test is conceptually simple, yet fast and flexible. It controls the Type I error across realistic scenarios and handles data of arbitrary dimensions by leveraging large-sample asymptotics. Through extensive simulations and a comparison against unsupervised tests of stratification based on random matrix theory, we find that our test compares favorably in various scenarios of interest. We apply the test to data from the 1000 Genomes Project, demonstrating how it can be employed to assess exchangeability of the genetic sample or find optimal linkage disequilibrium (LD) splits for downstream analysis. For exchangeability assessment we find that removing rare variants can substantially increase the p -value of the test statistic. For optimal LD splitting, the V test reports different optimal splits than previous approaches not relying on hypothesis testing. Software for our methods is available in R (CRAN: `flintyR`) and Python (PyPI: `flintyPy`).

1. Introduction. In many practical settings involving the analysis of multivariate samples, a fundamental question arising for the user is the exchangeability of the sample: is the joint distribution of the units making up the sample invariant to the ordering of the underlying units? Stated mathematically, if the sample is $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ with each unit \mathbf{x}_i lying in \mathbb{R}^P , then does the following equation hold true for all permutations π of the index set $[N] = \{1, \dots, N\}$:

$$(1) \quad (\mathbf{x}_1, \dots, \mathbf{x}_N) \stackrel{d}{=} (\mathbf{x}_{\pi(1)}, \dots, \mathbf{x}_{\pi(N)})?$$

To motivate our abstract question with a real example, suppose a geneticist is handed a multivariate single nucleotide polymorphism (SNP) sample consisting of individuals drawn from one or more populations, as depicted in Figure 1, using SNPs for Yoruba and Luhya individuals from the 1000 Genomes Project (1000 Genomes Project Consortium et al. (2015)). An important consideration is whether the sample should be treated as originating from a single population or be considered originating from two or more distinct populations; that is, the sample is structured. In other words, the geneticist ought to consider whether the data: (a) fits a probabilistic model with independent features, possibly with mixture distributions for

Received November 2022; revised March 2023.

Key words and phrases. Exchangeability, feature independence, nonparametric test, population stratification, LD splitting.

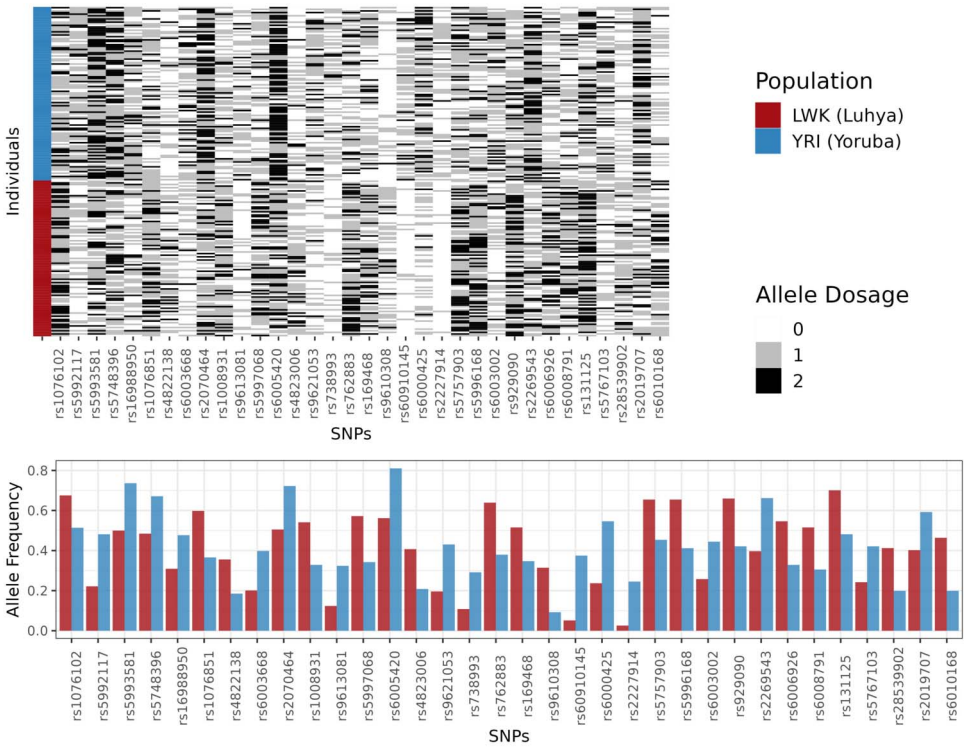


FIG. 1. Heat map of allele dosages (0, 1 or 2) across 34 approximately independent SNP markers from Chromosome 22 for a sample of $N = 205$ African individuals, who are either Yoruba in Ibadan, Nigeria ($N_{YRI} = 108$) or Luhya in Wubye, Kenya ($N_{LWK} = 97$). Population-specific allele frequencies of each marker are depicted in the bar plot below. The user must decide, on the basis of differences in observed allele frequencies, whether the African sample should be treated as a single panmictic population.

each independent feature, or (b) better fits two population-specific probabilistic models, one for Luhya and one for Yoruba, with independent features *conditioned on* population membership, or (c) does not fit either model, because of latent structure not fully captured by the Luhya-Yoruba labeling. This consideration is relevant to downstream genetic analyses that might include fitting a demographic model or computing reference panel statistics for use on future samples. If the sample were structured but the structure was overlooked by the analyst, then the resulting fitted models and reference panels may suffer from poor predictive accuracy on future samples from specific subgroups. A statistical approach to addressing this *structuredness* or *stratification* issue requires stating a null model, which in this setting is the assumption that if features are independent conditioned on population, then draws of units from the same population would produce an exchangeable sample with independent features. However, if units are drawn from different populations, then either the features are dependent (in case population labels are unknown) or the sample must be nonexchangeable (in case population labels are known and fixed, see Figure 1). As shuffling the rows of the array in Figure 1 might suggest, the sample can be made exchangeable by randomly permuting the units, but this will necessarily introduce dependence between the features. Consequently, if a method for addressing stratification assumes that a given sample has independent features—as is customary for fitting population demographic models—then that method cannot simultaneously assume sample exchangeability without implicitly assuming that the units all come from a single “homogeneous” population.

The example above reveals at least two reasons why exchangeability is important. First, from a modeling perspective, exchangeability justifies the use of a statistical model accord-

ing to which each unit is marginally *identically* distributed, despite the potential *statistical dependency between the units*. This can be shown as a mathematical consequence of (1),

$$(2) \quad (1) \implies \forall i \text{ and } j, \quad \mathbb{P}(\mathbf{x}_i \in A) = \mathbb{P}(\mathbf{x}_j \in A)$$

for any Borel-measurable set $A \subseteq \mathbb{R}^P$, with $\mathbb{P} \approx \mathbb{P}_{\mathcal{M}}$ for some statistical model \mathcal{M} in practice. (For a simple argument of (2), see Section 2 of [Kuchibhotla \(2020\)](#).)

Second, exchangeability is a sufficiently weak assumption mirroring realistic sampling procedures, which provides users with statistically valid downstream procedures. In a conformal prediction setting, where it is prudent not to rely on a(n) (over)fitted model to quantify uncertainty on new data, the user falls back on exchangeability to construct valid confidence intervals for model-based predictions on unseen data ([Kuchibhotla \(2020\)](#)). In causal inference studies, when performing counterfactual estimation and inference based on randomization of treatment, individuals are assumed exchangeable across different treatment levels in computing unbiased treatment effect estimates ([Hernán and Robins \(2020\)](#)).

In statistical genomics, which broadly encompasses the development and implementation of statistical models to mine insights from genetic data, the problem of exchangeability is well-recognized and more commonly referred to as the “population structure problem.” Population structure or stratification—more precisely its lack thereof—is appreciated as an implicit requirement for fitting statistical models of evolution to individuals sampled from single, panmictic populations (see Section 5 of [Kingman \(1978\)](#), e.g.). Methods have also been developed to detect stratification from large genetic samples, which recognize that the biological process of recombination generates features that have a block-like dependency structure. This precisely means that the features can be grouped into disjoint blocks such that correlations between blocks are close to zero while those within the same block are bounded away from zero. (The grouping itself is also assumed reasonably well-approximated.)

The foundational role played by recombination in the design of statistical genomic methods has also led to a biologically important dual question: given that a sample of individuals originates from an unstratified population, how can a practitioner adequately identify split points that partition the genome into the dependency blocks described earlier? Known as the “optimal LD splitting problem,” this challenge has motivated the development of LD splitting algorithms, which are critical to downstream tasks of clinical relevance such as polygenic risk score construction.

In this paper we propose a hypothesis testing approach to assessing exchangeability in settings where features can be partitioned into disjoint subsets of features with independence between subsets. Specifically, let $\mathbf{X} = (x_{np}) \in \mathbb{R}^{N \times P}$ be a individual-by-feature dataset whose N units originate in some finite population, with there being no unit labels,

$$\mathbf{X} = \begin{pmatrix} - & \mathbf{x}_1^T & - \\ - & \vdots & - \\ - & \mathbf{x}_N^T & - \end{pmatrix} = \begin{pmatrix} | & \cdots & | \\ \mathbf{x}^{(1)} & \cdots & \mathbf{x}^{(P)} \\ | & \cdots & | \end{pmatrix},$$

with $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ denoting the sample and $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(P)}$ denoting the features. We provide flexible nonparametric tests for the following hypothesis:

(H1) Assuming the dependencies between $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(P)}$ are known and also in the form of groupings of features into disjoint subsets, *the units $\mathbf{x}_1, \dots, \mathbf{x}_N$ are exchangeable* (i.e., (1) holds). Here groupings of features into disjoint subsets means that there exists some integer $B \leq P$ and some known surjective map $f : [P] \rightarrow [B]$ such that $f(i) = f(j)$ if and only if $\mathbf{x}^{(i)} \perp\!\!\!\perp \mathbf{x}^{(j)}$ for any pair of feature indices i and j such that $i \neq j$.

Additionally, our approach can also test the following dual hypothesis, which is relevant to the optimal LD splitting problem arising in statistical genomics.

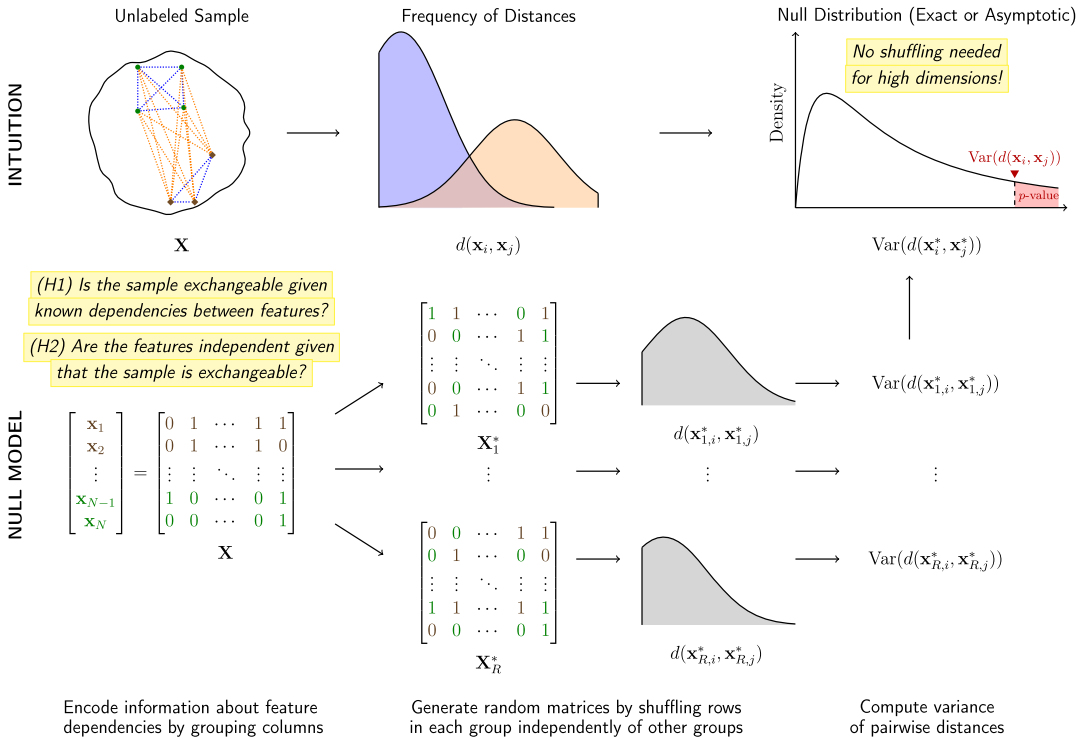


FIG. 2. Overview of our method for detecting sample nonexchangeability or dependence between features.

(H2) Assuming $\mathbf{x}_1, \dots, \mathbf{x}_N$ are exchangeable, the features, or groups of features, are independent.

Our tests are built on the straightforward idea that an exchangeable sample, after accounting for feature dependencies, should have small spread of pairwise distances for a distance metric chosen by the user (see Figure 2). But before elaborating on our approach, we first review existing work.

1.1. *Related work.* The general history of exchangeability tests comprises multiple threads, which we summarize in Appendix A. Within statistical genomics Patterson et al. (2006) were first to propose a test of population structure, based on the spectral theory of random matrices (Bai and Silverstein (2010)). This test relies on the celebrated result that the largest eigenvalue of the covariance matrix of N i.i.d. P -variate sub-Gaussian points has a distribution that is asymptotically Tracy–Widom as P and N tend to infinity with $N/P \rightarrow c \in (0, \infty)$ (Soshnikov (2002)). However, this theory requires selection of approximately independent markers, in practice, and does not leverage LD structure that is prevalent in genetic data. Recently, Zhou, Marron and Wright (2018) propose a computationally intensive block permutation approach that preserves local LD structure while testing for eigenvalue significance. Specifically, given the genotype matrix \mathbf{X} , the authors propose residualizing \mathbf{X} by a singular value-thresholded approximation \mathbf{X}' , attributed to LD, before performing block permutations to the residualized matrix $\mathbf{X} - \mathbf{X}'$ and applying Tracy–Widom theory to determine population stratification.

Closely related to stratification detection in statistical genomics is the dual problem of optimal linkage disequilibrium (LD) splitting. Optimal LD splitting has recently received attention by the genomics community, because many downstream applications require computationally infeasible mathematical operations to be performed on the ultra high-dimensional

LD matrix, for example, in simulation studies (Mancuso et al. (2019)) and in polygenic risk score construction (Mak et al. (2017), Privé et al. (2021), Spence et al. (2022)). As stated in the **Introduction**, the objective is to leverage the banded property of the LD matrix, a consequence of genetic recombination in a panmictic population, to split the $P \times P$ LD matrix (with $P \approx 10^6$ across the entire genome) into approximately independent blocks, thereby allowing mathematical operations to be performed in parallel on the resulting smaller LD submatrices. Existing methods of performing LD splitting include Berisa and Pickrell (2016), Kim et al. (2018), Privé (2022). These methods are all deterministic, relying on optimizing an objective function to produce optimal split points.

1.2. Our contributions. We propose a permutation resampling approach, called the V test, to test sample exchangeability (i.e., Hypothesis H1), given a multivariate dataset \mathbf{X} . We assume that the multivariate features are binary or binarizable to facilitate exposition, but Section 4.2 discusses how our approach can be immediately extended to all types of multivariate features, including those lying in abstract metric spaces. We let the user designate groups for the features and treat different groups as independent (thus, conditioned on population, groups of features are independent of one another). In doing so, the user assumes that the feature grouping captures all feature dependencies and tests (Hypothesis H1). Although this requirement is fairly strong, we show that it is particularly applicable to statistical genomics applications, where the groupings capture local LD structure.

Additionally, because panmictic populations produce exchangeable samples, the V test can also test feature independence (i.e., Hypothesis H2) on genetic data. In particular, since H2 precisely describes the objective of optimal LD splitting, that is, to obtain a partition of features into approximately independent groups, using individuals assumed exchangeable, our test can be used as a post-hoc diagnostic for existing optimal LD splitting algorithms.

Unlike random matrix theory, which principally relies on sample covariances and underlies the works described earlier, we use between-individual distances to construct our test statistic. Unlike previous works that do not address computational limitations of permutation resampling, we use asymptotic theory to obtain large-dimensionality and large-sample approximations of our permutation null distribution, which allows our testing procedure to scale to high-dimensional datasets. Our approach also adapts to feature-feature dependencies in an interpretable fashion: similar to the block permutation approach of Zhou, Marron and Wright (2018), dependent features can be grouped in blocks before performing permutations, with the user choosing the block groupings. The user chooses groupings guided by available domain knowledge or an external procedure, as is the case in Zhou, Marron and Wright (2018), making the test transparent and contingent on interpretable assumptions. Moreover, unlike Zhou, Marron and Wright (2018), we prove that our large-dimensionality approximation works even under this dependent feature setting, allowing our resampling approach to surmount computational difficulties faced by block permutation tests. Finally, focusing on realistic binary or binarizable datasets with characteristics commonly encountered in practical scenarios, we perform an extensive simulation study for evaluating the efficacy of domain-agnostic tests of stratification. Through evaluating both our approach and a random matrix theory approach using this framework, we find that our approach remains well-powered and well-calibrated, even under extreme sample imbalance and other features reflective of real datasets. Moreover, we also identify practical scenarios where using one approach might be better than the other.

The remainder of the paper is organized as follows. In Section 2 we state our test and formulate our algorithm in the ideal scenario where the features are assumed independent. This is written from the point of view of verifying sample exchangeability (Hypothesis H1). We also state our asymptotic results that allow our framework to scale to high-dimensional

datasets. Section 3 reports Type I error control of our test (Section 3.1) as well as the simulation study we perform to evaluate the power (Section 3.2) and efficacy (Section 3.3) of our test on realistic datasets. In Section 4 we state how our approach can be adapted to scenarios where features are dependent, showing that it still scales to high-dimensional datasets and remains valid. Furthermore, we describe how our test generalizes to arbitrary nonbinary datasets. Section 5, largely technical, reports the accuracy of the approximations stated in Section 2. Finally, in Section 6 we demonstrate the dual use of our test—to detect population structure or stratification and to verify independence between groups of features—on genetic data. We conclude with a discussion of our approach, including limitations that motivate potential avenues for future research.

To guide users interested in applying our methods to their work, we provide open-source software and accessible vignettes for all data analyses reported in this paper. Our software is named *flinty* (*flexible and interpretable nonparametric tests of exchangeability*) and is available in both R (*flintyR*) and Python (*flintyPy*).

2. Permutation test of sample exchangeability and feature independence. Let $\mathbf{X} = (x_{np})$ be our $N \times P$ dataset. For exposition we assume the features $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(P)}$ are binary or binarizable so that each entry x_{np} of \mathbf{X} is either 0 or 1. Section 4.2 describes a generalization of our treatment to arbitrary-valued features. Intuitively, if the sample $\mathbf{x}_1, \dots, \mathbf{x}_N$ were exchangeable, then by comparing every subsample of size $M < N$, we should expect small differences between them. We can measure the overall difference between M -subsets by comparing how a summary statistic of an M -subset of $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ differs from the average value of the summary statistic computed across all M -subsets of \mathbf{X} .

2.1. *Test statistic.* To formalize the intuition above, we define the test statistic

$$(3) \quad V_f(\mathbf{X}) = \frac{1}{P \binom{N}{M}} \sum_{S \in \binom{[N]}{M}} [f(\mathbf{X}|_S) - \mu_f]^2,$$

where $f(\cdot)$, which takes on scalar values, is a summary statistic chosen by the user and $\mu_f = \frac{1}{\binom{N}{M}} \sum_{S \in \binom{[N]}{M}} f(\mathbf{X}|_S)$ denotes the average of f computed across all M -subsamples S of \mathbf{X} . Here $\binom{[N]}{M}$ denotes the family of all M -subsets of $[N] = \{1, 2, \dots, N\}$, and $\mathbf{X}|_S$ denotes the $M \times P$ array obtained by including only observations belonging to the M -subset S .

For our present work, we set $M = 2$ and let f be the Hamming distance function $d_H(\cdot, \cdot)$, which counts the number of differences between a pair of individuals considered. That is,

$$f(\mathbf{X}|_S) = d_H(\mathbf{x}_i, \mathbf{x}_j) = \sum_{p=1}^P \mathbb{1}(x_{ip} \neq x_{jp}),$$

where $S = \{i, j\}$ is an arbitrary 2-subset of $\{1, \dots, N\}$. Dropping the subscript f in V_f , this gives

$$(4) \quad V(\mathbf{X}) = \frac{1}{P \binom{N}{2}} \sum_{i < j} [d_H(\mathbf{x}_i, \mathbf{x}_j) - \mu]^2,$$

$$(5) \quad \mu = \frac{1}{\binom{N}{2}} \sum_{i < j} d_H(\mathbf{x}_i, \mathbf{x}_j).$$

Given the test statistic in (4), we now describe its null distribution. Let the sequence of vector-valued observations $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ have a (possibly unknown) joint distribution $\mathcal{L}(\mathbf{X})$. Recall that the distribution \mathcal{L} is exchangeable if it is invariant to permutations of

the index set $[N]$: $\mathcal{L}(\mathbf{X}) = \mathcal{L}(\mathbf{\Pi X})$ for any $N \times N$ permutation matrix $\mathbf{\Pi}$; in other words, $(\mathbf{x}_1, \dots, \mathbf{x}_N) \stackrel{d}{=} (\mathbf{x}_{\pi(1)}, \dots, \mathbf{x}_{\pi(N)})$ for any permutation π . If we further assume that the P features are statistically independent, then the distribution of sequences satisfies a stronger permutation invariance hypothesis, which we call the Exchangeable Sample and Independent Features (ES&IF) null (this combines the two hypotheses described in the [Introduction](#)).

DEFINITION 2.1 (ES&IF Null Hypothesis). Given a multivariate sample $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, the following equality of distributions holds: $(\mathbf{x}_1 = (x_{11}, \dots, x_{1P}), \dots, \mathbf{x}_N = (x_{N1}, \dots, x_{NP})) \stackrel{d}{=} (\mathbf{x}_1^\pi = (x_{\pi_1(1)1}, \dots, x_{\pi_1(1)P}), \dots, \mathbf{x}_N^\pi = (x_{\pi_1(N)1}, \dots, x_{\pi_1(N)P}))$, where π_1, \dots, π_P are P independent permutations, and we denote by \mathbf{x}_i^π the result of applying the P independent permutations to each respective component of observation \mathbf{x}_i .

The ES&IF null hypothesis captures a subtle but important intuition about sample exchangeability: the greater the number of *independent* features measured on N units, the more information there is available about the N units, which makes the assessment of their exchangeability more straightforward. ES&IF also implies that any array \mathbf{X}^* , obtained by permuting the positions of 1s and 0s along each independent feature $\mathbf{x}^{(P)}$, has equal probability of being observed as \mathbf{X} itself. Thus, to arrive at the null distribution of the test statistic under ES&IF, denote the column sums of \mathbf{X} by $\mathbf{c}(\mathbf{X}) = (c_1, \dots, c_P)$ so that c_p counts the number of ones appearing in $\mathbf{x}^{(p)}$. Conditioning on the column sums $\mathbf{c} = \mathbf{c}(\mathbf{X})$ being fixed, the null distribution of V , denoted F_{perm} , is the distribution induced on V by uniformly sampling from all permissible arrays \mathbf{X}^* .

Let $\mathcal{Y}_{\mathbf{c}}$ be the set of all permutation resampled arrays \mathbf{X}^* conditioned on fixing the column sums, that is,

$$\mathcal{Y}_{\mathbf{c}} = \left\{ (\mathbf{x}_1, \dots, \mathbf{x}_N) \in \{0, 1\}^{N \times P} : \sum_{i=1}^N \mathbf{x}_i = \mathbf{c} \right\}.$$

A counting argument shows that the cardinality of $\mathcal{Y}_{\mathbf{c}}$ is given by the quantity

$$|\mathcal{Y}_{\mathbf{c}}| = \prod_{p=1}^P \binom{N}{c_p},$$

which could grow exponentially in N and render the enumeration of all permutations infeasible. Hence, in implementing our test, we allow the user to specify a resampling number R , which sets the number of permuted arrays resampled to approximate the distribution. This conditional Monte Carlo strategy effectively makes our test an *approximate* permutation test, as is typical of many permutation tests. (In our implementations we typically set $R \geq 2000$.) As shown in earlier work (e.g., [Hemerik and Goeman \(2018\)](#), [Phipson and Smyth \(2010\)](#)), the conditional Monte Carlo approach provides an unbiased estimate of the true p -value but suffers from inflated Type I Error for extremely stringent significance cutoffs α (typical of analyses that involve multiple testing). In Supplementary Material B ([Aw, Spence and Song \(2024\)](#)), we show how to leverage permutation invariance to provide a valid test that is more conservative. We describe our implementation of the test in Algorithm 1 (“V Test” of exchangeability).

Note that under ES&IF, the sample is exchangeable, and the features are independent. As a result, rejecting the null indicates that at least one of these assumptions is false: either the sample is nonexchangeable, or the features are nonindependent, or both. If domain knowledge can rule out feature dependence, then this is a test of Hypothesis H1. If domain knowledge can rule out sample nonexchangeability, then this is a test of Hypothesis H2. In other words, the same test statistic and permutation scheme is used, either as a test of H1 or as a test of H2.

Algorithm 1 “V Test”

- 1: **Input:** Individual-by-feature array $\mathbf{X}_{N \times P}$, resampling number R , type of p -value approximation (*unbiased* or *valid*)
 - 2: Record $\mathbf{c} = \mathbf{c}(\mathbf{X})$, μ and $V_{\text{obs}} = V(\mathbf{X})$ (see (4) and (5))
 - 3: Set $r = 0$, $\mathcal{V}^* = \emptyset$
 - 4: **while** $r < R$ **do**
 - 5: Generate resampled array \mathbf{X}^* from permutation null
 - 6: Compute $V^* = V(\mathbf{X}^*)$
 - 7: $\mathcal{V}^* \leftarrow \mathcal{V}^* \cup \{V^*\}$
 - 8: $r \leftarrow r + 1$
 - 9: **end while**
 - 10: **if** type is *unbiased* **then**
 - 11: **Output:** $p = \frac{1}{R} \cdot \#[V^* > V_{\text{obs}}]$
 - 12: **else**
 - 13: **Output:** $p = \frac{1}{R+1} \cdot (\#[V^* \geq V_{\text{obs}}] + 1)$
 - 14: **end if**
-

We conclude with two remarks about our test statistic. First, we find empirically (Section 3) that not only is it powerful when the multiple populations making up the sample are highly heterogeneous but it is also particularly robust to scenarios where there is uneven representation of multiple populations that make up the sample. Second, being the empirical variance of pairwise distances, our test statistic might suggest our approach is testing for homogeneity. This view is mistaken, because homogeneity is difficult to rigorously define in the setting of a single unlabeled sample. In Appendix B we explore a more colloquial interpretation of statistical homogeneity and show that, even for what might be considered a heterogeneous sample, our test correctly identifies it as exchangeable.

2.2. *Asymptotic null distributions.* Running the V Test (Algorithm 1) requires performing R independent resampling routines, with each routine performing independent permutations across P features and then computing $O(N^2)$ pairwise Hamming distances to calculate the test statistic. These amount to $R \times (NP + O(N^2P)) = O(N^2PR)$ operations, which can be slow when P or N is large. To speed things up, we propose three approximations to the null distribution that correspond to three limiting regimes: (1) P is large, (2) both N and P are large and (3) N is large. Approximations (1) and (2) provide exact analytical expressions for the null distribution of our test statistic, which enable the use of much faster numerical integration methods to compute p -values. Approximation (3) is based on the bootstrap. We evaluate the accuracy and speed of our approximations using theory and simulations. To facilitate the exposition of our main results, we defer this evaluation to Section 5. We defer all proofs to Supplementary Material G (Aw, Spence and Song (2024)).

Let N binary vectors with P features be collected, and define the test statistic

$$V^{(N,P)}(\mathbf{X}) := \frac{1}{P \binom{N}{2}} \sum_{i < j} (d_H(\mathbf{x}_i, \mathbf{x}_j) - \mu)^2.$$

Approximation 1 (Large P): The following theorem provides an approximation to the null distribution of the permutation-induced random variable $V^{(N,P)*}$ associated with the test statistic when P is large. It says that $V^{(N,P)*}$ is approximately distributed as a weighted sum of two chi-square random variables, with weights determined by the column sums of the dataset.

THEOREM 2.2 (Large- P Limit). *Let $V^{(N,P)*}$ be the random variable with the distribution of $V^{(N,P)}$ under the ES&IF null (see Definition 2.1). Define the random variable*

$$V^{(N,\infty)} = \frac{a_1^N \chi_{N-1}^2 + a_2^N \chi_{\binom{N-1}{2}-1}^2}{\binom{N}{2}},$$

where a_1^N and a_2^N are large- P limits of quantities depending on the column sums of the dataset and χ_{N-1}^2 and $\chi_{\binom{N-1}{2}-1}^2$ denote independent chi-square random variables with $N - 1$ and $\binom{N-1}{2} - 1$ degrees of freedom, respectively. Then $V^{(N,P)*} \xrightarrow{d} V^{(N,\infty)}$ as $P \rightarrow \infty$.

Theorem 2.2 implies that, for P large, $V^{(N,P)*}$ is approximately equal in distribution to $V^{(N,\infty)}$. We report how quantities a_1^N and a_2^N can be computed from the dataset in Supplementary Material C.1 (Aw, Spence and Song (2024)).

Approximation 2 (Large P and large N): We show in Supplementary Material C.2 (Aw, Spence and Song (2024)) that the null distribution of V converges to a Gaussian distribution as $N, P \rightarrow \infty$.

Approximation 3 (Large N): We show in Supplementary Material A (Aw, Spence and Song (2024)) that the exact distribution of V is a quadratic mapping of an exponential family distribution $f(x|\theta)$ conditioned on a sufficient statistic, where the P -parameter exponential family distribution is given by equation (S1) of Supplementary Material (Aw, Spence and Song (2024)). Differentiating the log-partition function reveals that the MLE of the exponential family parameter θ is $\hat{\theta} = c/N$, which is the column frequency vector of the dataset. Owing to the consistency of the MLE, for large N we may use the MLE $\hat{\theta}$ (obtained from the dataset) to obtain maximum likelihood estimates of the probability mass function of each P -dimensional binary vector \mathbf{x} , and plug these latter estimates into equation (S1) of Supplementary Material (Aw, Spence and Song (2024)) to obtain the parametric bootstrap distribution. Another way to view this distribution is that we resample datasets \mathbf{X}^* by drawing each sample as a realization of a product of Bernoulli distributions, where the parameters of these Bernoulli distributions are estimated as $\hat{\theta} = c/N$ from the dataset.

In practice, Approximation 1 works well, even for surprisingly small P (≈ 50 ; see Section 5). Since both Approximation 1 and Approximation 2 rely on highly efficient numerical integration routines, we find no substantial difference in our results when applying Approximation 1 over Approximation 2, even in situations where Approximation 2 is appropriate. In our simulations and analyses of real datasets, we rely on Approximation 1 whenever applicable. Algorithm 2 describes our implementation of the V Test in our open-source software.

Algorithm 2 Efficient Computation of p -value from Data

- 1: **Input:** Individual-by-feature array $\mathbf{X}_{N \times P}$, resampling number R , type of p -value approximation (*unbiased or valid*)
 - 2: **if** $P \geq 50$, **then**
 - 3: Apply Approximation 1 (see Theorem 2.2)
 - 4: **else**
 - 5: Run Algorithm 1 (“V Test”) with same inputs as in Line 1.
 - 6: **end if**
-

3. Statistical calibration, power and robustness. We evaluate the V test by considering its control of false positive rate (FPR) and its statistical power on simulated data. We consider a variety of simulation scenarios when evaluating statistical power, effectively providing a systematic framework for measuring the robustness of *any* unsupervised test of exchangeability. We study the robustness of V using this framework and report the area under the receiver-operating curve (AUROC) obtained by pairing a null model with a nonexchangeable alternative model. To allow for comparison, we also evaluate the performance of a “Tracy–Widom” (TW) approach, based on the largest eigenvalue of the centered Gram matrix of \mathbf{X} , which we now describe.

Assume that $\mathbf{X} \in \mathbb{R}^{N \times P}$ consists of N i.i.d. sub-Gaussian vectors in \mathbb{R}^P , where for each vector the P components are independent and each is distributed with zero mean and unit variance. A celebrated result in random matrix theory says that under the assumptions: (i) $N \rightarrow \infty$, $P \rightarrow \infty$, and (ii) the ratio P/N stays uniformly bounded by a constant lying in $(0, \infty)$, the normalized maximum singular value $s = \sigma_{\max}(\mathbf{X})$ satisfies

$$\frac{s^2 - (\sqrt{N-1} + \sqrt{P})^2}{(\sqrt{N-1} + \sqrt{P})(\frac{1}{\sqrt{N-1}} + \frac{1}{\sqrt{P}})^{1/3}} \xrightarrow{d} F_1(s),$$

where F_1 is the Tracy–Widom distribution with ensemble index 1 (Tracy and Widom (2002)), that is,

$$(6) \quad F_1(x) = \exp\left(-\int_x^\infty [u(s) + (s-x)u^2(s)] ds\right) \quad \text{for } x \in \mathbb{R},$$

with $u(s)$ defined as the solution to the nonlinear ordinary differential equation $u'' = 2u^3 + su$ with asymptotic condition $u(s) \sim \frac{1}{2\sqrt{\pi}s^{1/4}} \exp(-\frac{2}{3}s^{3/2})$ as $s \rightarrow \infty$. (The ODE is called the Painlevé II equation and its solution the Hastings–McLeod solution.)

Since the square of the maximum singular value $\sigma_{\max}(\mathbf{X})$ is just the eigenvalue of the Gram matrix $\mathbf{X}\mathbf{X}^T$, an asymptotic test can be devised immediately. Let $\mathbf{M} := \mathbf{X}_\circ \mathbf{X}_\circ^T$, where \mathbf{X}_\circ denotes the column-centered and column-scaled version of \mathbf{X} . This test, a variant of which was proposed by Patterson et al. (2006) in population-genetic studies, works as follows. Given an individual-by-feature array \mathbf{X} , for each column $j \in \{1, \dots, P\}$, subtract column means c_j/N from each entry and divide each entry by the normalizing factor, $\sqrt{\frac{c_j}{N}(1 - \frac{c_j}{N})}$. Then an approximate (two-sided) p -value, under the assumption that N observations are independently generated, is given by

$$(7) \quad p = 2 \times \min \left\{ F_1^{-1} \left(\frac{\lambda(\mathbf{M}) - (\sqrt{N-1} + \sqrt{P})^2}{(\sqrt{N-1} + \sqrt{P})(\frac{1}{\sqrt{N-1}} + \frac{1}{\sqrt{P}})^{1/3}} \right), \right. \\ \left. 1 - F_1^{-1} \left(\frac{\lambda(\mathbf{M}) - (\sqrt{N-1} + \sqrt{P})^2}{(\sqrt{N-1} + \sqrt{P})(\frac{1}{\sqrt{N-1}} + \frac{1}{\sqrt{P}})^{1/3}} \right) \right\},$$

where $\lambda(\mathbf{M})$ is the largest eigenvalue of \mathbf{M} and F_1 is the cumulative distribution function in (6).

We refer to (6) as the TW null distribution, and call p -values computed using (7) the TW test. Anticipating readers who might suspect a “straw man” in the midst of our comparison, we note that some of the approaches mentioned in Section 1.1 have proposed modifications to the TW test to deal with idiosyncrasies, like feature dependencies and finite-sample bias. These include using method of moments estimates, pruning or performing regression on features, and fitting reasonably flexible parametric models before performing the test. Here we are interested in comparing two equally straightforward approaches requiring as few modifications to the original dataset as possible. We also want to provide an honest and helpful

evaluation of “folk wisdom” that the TW approximation, per se, is “surprisingly good,” which we believe benefits the broader scientific community.

For the rest of this section, we describe our choice of null and nonnull simulation models and report the AUROCs computed from a null and nonnull pair. Results for statistical power and false positive rate analyses are included in Supplementary Material H (Aw, Spence and Song (2024)).

3.1. Null models to estimate type I error. We simulate binary datasets under three simple generative models corresponding to three scenarios: (i) markers have uniformly low population frequencies, (ii) markers have varying population frequencies and (iii) markers have uniformly high population frequencies. Concretely, each sampled row we draw to form the array is a realization of a product of Bernoulli’s, $\text{Bern}(\theta_1) \times \cdots \times \text{Bern}(\theta_P)$, where the vector of parameters $\vec{\theta} = (\theta_j : j = 1, \dots, P)$ is fixed and determined by the scenario as follows: (i) Low frequencies: Each $\theta_j \in [0.1, 0.2]$; (ii) Varying frequencies: Each $\theta_j \in [0.2, 0.55]$; (iii) High frequencies: Each $\theta_j \in [0.8, 0.9]$.

To demonstrate the performance of our approach on a range of possible numbers of features present in datasets, we also vary P by allowing $P \in \{10, 100, 1000\}$. Note scenario (i) produces sparse arrays by which we mean that the number of nonzero entries in \mathbf{X} is very small compared to the size $N \times P$ of \mathbf{X} . In contrast, scenario (iii) produces dense arrays.

3.2. Nonexchangeable models to estimate power. We simulate datasets under a simple hierarchical generative model, incorporating various sampling designs, parameter choices and data processing or collection artifacts that reflect realistic datasets. Our general model assumes that there are $K \geq 2$ distinct populations from which N_k observations are drawn from Population k ($1 \leq k \leq K$) to make up a sample of size N . These populations are distinct owing to the frequency of each binary feature being distinct at the population level. To produce these distinct population frequencies in turn, we generate them as realizations of uniform distributions. The entire generative process can be described concretely as follows (see Figure S1 in the Supplementary Material (Aw, Spence and Song (2024)) for a plate diagram):

1. Fix ε , a hyperparameter that controls the range of marker frequencies for the population and also determines overall how discerning the P markers are between distinct populations.

2. For a population k ($1 \leq k \leq K$), independently draw P realizations from a uniform distribution parametrized by ε and dependent on k . For example, $\theta_j^{(k)} \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}[0.5 + 0.075 \cdot (-1)^k - \varepsilon, 0.5 + 0.075 \cdot (-1)^k + \varepsilon]$. This produces marker frequencies for Population k . (Details on dependency of the uniform distribution on k are described in the Supplementary Material H.2 (Aw, Spence and Song (2024)).)

3. To draw a sample of size N_k from Population k , independently draw P realizations of Bernoulli distributions, where each Bernoulli distribution j is parametrized by $\theta_j^{(k)}$. In other words, for $i = 1, \dots, N_k$, $\mathbf{x}_i \stackrel{\text{i.i.d.}}{\sim} \text{Bern}(\theta_1^{(k)}) \times \cdots \times \text{Bern}(\theta_P^{(k)})$.

Our sampling designs, parameter choices and data processing artifacts fall under seven scenarios (Table 1). To compare statistical power, we generate datasets by pairing Scenarios 3–7 with Scenarios 1 and 2 in Table 1, illustrating the impact of the sample size N and the closeness of population features on the particular former scenario. To investigate the performance of our approach on a range of possible numbers of features present in datasets, we also choose $P \in \{10, 100, 1000\}$. We estimate power by averaging the true positive rate. Since Section 5.1 shows that the large- P approximation is good for $P \geq 50$, we apply the large- P approximate test whenever $P \in \{100, 1000\}$. Altogether, we perform $[5 \times 4 \times (4 + 9 + 9 + 1 + 1)] \times 3 \times 2 = 2880$ sets of simulations and power estimations, across the two test types (TW vs. V);supplementary Material H.2 (Aw, Spence and Song (2024)) for how we arrive at this count.

TABLE 1
Seven scenarios we consider when generating nonexchangeable samples

Scenario	Relevance or meaning
1. Number of observations	The sample size N of the dataset on which the test is to be performed.
2. Closeness of population features or parameters	How close the true marker frequencies are between the populations whose representatives make up the sample.
3. Number of populations	The number of distinct true populations, K , from which observations were drawn to make up the sample.
4. Sparsity of discerning features	The number of features among all P features that truly discern between the populations whose representatives make up the sample.
5. Evenness of sampling	How evenly represented the various distinct populations are in the sample.
6. Different sources of heterogeneity	How differences in population marker frequencies affect row sums.
7. Column flipping	For binary or binarizable markers, where the binarization provides an interpretation of ‘1’ and ‘0’ for the resulting binary array, the existence or absence of erroneous binarization.

3.3. ROC analysis reveals robustness of nonparametric test. As we report in Supplementary Material H.1 and H.2 (Aw, Spence and Song (2024)), results from running simulations, described above, reveal complex performances of the V test and the TW test. To provide a holistic comparison of V against TW, we consider the test as a binary decision procedure, whereby a dataset is assumed to be drawn uniformly at random from exactly one of a specified pair of generative models and classified as exchangeable or nonexchangeable based on a user-specified significance level α . We pair our null models from Section 3.1 against the nonnull generative models, considered in Section 3.2, and generate receiver-operating characteristic (ROC) curves by sliding the user-specified significance level α from 0 to 1. We evaluate classification accuracy by computing the area under the ROC curve (AUROC). A total of $(3 \times 3) \times (4 + 9 + 9) \times 3 \times 4 \times 2 = 4752$ AUROCs are computed across all pairings (null with nonnull), population closeness parameters (four choices of ε) and test types (TW vs. V); see Supplementary Material H.3 (Aw, Spence and Song (2024)) for how we arrive at this count.

We find that V achieves AUROC at least 0.5, across all sample sizes N , numbers of features P and pairings considered; see Figure 3A. This shows that V performs at least as well as a random classifier, *regardless of* the choice of nonexchangeable model, which is one indication of robustness. The same is not true for TW. That many AUROCs for the TW test lie below 0.5 leads to V being a better classifier on the whole. (See Figures S13–S15 for AUROCs plotted against the various nonexchangeable models considered.) More precisely, we also find that V is particularly robust to sampling unevenness. Figure 3B shows that V, on average, has a higher AUROC and less variability than TW when varying the degree of evenness while holding all other scenario variables constant. In fact, as Figure 3C shows, in case the representation of populations in the sample is very uneven, V still has reasonably high AUROC, but TW has a markedly smaller AUROC.

Finally, we find that V is a relatively weak classifier in cases where the number of discerning features is small; see Figure 3D, for example. (Figure S14 reports all AUROCs against this scenario.). In such cases TW achieves higher classification accuracy overall, even though

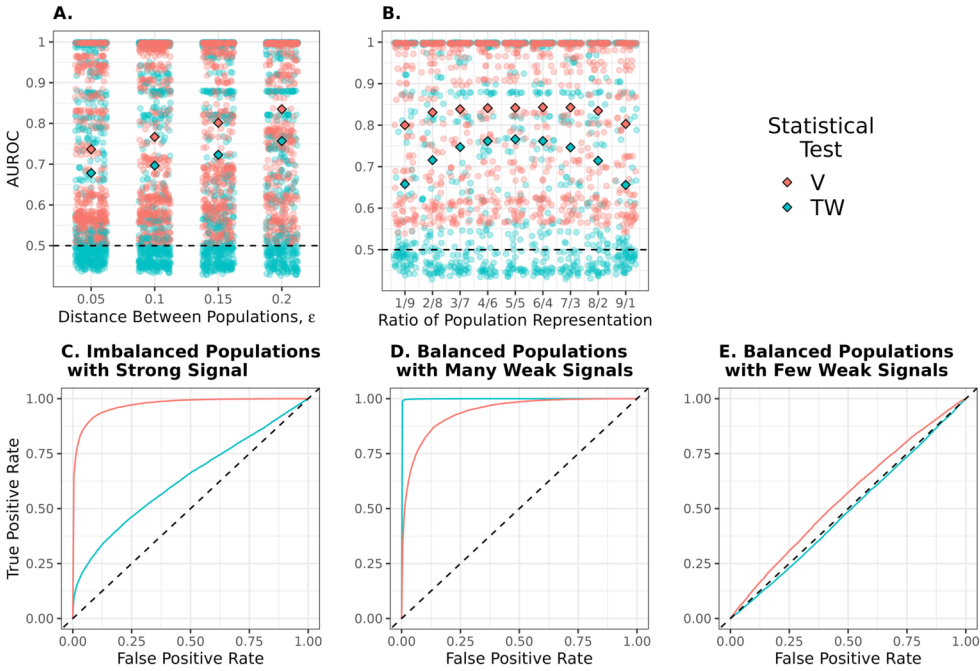


FIG. 3. Top Row shows AUROCs of the V test and of the TW test for pairings of a null model and a nonexchangeable model, with solid diamond points reporting the mean AUROCs for the particular test. Bottom Row shows ROCs generated from pairing a null model and a nonexchangeable model, both of which generate samples containing $N = 50$ observations and P features. A. AUROC points are split into different distances between populations (Scenario 2, Table 1). B. AUROC points are split into different choices of sampling unevenness (Scenario 5, Table 1). C. $P = 100$ features generated. For the nonexchangeable model, individuals are drawn from $K = 2$ populations such that five individuals are drawn from Population 1 and the remaining 45 are drawn from Population 2; population closeness set to $\epsilon = 0.2$. D. 25 individuals are drawn each from $K = 2$ populations, with $P = 1000$ features only 20% of which truly discern between the two source populations; population closeness set to $\epsilon = 0.2$. E. 25 individuals are drawn each from $K = 2$ populations, with $P = 100$ features only 20% of which truly discern between the two source populations; population closeness set to $\epsilon = 0.2$.

for small to moderate number P of features, V is more efficacious, owing to TW having AUROC less than 0.5; see Figure 3E, for example.

4. Adapting to feature dependencies. Statistical independence between features does not hold in many realistic settings. There are many ways in which the P features of \mathbf{X} can depend on each other, for instance, as observations of an undirected graphical model, as draws from a stochastic process or as blocks satisfying between-block independence and within-block dependence. In our present work we consider the setting where the P features are *partitionable*; that is, they can be partitioned into B disjoint sets or blocks $\{1, \dots, P_1\}, \dots, \{P_{B-1} + 1, \dots, P_B\}$ with block delimiters $1 \leq P_1 < \dots < P_B = P$ so that features within the same block are not statistically independent, but features belonging to different blocks are. We modify our ES&IF null hypothesis to accommodate such dependencies as follows: instead of permuting the P features independently, we permute the B sets or blocks independently, keeping the configuration within each block of observations fixed. We call this resulting null distribution on resampled arrays the Exchangeable Sample and Independent Groups of Features (ES&IGF) null (cf. Definition 2.1). This procedure is formalized as Algorithm 1 in Supplementary Material F (Aw, Spence and Song (2024)).

4.1. Asymptotic null distribution. Our asymptotic theory carries over to this setting when $B \rightarrow \infty$ as $P \rightarrow \infty$: as in the independent features case (cf. Theorem 2.2), we may approx-

imate the block permutation null distribution with a convolution of two scaled chi-square distributions. This enables our approach to scale to wide datasets ($P \gg N$), even when the features of the dataset are dependent as long as the number of independent blocks B is large enough. We report this theoretical result in Supplementary Material C.3 (Aw, Spence and Song (2024)).

We evaluate the accuracy of our large B and large P approximation (Theorem S4 in Supplementary Material C.3 (Aw, Spence and Song (2024))) in practice by empirically evaluating its control of FPR for simulated autoregressive time series data and simulated genomes. Details are in Supplementary Material D (Aw, Spence and Song (2024)). As shown in Figure S18, we find that our approximation largely controls FPR, with the null rejected more frequently than α only when $N = 10$. This provides evidence that our approximation is good for reasonably large sample sizes. (One can run the permutation test on datasets with few observations, which is not time-consuming.)

4.2. A general nonparametric test of exchangeability. Theorem S4 in Supplementary Material C.3 (Aw, Spence and Song (2024)) reveals that the large B and large P asymptotics apply directly to *pairwise distances between objects* rather than to the N objects themselves. Thus, our test and its efficient asymptotic counterpart can be applied to datasets where only pairwise distances across independent blocks of binary, real-valued or even abstract features are available. We list some examples in Supplementary Material E (Aw, Spence and Song (2024)).

In practical applications, caching pairwise distances also helps reduce the memory burden of performing the permutation test on ultra high-dimensional datasets with only a small number of independent blocks, where large P asymptotics are invalid. We apply this caching procedure in our assessment of exchangeability of populations from the 1000 Genomes Project in Section 6.1.

5. Speed and accuracy of asymptotic approximations. We justify the use of our approximate null distributions, in practice, by investigating both the accuracy of these approximations via theory and simulations as well as the speed gains by implementing these approximations over permutation resampling.

5.1. Theory and simulations verify accuracy of approximations. We find that the total variation distance between the permutation null distribution F_{perm} , as described in Section 2, and the large P distribution as described in Theorem 2.2, goes to zero at a rate proportional to the square root of the number of independent features, P .

THEOREM 5.1 (Large P Approximation Convergence Rate). *For any fixed sample size N , the rate of convergence of the permutation null distribution to its large P approximate distribution, measured by a bound in the total variation distance, is of order at most $O(P^{-1/2})$. Specifically, for a fixed sample size N , let $V^{(N,P)*}$ and $V^{(N,\infty)}$ be defined as in Theorem 2.2. Then there exists a positive constant C , which depends only on N , a_1^N and a_2^N , such that, for all $t \geq 0$,*

$$|\mathbb{P}(V^{(N,P)*} \leq t) - \mathbb{P}(V^{(N,\infty)} \leq t)| \leq \frac{C}{\sqrt{P}}.$$

In practice, for $P = 50$ independent features—regardless of the magnitude of the sample size N (Figure 4)—the approximation is accurate.

We also observe fast convergence, in practice, for the large P and large N approximation, as described in Theorem S3 of Supplementary Material C.2 and Figure S17 of Supplementary

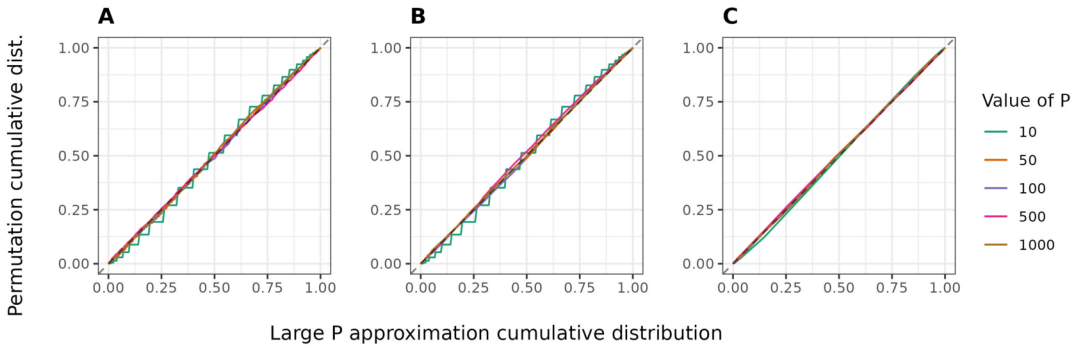


FIG. 4. Probability-probability plots of the permutation-based distribution, F_{perm} , against the large P approximation. **A.** $N = 10$. **B.** $N = 100$. **C.** $N = 1000$.

Material (Aw, Spence and Song (2024)). The parametric bootstrap, described in Section 2.2, converges slower to the null (Figure S16). Based on our simulations, we recommend using the chi-square approximation as long as $P \geq 50$ and recommend using the parametric bootstrap approximation only when $P < 50$ and $N \geq 500$. When $N \geq 50$ and $P \geq 50$, the normal approximation is also fine. These recommendations are based solely on the approximation accuracy; accounting for efficiency in Section 5.2 further narrows down our recommendations.

An analogous result to Theorem 5.1 holds for the large B and large P approximation to the block permutation null distribution described in Section 4. Concretely, owing to similar boundedness assumptions holding in the block permutation null setting, a convergence rate of $O(B^{-1/2})$, where B is the number of independent blocks, can be obtained.

5.2. Speed gains for wide and high-dimensional arrays. To compare the speed gains from running our approximations, we run our permutation test and its approximations on 100 simulated datasets with varying dimensionalities $(N, P) \in \{(50, 500), (500, 50), (500, 500)\}$, calculating the time it takes for each algorithm to compute 100 p -values from 100 generated arrays of varying dimensions. For both the exact and the parametric bootstrap resampling algorithms, we set the resampling number $R = 5000$. We run all algorithms on a Macbook Pro CPU with four cores, a 2.3 GHz processor and 16 GB memory.

Table 2 summarizes our runtime experiment, where we report the average runtime across all 100 p -value computations. We find that the chi-square test is, on average, at least 2000 times faster than the permutation test. We also find that the parametric bootstrap can surprisingly be slower than the permutation test for problem dimensionalities where it is applicable. This is likely to do with our optimized implementation of the permutation test, where we: (1) compute Hamming distances with C or C++ bitwise operations and (2) cache pairwise

TABLE 2

Average runtime (in seconds) for each algorithm to compute a single p -value from arrays with varying dimensionalities. **Boldfaced times** indicate that the algorithm is statistically appropriate for the problem's dimensionalities, as evidenced by the analysis in Section 5.1

Dimensionality	Permutation-based	Chi-square	Bootstrap	Normal
$N = 50, P = 50$	4.52	3.99×10^{-3}	3.20	9.40×10^{-4}
$N = 50, P = 500$	27.81	1.07×10^{-2}	8.30	7.87×10^{-3}
$N = 500, P = 50$	37.36	1.33×10^{-2}	97.81	1.11×10^{-2}
$N = 500, P = 500$	96.01	4.10×10^{-2}	81.68	3.78×10^{-2}

Hamming distances with their corresponding sample indices to avoid costly Hamming distance computations required per permutation.

Considering both the accuracy and the speed gains of our approximations, we find that the chi-square approximation is the most reliable in practice, and we recommend its use as long as $P \geq 50$. (In all other cases, use the permutation test.) The normal approximation is also reliable, but considering the practically insignificant differences in runtimes, we do not strongly recommend it.

6. Application to data. We apply our approach to two problems in statistical genomics: (1) stratification detection and (2) optimal LD splitting. The first problem, which we alluded to in the opening question and example in the [Introduction](#), allows us to apply V as a test of Hypothesis H1. The second problem, owing to partitionability of genetic features, allows us to apply V as a test of Hypothesis H2. In both applications we apply the general version of our test, which does not require that the features are binary. Code for all analyses is available through a zip file in the Supplementary Material and also online at: <https://github.com/songlab-cal/flinty>.

6.1. Stratification detection. In studies involving clustering human populations from genomic data, unstructuredness is often an implicit desideratum of a cluster. More broadly, in genetic association studies or analyses involving the fitting of demographic models to genomic data, population stratification can be a source of confounding, resulting in inaccurate inferences of evolutionary parameters of interest.

To evaluate the exchangeability of real genetic samples, we run the V test on the 26 populations comprising the 1000 Genomes Project. The sample sizes for these populations range from $N_{\min} = 55$ to $N_{\max} = 113$, and the number of diploid variants genome-wide is $P = 1,836,406$ after removal of variants not passing the Hardy–Weinberg test (see `process_1000G.txt` in Supplementary Material zip file for details of data pre-processing). We group variants within the same chromosome together and assume that variants from different chromosomes are independent, because genetic recombination (or “crossing over”), which breaks down linkage disequilibrium between variants, occurs at a rate directly proportional to physical distance. This procedure partitions the P variants into $B = 22$ sets on which we proceed to apply the independent blocks version of our test, as described in Section 4 (i.e., we test Hypothesis H1 under the null distribution given by ES&IGF). We use the Manhattan metric to compute pairwise distances between individuals, set the resampling number $R = 2000$ and cache pairwise distances within each set of variants owing to the large dimensionalities of each chromosome (see Section 4.2). As we run the V test, we successively remove rare variants from each population by applying a progressively larger allele frequency threshold r for variant inclusion within each population that increasingly restricts the number of variants included ($r \in \{0.00, 0.01, 0.02, 0.03, 0.04, 0.05\}$).

We find that, in the case where all the variants are included (i.e., $r = 0.00$), all but one population has very small p -values ($\ll 0.05$); the Yoruba population has p -value = 0.027. However, as we remove more and more rare variants, while most populations still report very small p -values ($\ll 0.05$), we observe a generally increasing trend (see Figure 5) in p -value for three populations: the Utah population carrying Northern and Western European ancestry (CEU), the Vietnamese Kinh population (KHV) and the Yoruba population (YRI). In particular, when an allele frequency threshold of $r = 0.05$ is applied, the Yoruba population has a p -value of 0.28, which is not only insufficient evidence to reject Hypothesis H1 at $\alpha = 0.05$ but also a >10 -fold increase from the p -value reported when to all variants are kept.

Our findings suggest that the geographically defined populations in the 1000 Genomes Project are structured, that is, not exchangeable. The insufficient evidence that the Yoruba




FIG. 5. Exchangeability test (Hypothesis H1) p -values for the Utah population (CEU), Kinh population (KHV) and Yoruba population (YRI) across progressively stringent allele frequency threshold choices, r . Raw p -values are log-transformed for better visualization.

population is structured upon removal of rare variants is consistent with reports of high levels of inbreeding in most 1000 Genomes populations (see Figure 2 of Gazal et al. (2015)), because close relatives tend to share family-specific rare variants (Shirts, Pritchard and Walsh (2016)). In other words, the removal of rare variants likely removes those variants arising from fine structure to do with inbreeding, thus attenuating or even removing the signal of genetic stratification.

6.2. *Optimal linkage disequilibrium (LD) splitting.* In many practical applications of statistical genomics, including summary statistics imputation and computation of polygenic scores for precision health, genome-wide correlation matrices, or “LD matrices” as they are more commonly known, are required as input to some algorithm. LD matrices are not only ultra high-dimensional (typically on the order of $10^6 \times 10^6$), presenting challenges in performing mathematical operations on them, but they also possess block-like structure. Taken into consideration together, these qualities have motivated the development of methods to split the LD matrix into blocks (Berisa and Pickrell (2016), Kim et al. (2018), Privé (2022)). The goal of such splits, presumably obtained by some LD splitting algorithm, is to obtain a set of submatrices, or LD blocks, whereby variants in distinct blocks are approximately independent of one another.

Popular LD splitting algorithms rely on minimizing a cost function associated with the LD blocks, without explicitly considering assumptions about the cohort from which the original genome-wide LD matrix was derived. This can be problematic, because genetically stratified cohorts often result from complex population histories (e.g., admixing, endogamy) that may produce long-range LD patterns, thus violating the block-like structure assumption (Price et al. (2008)). We show how our method complements existing splitting algorithms by formally testing the hypothesis that the variants between blocks are independent, while assuming the cohort from which the LD matrix is computed is exchangeable (i.e., Hypothesis H2). Concretely, we consider a sample of $N_{\text{Afr}} = 652$ individuals of African ancestry from the 1000 Genomes Project. We restrict to Chromosome 22 single nucleotide variants and include only variants satisfying minor allele frequency > 0.05 ($P_{\text{Afr}} = 18,791$) before computing the LD matrix. Using publicly available optimal splits produced by `ldetect` (Berisa and Pickrell (2016)), we partition the P_{Afr} features into $B_{\text{Afr}}^{\text{ldetect}} = 34$ blocks of features. We run the

independent-blocks version of our test on the individual-by-genotype matrix, with variants in the same block grouped together and use the Manhattan metric and $R = 3000$ permutations. We obtain a very small p -value ($\ll 0.05$). We next perform LD splitting on the $P_{\text{Afr}} \times P_{\text{Afr}}$ LD matrix using a dynamic programming approach, `snpldsplit`, recently proposed by Privé (2022) (parameter settings: `thr_r2 = 0.0`, `min_size = 500`, `max_size = 10,000` and `max_K = 40`). After obtaining the optimal split blocks ($B_{\text{Afr}}^{\text{snpldsplit}} = 6$), we run the independent blocks version of our test on the individual-by-genotype matrix, with variants in the same block grouped together and with the same test settings described above. We obtain a very small p -value again ($\ll 0.05$). We further perturb parameter settings for the LD splitting algorithm, which include a thresholding parameter to account for spurious finite-sample correlations and the minimum and maximum block sizes. These perturbations do not lead to meaningful changes in the small p -value returned by our test. Thus, we find that multiple LD splitting algorithms do not produce approximately independent blocks for the African sample, assuming that the African sample is exchangeable. A possible reason is that our assumption about the African sample being exchangeable—a sine qua non of formally testing H_2 —is actually false. As it turns out, the individuals making up the sample are from seven geographically distinct populations with previously reported population-specific differences in recombination frequencies (Spence and Song (2019)). Notably, the sample contains African-American individuals residing in the United States Southwest, who have varying degrees of admixed genetic components owing to non-African ancestral genetic contributions, resulting in (well-documented) long-range LD patterns that are unlikely present in other less admixed African populations (Mourad et al. (2011)). This suggests that the individuals are not exchangeable.

We next analyze a subsample of the African sample that consists of only $N_{\text{Yor}} = 108$ Yoruban individuals, who all reside in Ibadan, Nigeria. To our knowledge, there is no evidence that these individuals have detectable population substructure, so we assume that they are exchangeable. Again, after restricting to Chromosome 22 and keeping only variants satisfying minor allele frequency > 0.05 ($P_{\text{Yor}} = 18,376$), we compute the LD matrix for this smaller set of individuals. We first partition the P_{Yor} features into $B_{\text{Yor}}^{\text{lddetect}} = 34$ blocks of features using optimal splits, computed by `lddetect` for the African sample, and run the V test on the $N_{\text{Yor}} \times P_{\text{Yor}}$ matrix, with variants in the same block grouped together and using the same test settings described in the previous paragraph. We obtain a p -value of 0.06, which is insufficient evidence to reject Hypothesis H_2 at $\alpha = 0.05$. We next perform LD splitting with `snpldsplit`, using the same parameter settings described in the previous paragraph, except for `thr_r2`, which we increase to 0.2 to ensure that the algorithm accounts for spurious positive correlations. We obtain $B_{\text{Yor}}^{\text{snpldsplit}} = 15$ blocks of variants. Similar to how we run the V test using `lddetect` splits, we now run it using the `snpldsplit` splits. We obtain a p -value of 0.66, which is insufficient evidence to reject Hypothesis H_2 at $\alpha = 0.05$. To address the possibility of a smaller sample size reducing the power of our test, we further perform a subsampling analysis, where we repeat the procedure on 500 random 108-subsets of the 652 African individuals. We find that across all 500 subsets, either using `lddetect` splits for the African sample or after identifying optimal splits using `snpldsplit`, p -values returned by V are significant at a nominal 0.05 level. This is so, even after controlling the false discovery rate using the Benjamini–Hochberg procedure. Altogether, our analysis provides reasonable justification that H_2 holds for the optimal split identified for the set of Yoruban individuals.

To briefly investigate the utility of our test beyond merely providing post hoc verification of optimal LD splits, we additionally compare various splits returned by `snpldsplit` for the Yoruban individuals. The optimal split reported in the previous paragraph partitions Chromosome 22 into 15 disjoint blocks. However, we observe that suboptimal splits need




FIG. 6. Rotated heatmap of pairwise LD r^2 values within a 2000 b.p. region of Chromosome 22, with r^2 values less than 0.05 removed for better visualization. Superimposed on the heatmap are split points lying in the region, as identified by various LD splitting methods, including `ldetect` (split points for entire African sample) and `snp_ldsplit` (optimal split points and suboptimal split points), as described in Section 6.2. The split points are given by `ldetect`: (21419799, 22878110, 2317414023717987, 24488861, 25664408), optimal `snp_ldsplit`: (22579801, 23849683) and suboptimal `snp_ldsplit`: (22579801, 23849683, 25286983).

not lead to the rejection of Hypothesis H2. For example, we find a suboptimal split yielding 21 disjoint blocks for which the V test returns a p -value of 0.64 using the same test settings described earlier. This suboptimal split shares common split points with the optimal split but also identifies additional split points, such as *rs139729* at physical position 25, 286, 983 (see Figure 6).

In summary, our method complements existing LD splitting algorithms for generating optimal LD splits and emphasizes careful consideration of the assumptions about the cohort. Such latter considerations may be beneficial for downstream statistical genomics algorithms, such as variant imputation and polygenic score model fitting.

7. Discussion. We have presented an exact, nonparametric approach to testing if a multivariate sample is exchangeable or if its features are independent. We have shown that our approach scales to high dimensionalities of the dataset and flexibly accommodates feature dependencies obeying a partitionable dependency structure. We have also demonstrated, through extensive simulations, when our approach is robust, especially so by making comparisons with eigenanalysis approaches that have gained popularity in recent works. Through applications to simulated and real genetic datasets, we have provided concrete ways in which our approach can be used in statistical genomics. These include detecting population stratification and ensuring optimal LD splits do not violate the exchangeability assumption about the genomic sample.

One limitation of our approach is the need for feature dependencies to be partitionable when testing for exchangeability. While this requirement is reasonable in statistical genomics, it is not true for many other real datasets, where complex dependency structures underlie the observed feature-feature correlations. In such a setting, it is difficult to construct a permutation of observations that preserves the dependency structure. This is consistent with many methods in practice relying on parametric resampling, that is, empirical Bayes approach after fitting a graphical model on the features. From another perspective and in situations where it may be unclear how to choose a reasonable model to capture the dependency structure, our approach provides a clear preliminary approach for deciding exchangeability while accounting for feature dependencies to some degree.

Another limitation is that we have diagnosed the efficacy of our test only under the setting where the P features are independent and binarizable. Although the broad conclusions

derived from our simulation study will likely port over to the nonindependent version of our test (or even the most general version described in Section 4.2), it will be more revealing to perform a thorough diagnosis of our approach against real and simulated datasets with multivariate partitionable dependent features.

Limitations aside, our present work can be extended in many ways. First, we can modify the test statistic by: (1) exploring functionals other than the variance of the user-defined distance and (2) introducing weights to the features when computing differences. We surmise that such modifications may identify even more powerful tests but also suspect that obtaining asymptotic approximations will be challenging. Second, given the prominence of finite-sample tail bounds in the recent literature on high-dimensional statistics, it is possible that such tail bounds can be used to compute lower tails of our observed test statistic, providing an efficient “CDF integration”-free means to obtain p -values. We pursued this direction and encountered difficulty in obtaining tight and achievable upper bounds on sub-Gaussian and subexponential parameters, which matter in practice. Third, in our simulation study verifying robustness, our nonexchangeable models are characterized by samples drawn from different multivariate distributions, with each distribution characterized by a product of independent univariate distributions. It would be interesting to broaden the family of nonexchangeable models considered, especially in domains where such nonexchangeable models are known or well-studied.

Our current work does not explore our approach as a test of independence. There are existing approaches for testing independence of features, given multiple vector-valued observations. Some of these approaches rely on kernel methods (Pfister et al. (2018), Gretton and Györfi (2010)) or rank-based methods (Han, Chen and Liu (2017)), while others, like our present work, leverage sample distances (Heller and Heller (2016), Guo and Modarres (2020)). These methods largely assume that the observations are i.i.d., which is stricter than exchangeability, and do not include cases where variables are partitioned into mutually independent groups. It will be interesting to compare our approach against such approaches, especially in settings where the observations are exchangeable, but not i.i.d., and to consider how tests of independence of variables may be generalized to tests of independence of *groups of variables*. Finally, our work may find further uses in population and statistical genomics. For example, to evaluate the efficacy of our test at detecting genetic substructure or stratification more substantially, one may investigate the limits of our approach on samples drawn from various nonexchangeable demographic models, while varying ranges of parameters responsible for nonexchangeability (e.g., time since population split or admixture, change in recombination rates and inclusion or removal of rare variants). Another relevant investigation would be the impact of using optimal LD splits not violating Hypothesis H2 on downstream tasks, like variant imputation and polygenic risk score construction.

In summary, our work interrogates a fundamental but important assumption made in many areas of data analysis and contributes to the growing literature on applications of exchangeability to modern statistics. On top of our carefully exposed technical proofs, which may be of interest to some readers seeking to extend permutation testing methodology, we hope that our work will generate some discussion among the wider scientific community.

APPENDIX A: GENERAL REVIEW OF TESTS OF EXCHANGEABILITY

To our knowledge, the earliest test of sample exchangeability is the correlation test of randomness of Wald and Wolfowitz (1943). For a univariate real-valued N -sample $\{x_1, \dots, x_N\}$, in order to test that the joint distribution $F(x_1, \dots, x_N)$ is given by the product $F(x_1) \cdots F(x_N)$, the correlation test calculates the quantity $R_h = \sum_{n=1}^N x_n x_{n+h}$, called the lag h correlation coefficient, for some user-chosen h . Wald and Wolfowitz (1943) showed—without requiring that the x_n 's are i.i.d.—that under the null where $\{x_1, \dots, x_N\}$ is uniformly

sampled from one of the $N!$ permutations of the underlying values (this is an exchangeable null, also called the randomization hypothesis; see [Bartels \(1982\)](#), [Vovk \(2021\)](#)), R_h is approximately normally distributed.

Subsequent application-driven tests of exchangeability have arisen in a variety of other contexts, and they broadly fall under tests pertaining to the sample or tests pertaining to features. Tests pertaining to features are generally relevant in settings where repeated measurements are obtained across multiple subjects, such as in clinical trials where each measurement corresponds to a treatment or control. In such contexts, tests of exchangeability typically work with bivariate data, and are also called tests of bivariate symmetry (see [Modarres \(2008\)](#) and [Kalina and Janáček \(2022\)](#) for a comprehensive review). Recently, [Kalina and Janáček \(2022\)](#) propose a way of generalizing bivariate symmetry tests to settings with multivariate features, where p -values are obtained for each pair of features before being nonparametrically combined via a combining function (see Section 1.2 of [Bonnini et al. \(2014\)](#)). Although the null hypothesis in their work is not strictly joint exchangeability of the features but rather a composite null of bivariate symmetry across all pairs, the authors demonstrate that nonparametric combination methodology leads to more powerful tests against the composite null, when compared against standard multiple testing correction procedures (e.g., Benjamini–Hochberg and Benjamini–Yekutieli).

Tests pertaining to the sample, which include our contributions in this work, are driven by two applications. The first is statistical genomics, which we review in the [Introduction](#). The second, conformal prediction, is concerned with predictions made by machine learning algorithms, especially when the performance of the algorithm depends crucially on distributional similarities between already seen training data and new, unseen data ([Shafer and Vovk \(2008\)](#), [Angelopoulos and Bates \(2023\)](#)). Although many applications of conformal prediction defer the exchangeability assumption to user judgement, in settings where new data arrives in a stream (e.g., time series or online learning), methods relying on martingale techniques have been proposed to test exchangeability (see Chapter 5 of [Balasubramanian, Ho and Vovk \(2014\)](#)).

In concluding our review, we note that there are also settings where exchangeability cannot be verified, such as in studies where it is impossible to know or observe all potential confounders ([Tchetgen Tchetgen et al. \(2020\)](#)). However, even if all covariates relevant to potential confounding mechanisms are measured, covariate-covariate dependencies and correlations can generate spurious signals of sample nonexchangeability, for instance, through inflated spectral statistics computed on the sample covariance matrix ([Efron \(2009\)](#)). Unfortunately, in many practical settings, one does not know the exact feature dependencies to correctly account for them while deciding sample exchangeability. These considerations are thematically related to historical discussions ([Draper et al. \(1993\)](#), [Lindley and Novick \(1981\)](#)) on the limits of formally assessing exchangeability in data analysis, which are also important in their own right.

APPENDIX B: EXAMPLE OF HETEROGENEOUS BUT EXCHANGEABLE SAMPLE

Are exchangeability and homogeneity essentially the same qualities of a sample? Scientists frequently think about homogeneous populations, homogeneous proportions and homogeneous clusters. A common way of conceptualizing homogeneity is to relate to statistical properties that are similar across multiple groups within a sample, *in the presence of group labels*. This conceptualization is often taken to imply that, even in the *absence of group labels* in a given dataset, the statistical properties of any one part are the same as any other part. Below we show that the latter intuition, despite being regarded as “common sense,” is different from exchangeability.




FIG. 7. An exchangeable but heterogeneous sample. We set $P = 2$ in the model described in Appendix B and draw $N = 40$ vectors $\mathbf{x}_n \in \mathbb{R}^2$. Points are shaped by the number of coordinates that lie above or below 0.

Let us define a population for which there are P observable features. These features are independent and identically distributed according to the mixture distribution $\frac{1}{2}N(-2, 1) + \frac{1}{2}N(2, 1)$. This means that for each feature, we flip a fair coin to decide whether a feature is drawn from $N(-2, 1)$ or from $N(2, 1)$. Our generative model is one where the features are independent conditioned on a single population.

Drawing N i.i.d. vectors from this distribution, $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, we obtain a sample, which we can also view as a $N \times P$ matrix after stacking the vectors horizontally, as described in the Introduction.

This sample has 2^P clusters, which suggests that the sample is heterogeneous; see Figure 7. Yet the sample is also exchangeable: we can verify that the joint distribution of $(\mathbf{x}_1, \dots, \mathbf{x}_N)$ satisfies (1) in the Introduction. Consequently, we can find some distribution \mathbb{P} according to which each unit is marginally identically distributed; see (2). (The generative model we described would be a statistical model giving rise to one such distribution.)

This example also illustrates that for samples without group labels, if the downstream goal is to *fit a statistical model to data*, then exchangeability—rather than homogeneity—is arguably a clearer conceptualization.

Acknowledgments. We thank Dan Erdmann-Pham, Ziyue Gao, Iain Mathieson, Nick Patterson, Sebastián Prillo, Florian Privé and Clara Wong-Fannjiang for helpful discussions. All authors are affiliated with the Center for Computational Biology at UC Berkeley.

Funding. This research is supported in part by an NIH Grant R35-GM134922.

SUPPLEMENTARY MATERIAL

Supplementary information (DOI: [10.1214/23-AOAS1817SUPPA](https://doi.org/10.1214/23-AOAS1817SUPPA); .pdf). The Supplementary Information PDF includes technical details, proofs, and supplementary figures for our work.

Zip file (DOI: [10.1214/23-AOAS1817SUPPB](https://doi.org/10.1214/23-AOAS1817SUPPB); .zip). The zip archive (supplement.zip) contains code for reproducing our analyses of 1000 Genomes data, and instructions on accessing our open-source software and integrative tutorials.

REFERENCES

- ANGELOPOULOS, A. N. and BATES, S. (2023). Conformal prediction: A gentle introduction. *Found. Trends Mach. Learn.* **16** 494–591.
- AW, A. J., SPENCE, J. P. and SONG, Y. S. (2024). Supplement to “A simple and flexible test of sample exchangeability with applications to statistical genomics.” <https://doi.org/10.1214/23-AOAS1817SUPPA>, <https://doi.org/10.1214/23-AOAS1817SUPPB>
- BAI, Z. and SILVERSTEIN, J. W. (2010). *Spectral Analysis of Large Dimensional Random Matrices*, 2nd ed. *Springer Series in Statistics*. Springer, New York. MR2567175 <https://doi.org/10.1007/978-1-4419-0661-8>
- BALASUBRAMANIAN, V., HO, S.-S. and VOVK, V. (2014). *Conformal Prediction for Reliable Machine Learning: Theory, Adaptations and Applications*. Morgan Kaufmann, San Mateo.
- BARTELS, R. (1982). The rank version of von Neumann’s ratio test for randomness. *J. Amer. Statist. Assoc.* **77** 40–46.
- BERISA, T. and PICKRELL, J. K. (2016). Approximately independent linkage disequilibrium blocks in human populations. *Bioinformatics* **32** 283–285. <https://doi.org/10.1093/bioinformatics/btv546>
- BONNINI, S., CORAIN, L., MAROZZI, M. and SALMASO, L. (2014). *Nonparametric Hypothesis Testing: Rank and Permutation Methods with Applications in R*. *Wiley Series in Probability and Statistics*. Wiley, Chichester. MR3237553 <https://doi.org/10.1002/9781118763490>
- DRAPER, D., HODGES, J. S., MALLOW, C. L. and PREGIBON, D. (1993). Exchangeability and data analysis. *J. Roy. Statist. Soc. Ser. A* **156** 9–28. MR1207984 <https://doi.org/10.2307/2982858>
- EFRON, B. (2009). Are a set of microarrays independent of each other? *Ann. Appl. Stat.* **3** 922–942. MR2750220 <https://doi.org/10.1214/09-AOAS236>
- 1000 GENOMES PROJECT CONSORTIUM ET AL. (2015). A global reference for human genetic variation. *Nature* **526** 68.
- GAZAL, S., SAHBATOU, M., BABRON, M.-C., GÉNIN, E. and LEUTENEGGER, A.-L. (2015). High level of inbreeding in final phase of 1000 Genomes Project. *Sci. Rep.* **5** 17453. <https://doi.org/10.1038/srep17453>
- GRETTON, A. and GYÖRFI, L. (2010). Consistent nonparametric tests of independence. *J. Mach. Learn. Res.* **11** 1391–1423. MR2645456
- GUO, L. and MODARRES, R. (2020). Nonparametric tests of independence based on interpoint distances. *J. Nonparametr. Stat.* **32** 225–245. MR4077760 <https://doi.org/10.1080/10485252.2020.1714613>
- HAN, F., CHEN, S. and LIU, H. (2017). Distribution-free tests of independence in high dimensions. *Biometrika* **104** 813–828. MR3737306 <https://doi.org/10.1093/biomet/asx050>
- HELLER, R. and HELLER, Y. (2016). Multivariate tests of association based on univariate tests. *Adv. Neural Inf. Process. Syst.* **29**.
- HEMERIK, J. and GOEMAN, J. (2018). Exact testing with random permutations. *TEST* **27** 811–825. MR3878362 <https://doi.org/10.1007/s11749-017-0571-1>
- HERNÁN, M. A. and ROBINS, J. M. (2020). *Causal Inference: What If*. Chapman & Hall/CRC Press, Boca Raton.
- KALINA, J. and JANÁČEK, P. (2022). Testing exchangeability of multivariate distributions. *J. Appl. Stat.* 1–15.
- KIM, S. A., CHO, C.-S., KIM, S.-R., BULL, S. B. and YOO, Y. J. (2018). A new haplotype block detection method for dense genome sequencing data based on interval graph modeling of clusters of highly correlated SNPs. *Bioinformatics* **34** 388–397.
- KINGMAN, J. F. C. (1978). Uses of exchangeability. *Ann. Probab.* **6** 183–197. MR0494344 <https://doi.org/10.1214/aop/1176995566>
- KUCHIBHOTLA, A. K. (2020). Exchangeability, conformal prediction, and rank tests. Preprint. Available at [arXiv:2005.06095](https://arxiv.org/abs/2005.06095).
- LINDLEY, D. V. and NOVICK, M. R. (1981). The role of exchangeability in inference. *Ann. Statist.* **9** 45–58. MR0600531
- MAK, T. S. H., PORSCH, R. M., CHOI, S. W., ZHOU, X. and SHAM, P. C. (2017). Polygenic scores via penalized regression on summary statistics. *Genet. Epidemiol.* **41** 469–480.
- MANCUSO, N., FREUND, M. K., JOHNSON, R., SHI, H., KICHAEV, G., GUSEV, A. and PASANIUC, B. (2019). Probabilistic fine-mapping of transcriptome-wide association studies. *Nat. Genet.* **51** 675–682. <https://doi.org/10.1038/s41588-019-0367-1>
- MODARRES, R. (2008). Tests of bivariate exchangeability. *Int. Stat. Rev.* **76** 203–213.
- MOURAD, R., SINOQUET, C., DINA, C. and LERAY, P. (2011). Visualization of pairwise and multilocus linkage disequilibrium structure using latent forests. *PLoS ONE* **6** e27320. <https://doi.org/10.1371/journal.pone.0027320>
- PATTERSON, N. et al. (2006). Population structure and eigenanalysis. *PLoS Genet.* **2** e190.
- PFISTER, N., BÜHLMANN, P., SCHÖLKOPF, B. and PETERS, J. (2018). Kernel-based tests for joint independence. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **80** 5–31. MR3744710 <https://doi.org/10.1111/rssb.12235>

- PHIPSON, B. and SMYTH, G. K. (2010). Permutation p -values should never be zero: Calculating exact p -values when permutations are randomly drawn. *Stat. Appl. Genet. Mol. Biol.* **9** Art. 39, 14 pp. MR2746025 <https://doi.org/10.2202/1544-6115.1585>
- PRICE, A. L., WEALE, M. E., PATTERSON, N., MYERS, S. R., NEED, A. C., SHIANN, K. V., GE, D., ROTTER, J. I., TORRES, E. et al. (2008). Long-range LD can confound genome scans in admixed populations. *Am. J. Hum. Genet.* **83** 132–135.
- PRIVÉ, F. (2022). Optimal linkage disequilibrium splitting. *Bioinformatics* **38** 255–256.
- PRIVÉ, F., ARBEL, J., ASCHARD, H. and VILHJÁLMSSON, B. (2021). Identifying and correcting for misspecifications in GWAS summary statistics and polygenic scores. *HGG Advances* **3** 100136.
- SHAFFER, G. and VOVK, V. (2008). A tutorial on conformal prediction. *J. Mach. Learn. Res.* **9** 371–421. MR2417240
- SHIRTS, B. H., PRITCHARD, C. C. and WALSH, T. (2016). Family-specific variants and the limits of human genetics. *Trends Mol. Med.* **22** 925–934. <https://doi.org/10.1016/j.molmed.2016.09.007>
- SOSHNIKOV, A. (2002). A note on universality of the distribution of the largest eigenvalues in certain sample covariance matrices. *J. Stat. Phys.* **108** 1033–1056. MR1933444 <https://doi.org/10.1023/A:1019739414239>
- SPENCE, J. P., SINNOTT-ARMSTRONG, N., ASSIMES, T. and PRITCHARD, J. K. (2022). A flexible modeling and inference framework for estimating variant effect sizes from GWAS summary statistics. *bioRxiv*.
- SPENCE, J. P. and SONG, Y. S. (2019). Inference and analysis of population-specific fine-scale recombination maps across 26 diverse human populations. *Sci. Adv.* **5** eaaw9206.
- TCHETGEN TCHETGEN, E. J., YING, A., CUI, Y., SHI, X. and MIAO, W. (2020). An introduction to proximal causal learning. Preprint. Available at [arXiv:2009.10982](https://arxiv.org/abs/2009.10982).
- TRACY, C. A. and WIDOM, H. (2002). Distribution functions for largest eigenvalues and their applications. In *Proceedings of the International Congress of Mathematicians, Vol. I (Beijing, 2002)* 587–596. Higher Ed. Press, Beijing. MR1989209
- VOVK, V. (2021). Testing randomness online. *Statist. Sci.* **36** 595–611. MR4323055 <https://doi.org/10.1214/20-sts817>
- WALD, A. and WOLFOWITZ, J. (1943). An exact test for randomness in the non-parametric case based on serial correlation. *Ann. Math. Stat.* **14** 378–388. MR0009838 <https://doi.org/10.1214/aoms/1177731358>
- ZHOU, Y.-H., MARRON, J. S. and WRIGHT, F. A. (2018). Eigenvalue significance testing for genetic association. *Biometrics* **74** 439–447. MR3825330 <https://doi.org/10.1111/biom.12767>