# ON CONTAMINATION OF SYMBOLIC DATASETS* 

ANTONY PEARSON ${ }^{\dagger}$ AND MANUEL E. LLADSER ${ }^{\ddagger}$


#### Abstract

Data taking values on discrete sample spaces are the embodiment of modern biological research. "Omics" experiments produce millions of symbolic outcomes in the form of reads (i.e., DNA sequences of a few dozens to a few hundred nucleotides). Unfortunately, these intrinsically non-numerical datasets are often highly contaminated, and the possible sources of contamination are usually poorly characterized. This contrasts with numerical datasets where Gaussian-type noise is often well-justified. To overcome this hurdle, we introduce the notion of latent weight, which measures the largest expected fraction of samples from a contaminated probabilistic source that conform to a model in a well-structured class of desired models. We examine various properties of latent weights, which we specialize to the class of exchangeable probability distributions. As proof of concept, we analyze DNA methylation data from the 22 human autosome pairs. Contrary to what it is usually assumed, we provide strong evidence that highly specific methylation patterns are overrepresented at some genomic locations when contamination is taken into account.


Key words. Categorical data, Contamination, DNA methylation, Exchangeability, Hypothesis testing, p-value, Symbolic data, Truthiness

AMS subject classifications. 62-07, 62F03, 62F10, 62G10, 62P10, 92D20

1. Introduction. Symbolic data is the epitome modern biological datasets due to the advent of high-throughput sequencing assays. These assays generate millions of comparatively short DNA sequences of a few dozen to a few hundred nucleotides and allow scientists to assess various microscopic processes such as investigating the relative abundance of unculturable organisms in an environment [26, 14], or pinpointing the location of the enzymes that are actively transcribing DNA into RNA along a genome [24]-among many other possibilities [35, 6, 28, 32]. Unfortunately, these datasets are often very noisy, and it is often unclear how to describe the noise because the possible sources of data corruption can be so intricate that there is little motivation for any specific and let alone universal representation of it. In contrast, Gaussian errors are often well-justified with continuous numerical data.

There is a rich history of using mixtures to describe deviations from idealized continuous models [27, 33, 16, 17]. Such mixtures are usually of the form: $P=$ $(1-\epsilon) \cdot Q+\epsilon \cdot R$, where $P$ is the probabilistic source producing the data, $Q$ is a Gaussian distribution, and $R$ is some contaminating probability distribution from which an observation is drawn with some small probability $\epsilon$. When the mixture is unspecified, and $P$ ought to be estimated from data, a highly specific structure for $R$ is usually needed, e.g. a Gaussian with known mean, to make $\epsilon$ and $Q$ identifiable [29].

In this manuscript, we address the problem of assessing contamination (or its opposite, "purity,") in symbolic datasets-which are intrisically discrete. Like prior work on continuous data, we model contamination as a mixture. Unlike previous lines of work, we treat contamination as incidental; in particular, we do not commit to any prespecified form for it. To do so, we introduce the notion of latent weight with respect to a given structured class of probabilistic models (e.g., exchangeable probability distributions, which are the main focus of this manuscript). Broadly

[^0]speaking, this is the largest weight a model in the structured class can have as a component of the source producing the data. In particular, it describes the largest expected fraction of samples from a contaminated random sample which conform to a probabilistic model in the structured class.

We argue that latent weights are always identifiable, and allow one to represent unstructured probabilistic models as mixtures with a well-structured component; in particular, when this component carries a substantial latent weight, most samples from the mixture can be attributed to it. Latent weights offer therefore a measure of the truthiness of a hypothesis, which may not be strictly true.

To fix ideas, consider binary random variables $X$ and $Y$ with joint probability distribution given by the matrix

$$
P:=\left(\begin{array}{cc}
Y=0 & Y=1  \tag{1.1}\\
1 / 10 & 3 / 10 \\
1 / 10 & 1 / 2
\end{array}\right) \begin{aligned}
& x=0 \\
& x=1
\end{aligned} .
$$

Since $X \sim$ Bernoulli(3/5) and $Y \sim$ Bernoulli(4/5), it is straightforward to check that $X$ and $Y$ are not independent. Nevertheless, we can ask whether or not $P$ can be represented as a mixture with a component with independent marginals. This is related to determining if the latent weight of $P$ with respect to the class $\mathcal{Q}$ of product measures of the form $(\mu \otimes \nu)$, with $\mu$ and $\nu$ probability models supported on $\{0,1\}$, is positive or not (see Definition 2.1). It turns out the largest weight one can give to the model Bernoulli $(3 / 5) \otimes$ Bernoulli $(4 / 5)$ formed using the marginals of $(X, Y)$ is $5 / 6 \approx 83 \%$. Indeed:

$$
P=\frac{5}{6} \cdot \text { Bernoulli }(3 / 5) \otimes \text { Bernoulli }(4 / 5)+\frac{1}{6} \cdot\left(\begin{array}{cc}
1 / 5 & 1 / 5 \\
0 & 3 / 5
\end{array}\right) .
$$

The latent weight of $P$ w.r.t. $Q$ must be therefore at least $5 / 6$. In fact, a simple calculation reveals that

$$
P=\frac{24}{25} \cdot \text { Bernoulli }(5 / 8) \otimes \text { Bernoulli }(5 / 6)+\frac{1}{25} \cdot\left(\begin{array}{ll}
1 & 0 \\
0 & 0
\end{array}\right)
$$

i.e., the latent weight of $P$ with respect to $\mathcal{Q}$ is at least $24 / 25=96 \%$. Further analysis would show that it is precisely this; in particular, up to a hidden event with $96 \%$ probability, $X$ and $Y$ behave independently. This finding is noteworthy for many reasons.

On one hand, perhaps unexpectedly, the marginal distributions of $X$ and $Y$ are not associated with the latent weight of $P$ with respect to $\mathcal{Q}$. On the other hand, if we derived our beliefs about the independence of $X$ and $Y$ from a large sample from $P$, e.g. using a chi-squared test of independence, at moderate significance levels we would typically reject the hypothesis that $X$ and $Y$ are independent, missing that most of the time this is not the case. In this regard, latent weights could be used as a proxy for the truthiness of a hypothesis, without being tied to the absolutes of truth or falsity in classical hypothesis testing approaches. Finally, if $P$ were estimated from a large but corrupted sample, as is usually the case with modern biological datasets, the high weight of Bernoulli $(5 / 8) \otimes \operatorname{Bernoulli}(5 / 6)$ as a component of $P$ would suggest modeling $X$ and $Y$ as independent-as opposed to a more complicated model with a spurious correlation induced by a seemingly $4 \%$ contamination

Paper organization. Since our primary interest is on symbolic data, we restrict ourselves to the setting of finite sample spaces. From a technical point of view, this helps to develop the theory but without trivializing it. Section 2 introduces and analyzes the notion of latent weight associated with a general class of probability models. Sections 3-5 are exclusively devoted to the class of exchangeable probability models. Section 3.1 introduces some upper-bounds that may be useful for assessing an exchangeable latent weight when the sample space is too large relative to the number of observations from the model. Section 4 develops the statistical machinery necessary for inference of exchangeable latent weights when the probabilistic source (producing the data) is observed only indirectly through a random sample. Finally, Section 5 demonstrates the use of latent weights to assess the exchangeability of DNA methylation, a near-universal assumption in epigenomic analyses.
2. Latent weights. In what follows, $\mathcal{P}$ denotes the set of all probability measures over certain finite non-empty sample space $\Omega$. As such, $\mathcal{P}$ is compact under any norm induced metric; in particular, the total variation norm. Recall that for $P_{1}, P_{2} \in \mathcal{P}$, this norm is defined as [22]:

$$
\left\|P_{1}-P_{2}\right\|:=\max _{A \subset \Omega}\left|P_{1}(A)-P_{2}(A)\right|=\frac{1}{2} \sum_{x \in \Omega}\left|P_{1}(x)-P_{2}(x)\right|
$$

In what remains of this manuscript, $\mathcal{Q} \subset \mathcal{P}$ denotes a closed non-empty subset of probability measures. In particular, $\mathcal{Q}$ is a compact subset of $\mathcal{P}$. Instances like this include, for example, singletons, as well as models with independent marginals (when $\Omega$ is a product space), among various other possibilities.

Definition 2.1. Let $P \in \mathcal{P}$. We define the (latent) weight of $\mathcal{Q}$ in $P$ as the coefficient $\lambda_{\mathcal{Q}}(P):=\sup \{\lambda$ such that $P \geq \lambda \cdot Q$ for some $Q \in \mathcal{Q}\}$, where $P \geq \lambda \cdot Q$ means that $P(\omega) \geq \lambda \cdot Q(\omega)$, for all $\omega \in \Omega$.

This definition resembles that given in [5, 25] but for the very different purpose of representing a long-lasting Markov chain by shorter-lived independent chains.

Clearly, $0 \leq \lambda_{\mathcal{Q}}(P) \leq 1$. Latent weights have various other properties which we now state.

Theorem 2.2. $\lambda_{\mathcal{Q}}(P)=1$ if and only if $P \in \mathcal{Q}$.
Proof. If $\lambda_{\mathcal{Q}}(P)=1$ then there exists a sequence of real numbers $\left(\lambda_{n}\right)_{n \geq 1}$ such that $1-1 / n<\lambda_{n} \leq 1$, and $P \geq \lambda_{n} \cdot Q_{n}$ for some $Q_{n} \in \mathcal{Q}$. Without loss of generality assume that $\lambda_{n}<1$. In particular, if we define $R_{n}:=\left(P-\lambda_{n} \cdot Q_{n}\right) /\left(1-\lambda_{n}\right)$ then $R_{n} \in \mathcal{P}$ and $P=\lambda_{n} \cdot Q_{n}+\left(1-\lambda_{n}\right) \cdot R_{n}$. As a result: $\left(P-Q_{n}\right)=\left(1-\lambda_{n}\right) \cdot\left(R_{n}-Q_{n}\right)$, which implies that $\left\|P-Q_{n}\right\|=\left(1-\lambda_{n}\right) \cdot\left\|R_{n}-Q_{n}\right\|<1 / n$. Hence $P \in \mathcal{Q}$ because $\mathcal{Q}$ is closed. The converse is immediate because $\lambda_{\mathcal{Q}}(Q)=1$ for all $Q \in \mathcal{Q}$.

Theorem 2.3.

$$
\begin{equation*}
\lambda_{\mathcal{Q}}(P)=\sup _{Q \in \mathcal{Q}} \min _{\omega \in \Omega} \frac{P(\omega)}{Q(\omega)}, \tag{2.1}
\end{equation*}
$$

where any division by zero (including zero-over-zero) is to be interpreted as $+\infty$, and this supremum is achieved; in particular, there is $Q \in \mathcal{Q}$ such that $P \geq \lambda_{\mathcal{Q}}(P) \cdot Q$ and there is $R \in \mathcal{P}$ such that

$$
\begin{equation*}
P=\lambda_{\mathcal{Q}}(P) \cdot Q+\left(1-\lambda_{\mathcal{Q}}(P)\right) \cdot R . \tag{2.2}
\end{equation*}
$$

Further, if $\mathcal{Q}$ is convex and $\lambda_{\mathcal{Q}}(P)<1$ then $\lambda_{\mathcal{Q}}(R)=0$; in particular, $R \notin \mathcal{Q}$.

Proof. Fix $P \in \mathcal{P}$, and define $\lambda^{*}:=\sup _{Q \in \mathcal{Q}} \min _{\omega \in \Omega} P(\omega) / Q(\omega)$. For each $Q \in \mathcal{Q}$, note that $P \geq \lambda \cdot Q$ if and only if $\min _{\omega \in \Omega} P(\omega) / Q(\omega) \geq \lambda$; in particular, $\lambda^{*} \geq \lambda_{\mathcal{Q}}(P)$. Define $\epsilon:=\lambda^{*}-\lambda_{\mathcal{Q}}(P)$. If $\epsilon>0$ then, from the definition of $\lambda^{*}$, there would be $Q \in \mathcal{Q}$ such that $\min _{\omega \in \Omega} P(\omega) / Q(\omega) \geq \lambda^{*}-\epsilon / 2$. In particular, for all $\omega \in \Omega$, $P(\omega) \geq\left(\lambda^{*}-\epsilon / 2\right) Q(\omega)$, hence $\lambda_{\mathcal{Q}}(P) \geq \lambda^{*}-\epsilon / 2>\lambda_{\mathcal{Q}}(P)$, a contradiction. Therefore $\epsilon=0$ i.e. $\lambda^{*}=\lambda_{\mathcal{Q}}(P)$.

Next, note that for each $Q \in \mathcal{Q}$ :

$$
\begin{equation*}
\min _{\omega \in \Omega} \frac{P(\omega)}{Q(\omega)}=\min _{\omega: Q(\omega)>0} \frac{P(\omega)}{Q(\omega)} . \tag{2.3}
\end{equation*}
$$

Further, if $\lim _{n \rightarrow \infty} Q_{n}=Q$ in total variation distance then $\lim _{n \rightarrow \infty} Q_{n}(\omega)=Q(\omega)$ uniformly for all $\omega \in \Omega$. Hence, the transformation $Q \in \mathcal{Q} \rightarrow \min _{\omega \in \Omega} P(\omega) / Q(\omega)$ from $\mathcal{Q}$ to $\mathbb{R}$ is continuous; in particular, because $\mathcal{Q}$ is compact, the supremum in equation (2.1) is achieved, and there is $Q \in \mathcal{Q}$ such that $P \geq \lambda_{\mathcal{Q}}(P) \cdot Q$. For brevity, define $\lambda:=\lambda_{\mathcal{Q}}(P)$. Following an argument similar to the proof of Theorem 2.2, it follows that there is $R \in \mathcal{P}$ such that $P=\lambda \cdot Q+(1-\lambda) \cdot R$. Likewise, if $\alpha:=\lambda_{\mathcal{Q}}(R)$ then there is $Q^{\prime} \in \mathcal{Q}$ and $R^{\prime} \in \mathcal{P}$ such that $R=\alpha \cdot Q^{\prime}+(1-\alpha) \cdot R^{\prime}$. As a result, if $\mathcal{Q}$ is convex then $P \geq \lambda \cdot Q+(1-\lambda) \alpha \cdot Q^{\prime}=(\lambda+(1-\lambda) \alpha) \cdot Q^{\prime \prime}$ for some $Q^{\prime \prime} \in \mathcal{Q}$. But then $\lambda \geq \lambda+(1-\lambda) \alpha$, which implies that $(1-\lambda) \alpha=0$. Thus, if $\lambda<1$ then $\alpha=0$ as claimed.

We emphasize that the probability measure $Q$ in Theorem 2.3 is not necessarily unique. For instance, if $\mathcal{Q}$ is the space of probability measures over $\Omega=$ $\{0,1\}^{2}$ with i.i.d. marginals, and $P$ is the uniform probability measure over the set $\{(0,0),(1,1)\}$ then $\lambda_{\mathcal{Q}}(P)=1 / 2$, and $\delta_{(0,0)}$ and $\delta_{(1,1)}$, the point masses at $(0,0)$ and $(1,1)$, respectively, both achieve the supremum in equation (2.1). Likewise, since $P=\delta_{(0,0)} / 2+\delta_{(1,1)} / 2$, this counterexample shows the importance of convexity to guarantee that the probability measure $R$ in equation (2.2) does not belong to $\mathcal{Q}$.

The identity in equation (2.2) implies that any probability model $P$ over $\Omega$ admits a mixture representation with a component in $\mathcal{Q}$ of weight $\lambda_{\mathcal{Q}}(P)$. (This motivates the terminology of "latent weight.") The weight of $\mathcal{Q}$ in $P$ may be interpreted therefore as the largest expected fraction of observations from $P$ that can be attributed to a single model in $\mathcal{Q}$.

We also note that a large latent weight is indicative of closeness in total variation distance to $\mathcal{Q}$. In fact, if $Q$ and $R$ are as in equation (2.2) then (see proof of Theorem 2.2): $\|P-Q\|=\left(1-\lambda_{\mathcal{Q}}(P)\right) \cdot\|R-Q\|$. So, if $\lambda_{\mathcal{Q}}(P)$ is close to 1 then $P$ is close to $\mathcal{Q}$ in total variation distance. The converse is not necessarily true, however. For instance, consider $\Omega=\{0,1\}^{d}$, with finite $d>1$, and let $0^{d}$ and $1^{d}$ denote the sequences of $d$ zeros and ones, respectively. Define $P(x):=\left(2^{d}-2\right)^{-1}$ for $x \in \Omega \backslash\left\{0^{d}, 1^{d}\right\}$. In particular, $P\left(0^{d}\right)=P\left(1^{d}\right)=0$, and $\left\|P-\operatorname{Uniform}\left(\{0,1\}^{d}\right)\right\|=2^{2-d}$, i.e. $P$ is very close in total variation distance to $\mathcal{Q}$. Nevetheless, the latent weight of $P$ with respect to the class $\mathcal{Q}$ of i.i.d. distributions over $\{0,1\}^{d}$ is 0 . Indeed, for all $\lambda>0$ and $Q:=\otimes_{k=1}^{d} \operatorname{Bernoulli}(q)$, with $0 \leq q \leq 1$, it is not possible to have $P \geq \lambda \cdot Q$. Otherwise, because $Q\left(0^{d}\right)=(1-q)^{d}$ and $Q\left(1^{d}\right)=q^{d}$, it would follow that $q=1$ and $q=0$ simultaneously.
3. Exchangeable weights. The remaining of the manuscript focuses on the class of exchangeable probabilistic models. Recall that a finite sequence of random variables $X_{1}, \ldots, X_{d}$ is called exchangeable if for any permutation $\sigma$ of $(1, \ldots, d)$ the random vector $\left(X_{\sigma(1)}, \ldots, X_{\sigma(d)}\right)$ has the same distribution as $\left(X_{1}, \ldots, X_{d}\right)$.

Exchangeability is a common a priori assumption in Bayesian statistics [8], permutation hypothesis testing [12], and coalescent theory [20]. The class of exchangeable models contains all finite sequences of independent and identically distributed (i.i.d.) random variables, but is generally much larger. In fact, sampling colored balls from an urn without replacement produces a random sequence of colors which is exchangeable but not necessarily independent.

In what follows, $\left(X_{1}, \ldots, X_{d}\right)$, with $d>1$ finite, is a random vector with probability distribution $P$. Each $X_{i}$ is assumed to take values in a certain finite set $\mathcal{X}$ of cardinality $k>1$. In addition, $\mathcal{P}$ denotes the set of all probability models over $\mathcal{X}^{d}$, and $\mathcal{E} \subset \mathcal{P}$ the subset of exchangeable models, which is clearly closed. We refer to the latent weight of $P$ with respect to $\mathcal{E}, \lambda_{\mathcal{E}}(P)$, as the exchangeable weight of $P$.

Definition 3.1. For each $x \in \mathcal{X}^{d}$, let $[x] \subset \mathcal{X}^{d}$ denote the set of all vectors of the form $\left(x_{\sigma(1)}, \ldots, x_{\sigma(d)}\right)$, with $\sigma$ a permutation of $(1, \ldots, d)$.

The set $\mathcal{X}^{d}$ has $\binom{k+d-1}{d}$ permutation-equivalence classes. (The number of permutation equivalence classes equals the number of ways to place $d$ unlabelled balls in $k$ labelled urns.) In fact, $\mathcal{E}$ is a simplex, in particular, also a convex set, with $\binom{k+d-1}{d}$ extreme points, each of which is a probability model having uniform mass over a single permutation-equivalence class [9].

In what follows, $\left[\mathcal{X}^{d}\right]$ denotes the set of permutation equivalence classes of $\mathcal{X}^{d}$. Next we show various properties of exchangeable weights, starting with the following explicit formula.

Theorem 3.2. For all $P \in \mathcal{P}$ :

$$
\begin{equation*}
\lambda_{\mathcal{E}}(P)=\sum_{x \in \mathcal{X}^{d}} \min _{y \in[x]} P(y)=\sum_{z \in\left[\mathcal{X}^{d}\right]}|z| \cdot \min _{x \in z} P(x) \tag{3.1}
\end{equation*}
$$

If $\lambda_{\mathcal{E}}(P)>0$ then exactly one probability measure $Q \in \mathcal{E}$ is associated with the exchangeable weight of $P:$ for each $x \in \mathcal{X}^{d}, Q(x)=\min _{y \in[x]} P(y) / \lambda_{\mathcal{E}}(P)$.

Proof. Define $\lambda^{*}:=\sum_{x \in \mathcal{X}^{d}} \min _{y \in[x]} P(y)$. Suppose $\beta \geq 0$ and $Q^{\prime} \in \mathcal{E}$ are such that $P \geq \beta \cdot Q^{\prime}$. Since exchangeability implies $Q^{\prime}(y)=Q^{\prime}(x)$ for all $y \in[x]$, it follows that

$$
\beta \cdot Q^{\prime}(x) \leq \min _{y \in[x]} P(y)
$$

Moreover, because $Q^{\prime}$ is a probability measure, we have that

$$
\beta=\sum_{x \in \mathcal{X}^{d}} \beta \cdot Q^{\prime}(x) \leq \sum_{x \in \mathcal{X}^{d}} \min _{y \in[x]} P(y)=\lambda^{*}
$$

which implies that $\lambda_{\mathcal{E}}(P) \leq \lambda^{*}$. But observe that $P \geq \lambda^{*} \cdot Q$, where $Q(x):=$ $\min _{y \in[x]} P(y) / \lambda^{*}$. Since $Q$ is an exchangeable probability measure over $\mathcal{X}^{d}$, it follows that $\lambda_{\mathcal{E}}(P) \geq \lambda^{*}$, hence $\lambda_{\mathcal{E}}(P)=\lambda^{*}$. The second identity in equation (3.1) is now direct from this equality.

Finally, assume that $\lambda^{*}>0$ and suppose that $S \in \mathcal{E}$ is such that $P \geq \lambda^{*} \cdot S$. We show using an argument by contradiction that $S \leq Q$. Indeed, if there were $x \in \mathcal{X}^{d}$ such that $S(x)>Q(x)$ then, because $Q$ and $S$ are exchangeable, the definition of $Q$ would imply that there is $y \in[x]$ such that $S(y)>Q(y)$ and $\lambda^{*} \cdot Q(y)=P(y)$. In particular, $P(y)<\lambda^{*} \cdot S(y)$, which is not possible. Thus $S \leq Q$, which implies that $S=Q$ as claimed.

Due to Theorem 2.3, whenever $P$ is not itself exchangeable i.e. $\lambda_{\mathcal{E}}(P)<1, P$ also has a unique component $R:=(P-\lambda \cdot Q) /(1-\lambda)$, which we call the unexchangeable component, which admits no exchangeable component of its own. In this sense $P$ can be distilled entirely into its exchangeable and unexchangeable parts. This property allows one to combine an exchangeable probability model with one that is totally unexchangeable so that the resultant source has a desired exchangeable weight:

Corollary 3.3. If $P=\beta \cdot Q^{\prime}+(1-\beta) \cdot R^{\prime}$, where $0 \leq \beta \leq 1, Q^{\prime} \in \mathcal{E}$, and $R^{\prime} \in \mathcal{P} \backslash \mathcal{E}$ is such that $\lambda_{\mathcal{E}}\left(R^{\prime}\right)=0$, then $\lambda_{\mathcal{E}}(P)=\beta$.

Proof. Since $\lambda_{\mathcal{E}}\left(R^{\prime}\right)=0$, the identity in equation (3.1) implies that for each $z \in\left[\mathcal{X}^{d}\right]$ there is $y \in z$ such that $R^{\prime}(y)=0$. In particular, because $Q^{\prime}$ is exchangeable, reusing equation (3.1) we obtain that:

$$
\lambda_{\mathcal{E}}(P)=\sum_{x \in \mathcal{X}^{d}}\left(\beta \cdot Q^{\prime}(x)+(1-\beta) \cdot \min _{y \in[x]} R^{\prime}(y)\right)=\beta \cdot \sum_{x \in \mathcal{X}^{d}} Q^{\prime}(x)=\beta .
$$

3.1. Bounds on exchangeable weights. In what follows, for each non-empty $I \subset\{1, \ldots, d\}, P_{I}$ denotes the marginal distribution of $\left(X_{i}\right)_{i \in I}$. Clearly, $P_{I}$ is exchangeable when $P$ is exchangeable.

The following result may be useful to estimate an upper-bound on the exchangeable weight of a probabilistic source on $\mathcal{X}^{d}$. Indeed, estimating $\lambda_{\mathcal{E}}(P)$ when $P$ is known only indirectly through data requires estimating $\left(k^{d}-1\right)$ free parameters, which may be infeasible in practice. However, marginalizing $X$ to a random sub-vector of dimension $s<d$, one may significantly reduce the dimension of the estimation problem.

Theorem 3.4. The exchangeable weight of a full d-dimensional joint distribution is a lower bound on the exchangeable weight of any marginal, i.e. $\lambda_{\mathcal{E}}(P) \leq \lambda_{\mathcal{E}}\left(P_{I}\right)$, for all $P \in \mathcal{P}$ and non-empty $I \subset\{1, \ldots, d\}$.

In particular, if $\lambda_{\mathcal{E}}\left(P_{I}\right)$ is small for some $I$ then so is $\lambda_{\mathcal{E}}(P)$; in which case, only very few of the data produced by $P$ could be attributed to an exchangeable source.

Proof. For each $x \in \mathcal{X}^{|I|}$ and $\alpha \in \mathcal{X}^{d-|I|}$, let $x \alpha$ be the vector $y \in \mathcal{X}^{d}$ such that $\left(y_{i}\right)_{i \in I}=x$ and $\left(y_{i}\right)_{i \notin I}=\alpha$. If $Q$ is the exchangeable probability measure given in equation (3.1) then

$$
P_{I}(x)=\sum_{\alpha \in \mathcal{X}^{d-|I|}} P(x \alpha) \geq \sum_{\alpha \in \mathcal{\mathcal { X } ^ { d - | I | }}} \lambda_{\mathcal{E}}(P) \cdot Q(x \alpha)=\lambda_{\mathcal{E}}(P) \cdot Q_{I}(x),
$$

for each $x \in \mathcal{X}^{|I|}$. Hence, since $Q_{I}$ is exchangeable, $\lambda_{\mathcal{E}}(P) \leq \lambda_{\mathcal{E}}\left(P_{I}\right)$.
The tightness of the inequality in the theorem is related to the notion of extendibility [10]. Let $1 \leq s \leq d$. An (exchangeable) probability model $\mu$ on $\mathcal{X}^{s}$ is called $d$-extendible if there is an exchangeable probability measure $\nu$ on $\mathcal{X}^{d}$ and $I \subset\{1, \ldots, d\}$ of cardinality $s$ such that $\mu=\nu_{I}$. Necessary and sufficient conditions for extendibility can be found in [11].

Corollary 3.5. For any $P \in \mathcal{P}$ with $\lambda_{\mathcal{E}}(P)>0$, and $I \subset\{1, \ldots, d\}$, if $\lambda_{\mathcal{E}}(P)=$ $\lambda_{\mathcal{E}}\left(P_{I}\right)$ then the exchangeable component of $P_{I}$ is d-extendable.

Proof. Suppose that $\lambda_{\mathcal{E}}(P)=\lambda_{\mathcal{E}}\left(P_{I}\right)>0$, with $|I|=s$. Let $Q$ and $\tilde{Q}$ denote the exchangeable components of $P$ and $P_{I}$, respectively. From the proof of Theorem 3.4, $P_{I} \geq \lambda_{\mathcal{E}}(P) \cdot Q_{I}=\lambda_{\mathcal{E}}\left(P_{I}\right) \cdot Q_{I}$; in particular, due to Theorem 3.2, $\tilde{Q}=Q_{I}$.

The converse in the corollary is not necessarily true, however. For a counterexample, consider $P \in \mathcal{P}\left(\{0,1\}^{3}\right)$ such that $P(1,0,1)=P(1,1,0)=P(1,1,1)=1 / 3$; in particular, $\lambda(P)=1 / 3$. If $I=\{1,2\}$ then $P_{I}(1,0)=1 / 3$ and $P_{I}(1,1)=2 / 3$, hence $\lambda\left(P_{I}\right)=2 / 3$, and $\tilde{Q}=\delta_{(1,1)}$, which is 3 -extendible to $\delta_{(1,1,1)}$ yet $\lambda(P) \neq \lambda\left(P_{I}\right)$.

Finally, another way to bound the exchangeable weight of a probability measure on $\mathcal{X}^{d}$ is to lump states in $\mathcal{X}$ as follows.

Theorem 3.6. Let $\mathcal{Y}$ be a finite set and $r: \mathcal{X} \rightarrow \mathcal{Y}$ a function, and define $\Phi: \mathcal{X}^{d} \rightarrow \mathcal{Y}^{d}$ as $\Phi(x):=\left(r\left(x_{1}\right), \ldots, r\left(x_{d}\right)\right)$. Then, for each $P \in \mathcal{P}\left(\mathcal{X}^{d}\right), \lambda_{\mathcal{E}}(P) \leq$ $\lambda_{\mathcal{E}}\left(P \circ \Phi^{-1}\right)$, where $P \circ \Phi^{-1}$ is the forward measure of $P$ by $\Phi$.

Proof. If $Q$ denotes the exchangeable component of $P$ then

$$
\begin{aligned}
\lambda_{\mathcal{E}}\left(P \circ \Phi^{-1}\right) & =\sum_{[y] \subset \mathcal{Y}^{d}}|[y]| \cdot \min _{z \in[y]} P\left(\Phi^{-1}(z)\right) \\
& \geq \sum_{[y] \subset \mathcal{Y}^{d}}|[y]| \cdot \min _{z \in[y]} \lambda_{\mathcal{E}}(P) \cdot Q\left(\Phi^{-1}(z)\right) \\
& =\lambda_{\mathcal{E}}(P) \cdot \sum_{[y] \subset \mathcal{Y}^{d}}|[y]| \cdot \min _{z \in[y]} Q\left(\Phi^{-1}(z)\right)=\lambda_{\mathcal{E}}(P),
\end{aligned}
$$

where for the very last identity we have used that $Q$ is exchangeable; in particular, $\min _{z \in[y]} Q\left(\Phi^{-1}(z)\right)=Q\left(\Phi^{-1}(y)\right)$, for each $y \in \mathcal{Y}^{d}$.
4. Estimation of exchangeable weights. In this section, $X_{1}, \ldots, X_{n}$ denote $d$-dimensional i.i.d. samples from a probability measure $P$ defined over $\mathcal{X}^{d}$. Define $\lambda:=\lambda_{\mathcal{E}}(P)$.

Let $\hat{P}_{n}:=\sum_{i=1}^{n} \delta_{X_{i}} / n$ denote the empirical measure associated with the sample. A natural estimator of $\lambda$ is $\hat{\lambda}_{n}:=\lambda_{\mathcal{E}}\left(\hat{P}_{n}\right)$. Since $\hat{P}_{n}$ is a maximum likelihood estimator (MLE) of $P, \hat{\lambda}_{n}$ is a MLE of $\lambda$. Moreover, as noted in the proof of Theorem 2.3, $\lambda_{\mathcal{E}}(\cdot)$ is continuous. In particular, since $\hat{P}_{n} \rightarrow P$ almost surely, $\hat{\lambda}_{n} \rightarrow \lambda$ also almost surely. Since $0 \leq \hat{\lambda}_{n} \leq 1$ for all $n \geq 1$, it follows that $\hat{\lambda}_{n}$ is an asymptotically unbiased estimator of $\lambda$. Nevertheless, because for each $z \in\left[\mathcal{X}^{d}\right]$ the transformation $P \rightarrow \min _{y \in z} P(y)$ is concave down, equation (3.1) implies that $\lambda_{\mathcal{E}}(\cdot)$ is concave down. Thus, by Jensen's inequality, $E\left(\hat{\lambda}_{n}\right) \leq \lambda$, i.e. $\hat{\lambda}_{n}$ is a negatively biased estimator of $\lambda$.

For each $x \in \mathcal{X}^{d}$, consider the quantities

$$
m_{x}:=\min _{y \in[x]} P(y) ; \quad \sigma_{x}^{2}:=m_{x}\left(1-m_{x}\right)
$$

Since these quantities remain constant within each permutation equivalence class, we sometimes abuse the notation and write for $z \in\left[\mathcal{X}^{d}\right]: m_{z}$ and $\sigma_{z}^{2}$ to mean $m_{x}$ and $\sigma_{x}^{2}$, with $x \in z$, respectively.

Our next result characterizes implicitly the asymptotic distribution of $\hat{\lambda}_{n}$.
Theorem 4.1. If for each $z \in\left[\mathcal{X}^{d}\right], C_{z}:=\left\{x \in z\right.$ such that $\left.P(x)=m_{x}\right\}$ and $c_{z}:=\left|C_{z}\right|$ then

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \sqrt{n}\left(\hat{\lambda}_{n}-\lambda\right) \stackrel{d}{=} \sum_{z \in\left[\mathcal{X}^{d}\right]}|z| \cdot \min _{x \in C_{z}} Z_{z}^{(x)} \tag{4.1}
\end{equation*}
$$

where $\left(Z_{z}^{(x)}\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}$ is a normal random vector such that, for each $z \in\left[\mathcal{X}^{d}\right]$, $\left(Z_{z}^{(x)}\right)_{x \in C_{z}}$ is a $c_{z}$-dimensional zero-mean exchangeable normal random vector with
variance-covariance matrix $\Sigma_{z}$ such that $\Sigma_{z}(x, x)=\sigma_{z}^{2}$ and $\Sigma_{z}(x, y)=-m_{z}^{2}$, for all $x, y \in C_{z}$ with $x \neq y$, and for $z_{1}, z_{2} \in\left[\mathcal{X}^{d}\right]$ with $z_{1} \neq z_{2}, \operatorname{cov}\left(Z_{z_{1}}^{(x)}, Z_{z_{2}}^{(y)}\right)=-m_{z_{1}} m_{z_{2}}$, for all $x \in z_{1}$ and $y \in z_{2}$.

Proof. Our arguments follow closely those in [13].
For the sake of notation, we remove the sub-index $n$ from quantities defined in terms of $\hat{P}_{n}$; in particular, we write $\hat{P}$ instead of $\hat{P}_{n}$ and $\hat{\lambda}$ instead of $\hat{\lambda}_{n}$. For $x \in \mathcal{X}^{d}$, define $\hat{M}_{x}:=\min _{y \in C_{[x]}} \hat{P}(y)$ and $\hat{m}_{x}:=\min _{y \in[x]} \hat{P}(y)$. According to the Law of Large Numbers, $\hat{P} \rightarrow P$ almost surely; in particular:

$$
\begin{equation*}
\mathbb{P}\left(\text { there is } x \in \mathcal{X}^{d} \text { such that } \hat{M}_{x} \neq \hat{m}_{x}\right)=o(1) \tag{4.2}
\end{equation*}
$$

Define $\delta:=\sqrt{n}(\hat{\lambda}-\lambda)$ and $\Delta:=\sqrt{n}\left(\sum_{[x] \subset \mathcal{X}^{d}}|[x]| \cdot \min _{y \in C_{x}} \hat{P}(y)-\lambda\right)$. That is, $\Delta$ is proportional to the difference between $\lambda$ and an estimator of it computed from only those $\hat{P}(x)$ corresponding to outcomes $x \in \mathcal{X}^{d}$ for which $P(x)=m_{x}$.

Fix $t \in \mathbb{R}$. By equations (3.1) and (4.2),

$$
\begin{equation*}
|\mathbb{P}(\delta \leq t)-\mathbb{P}(\Delta \leq t)|=o(1) \tag{4.3}
\end{equation*}
$$

Furthermore, due to the well-known Central Limit Theorem for the multinomial distribution:

$$
\begin{align*}
\mathbb{P}(\Delta \leq t)=\mathbb{P} & \left(\sum_{[x] \subset \mathcal{X}^{d}}|[x]| \cdot \min _{y \in C_{x}} \sqrt{n} \cdot\left(\hat{P}(y)-m_{x}\right) \leq t\right) \\
& \longrightarrow \mathbb{P}\left(\sum_{[x] \subset \mathcal{X}^{d}}|[x]| \cdot \min _{y \in C_{[x]}} Z_{[x]}^{(y)} \leq t\right), \tag{4.4}
\end{align*}
$$

where $\left(Z_{z}^{(x)}\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}$ is a zero-mean normal random vector with variance-covariance matrix as described above. The theorem is now a direct consequence of equations (4.3) and (4.4).

Let $\mathcal{P}_{U} \subset \mathcal{P}$ denote the set of $P \in \mathcal{P}$ such that $c_{x}=1$ for each $x \in \mathcal{X}^{d}$, i.e. for each $z \in\left[\mathcal{X}^{d}\right]$ there is a unique $y \in z$ such that $P(y)=m_{z}$. The following result is an almost direct consequence of the previous theorem. This result also follows, albeit less directly, from the multivariate delta method [34].

Corollary 4.2. If $P \in \mathcal{P}_{U}$ then

$$
\lim _{n \rightarrow \infty} \sqrt{n}(\hat{\lambda}-\lambda) \stackrel{d}{=} Z
$$

where $Z$ is a zero-mean normal random variable with variance

$$
V(Z)=\sum_{[x] \subset \mathcal{X}^{d}}|[x]|^{2} \sigma_{x}^{2}-\sum_{[x] \neq[y]}|[x]||[y]| m_{x} m_{y} .
$$

Proof. Due to the hypothesis on $P$,

$$
\sqrt{n}(\hat{\lambda}-\lambda) \xrightarrow{d} Z=\sum_{[x] \subset \mathcal{X}^{d}}|[x]| \cdot Z_{[x]},
$$

where $\left(Z_{[x]}\right)_{[x] \subset \mathcal{X}^{d}}$ is an $\binom{k+d-1}{d}$-dimensional exchangeable normal random vector such that $E\left(Z_{[x]}\right)=0, V\left(Z_{[x]}\right)=m_{x}\left(1-m_{x}\right)$, and for $[x] \neq[y], \operatorname{cov}\left(Z_{[x]}, Z_{[y]}\right)=$ $-m_{x} m_{y}$. In particular, $Z$ has a normal distribution, from which the corollary follows.口

Let $\mathbb{X}:=\left(X_{1}, \ldots, X_{n}\right)$ denote an i.i.d. sample from $P \in \mathcal{P}, \mathbb{X}^{*}:=\left(X_{1}^{*}, \ldots, X_{n}^{*}\right)$ denote a single resample with replacement from $\mathbb{X}$, and $\hat{\lambda}^{*}$ denote the exchangeable weight associated with the empirical measure $\hat{P}^{*}:=\sum_{i=1}^{n} \delta_{X_{i}^{*}} / n$. In the next theorem and corollary we characterize the asymptotic distribution of the bootstrap distribution estimator of $\sqrt{n}\left(\hat{\lambda}^{*}-\hat{\lambda}\right)$. In what follows, $Z=\left(Z_{z}^{(x)}\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}$ denotes the normal random vector described in Theorem 4.1, which has dimension $\sum_{z \in\left[\mathcal{X}^{d}\right]} c_{z}$.

ThEOREM 4.3. Fix $t \in \mathbb{R}$, and associate with each vector $v:=\left(v_{z}^{(x)}\right)_{z \in[\mathcal{X} d], x \in C_{z}}$, the function:

$$
\psi(v):=\mathbb{P}\left(\sum_{z \subset\left[\mathcal{X}^{d}\right]}|z| \cdot\left[\min _{x \in C_{z}}\left(Z_{z}^{(x)}+v_{z}^{(x)}\right)-\min _{x \in C_{z}} v_{z}^{(x)}\right] \leq t\right)
$$

Note that $\psi(0)=\lim _{n \rightarrow \infty} \mathbb{P}(\sqrt{n}(\hat{\lambda}-\lambda) \leq t)$, the C.D.F. of the limiting random variable described in Theorem 4.1. The bootstrap estimator of $\psi(0)$ is

$$
\hat{\psi}:=\mathbb{P}\left(\sqrt{n}\left(\hat{\lambda}^{*}-\hat{\lambda}\right) \leq t \mid \mathbb{X}\right)
$$

and

$$
\begin{equation*}
\hat{\psi} \xrightarrow{d} \psi(Y), \tag{4.5}
\end{equation*}
$$

where $Y$ is an independent copy of $Z$.
The proof of this theorem resembles closely the arguments given in [13]. We first require the following lemma.

LEMMA 4.4. Define $\hat{M}_{x}^{*}:=\min _{y \in C_{x}} \hat{P}^{*}(y)$ and $\hat{m}_{x}^{*}:=\min _{y \in[x]} \hat{P}^{*}(y)$, i.e. $\hat{m}_{x}^{*}$ is the minimum probability estimated from a bootstrap resample in each equivalence class, and $\hat{M}_{x}^{*}$ is the same, estimated only from $C_{x}$. Then

$$
\mathbb{P}\left(\text { there is } x \in \mathcal{X}^{d} \text { such that } \hat{M}_{x}^{*} \neq \hat{m}_{x}^{*} \mid \mathbb{X}\right)=o_{p}(1) .
$$

Proof. If $P \in \mathcal{E}, \hat{M}_{x}=\hat{m}_{x}$, and $\hat{M}_{x}^{*}=\hat{m}_{x}$ with probability 1 for each $x$, and the claim follows.

Instead, for $P \in \mathcal{P} \backslash \mathcal{E}$, define

$$
\xi:=\min _{[x] \subset \mathcal{X}^{d}} \sin _{\text {s.t. }|[x]|<c_{x}}\left\{\min _{y \in[x] \backslash C_{x}} p_{y}-m_{x}\right\}
$$

that is, the smallest difference between $p_{x}$, for $x \notin C_{x}$, and $m_{x}$. Note that $\xi$ is positive. Similarly, define $\hat{\xi}:=\min _{[x] \subset \mathcal{X}^{d}}$ s.t. $|[x]|<c_{x}\left\{\min _{y \in[x] \backslash C_{x}} \hat{p}_{y}-\hat{M}_{x}\right\}$. (If $\hat{\xi}<0$, then $\hat{M}_{x}>\hat{m}_{x}$ for some $\left.[x].\right)$ The Law of Large Numbers guarantees that $\mathbb{P}(\hat{\xi} \geq$ $\xi / 2)=(1-o(1))$.

Define $Y_{n}:=\mathbb{P}\left(\right.$ there is $x \in \mathcal{X}^{d}$ such that $\left.\hat{M}_{x}^{*} \neq \hat{m}_{x}^{*} \mid \mathbb{X}\right)$, a random variable taking values in $[0,1]$. Fix $\epsilon>0$. Then

$$
\mathbb{P}\left(Y_{n} \geq \epsilon\right)=\mathbb{P}\left(Y_{n} \geq \epsilon \mid \hat{\xi} \geq \xi / 2\right) \cdot(1-o(1))+\mathbb{P}\left(Y_{n} \geq \epsilon \mid \hat{\xi}<\xi / 2\right) \cdot o(1)
$$

A second application of the Law of Large Numbers now yields $\mathbb{P}\left(Y_{n} \geq \epsilon \mid \hat{\xi} \geq \xi / 2\right) \rightarrow$ 0 , proving the lemma.

Proof of Theorem 4.3. Define

$$
\begin{aligned}
\delta^{*} & :=\sqrt{n}\left(\hat{\lambda}^{*}-\hat{\lambda}\right) \\
\Delta^{*} & :=\sqrt{n}\left(\sum_{z \in\left[\mathcal{X}^{d}\right]}|z| \cdot\left(\hat{M}_{z}^{*}-\hat{M}_{z}\right)\right) .
\end{aligned}
$$

Due to Lemma 4.4:

$$
\begin{equation*}
\mathbb{P}\left(\delta^{*} \leq t \mid \mathbb{X}\right)=\mathbb{P}\left(\Delta^{*} \leq t \mid \mathbb{X}\right)+o_{p}(1) \tag{4.6}
\end{equation*}
$$

Then it follows:

$$
\begin{align*}
\mathbb{P}\left(\Delta^{*} \leq t \mid \mathbb{X}\right)= & \mathbb{P}\left(\sqrt{n} \sum_{z \in\left[\mathcal{X}^{d}\right]}|z| \cdot\left(\hat{M}_{z}^{*}-\hat{M}_{z}\right) \leq t \mid \mathbb{X}\right) \\
= & \mathbb{P}\left(\sum _ { z \in [ \mathcal { X } ^ { d } ] } | z | \cdot \left[\min _{x \in C_{z}}\left\{\sqrt{n}\left(\hat{p}_{x}^{*}-\hat{p}_{x}\right)+\sqrt{n}\left(\hat{p}_{x}-m_{z}\right)\right\}\right.\right.  \tag{4.7}\\
& \left.\left.\quad-\min _{x \in C_{z}} \sqrt{n}\left(\hat{p}_{x}-m_{z}\right)\right] \leq t \mid \mathbb{X}\right) \\
= & \psi\left(\left(\sqrt{n}\left[\hat{p}_{x}-m_{z}\right]\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}\right)+o_{p}(1) . \tag{4.8}
\end{align*}
$$

The last equality is due to the fact that, for almost every sample sequence $X_{1}, X_{2}, \ldots$, $\sqrt{n}\left(\hat{P}^{*}-\hat{P}\right)$ has the same conditional limiting distribution given $\mathbb{X}$ that $\sqrt{n}(\hat{P}-P)$ does unconditionally, because $\hat{P}$ has finite second moments and mean $P$ [4, Theorem 2.2]. That is, for any $s \in \mathbb{R}^{k^{d}}, \mathbb{P}\left(\sqrt{n}\left(\hat{P}^{*}-\hat{P}\right) \leq s \mid \mathbb{X}\right) \rightarrow \lim _{n \rightarrow \infty} \mathbb{P}(\sqrt{n}(\hat{P}-P) \leq s)$ almost surely (and therefore in probability). For this reason, we can replace $\sqrt{n}\left(\hat{\bar{P}}^{*}-\right.$ $\hat{P}) \mid \mathbb{X}$ in line (4.7) by its limiting random vector, which introduces some random perturbation to the conditional C.D.F. which converges in probability to 0 .

Finally, because $\psi(v)$ is a continuous function of $v$, we take equation (4.8) with equation (4.6) to find that

$$
\mathbb{P}\left(\delta^{*} \leq t \mid \mathbb{X}\right)=\psi\left(\left(\sqrt{n}\left[\hat{p}_{x}-m_{z}\right]\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}\right)+o_{p}(1) \xrightarrow{d} \psi\left(\left(Y_{z}^{(x)}\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}\right)
$$

as claimed.
The above convergence is weak, however, by requiring $P \in \mathcal{P}_{U}$, we can ensure convergence in probability of the bootstrap distribution estimator. The following corollary follows simply from the fact that when $c_{z}=1$ for each $z \in\left[\mathcal{X}^{d}\right]$, $\min _{x \in C_{z}}\left(Z_{z}^{(x)}-Y_{z}^{(x)}\right)-\min _{x \in C_{z}} Y_{z}^{(x)}=Z_{z}^{(x)}$.

Corollary 4.5. When $P \in \mathcal{P}_{U}$,

$$
\mathbb{P}\left(\sqrt{n}\left(\hat{\lambda}^{*}-\hat{\lambda}\right) \leq t \mid \mathbb{X}\right) \xrightarrow{p} \mathbb{P}(\sqrt{n}(\hat{\lambda}-\lambda) \leq t)
$$

for each $t \in \mathbb{R}$.
Let $\mathcal{P}_{N}:=\mathcal{P} \backslash \mathcal{P}_{U}$ denote the set of sources which have at least one permutation class $z \in\left[\mathcal{X}^{d}\right]$ where $c_{z}>1$, i.e., sources for which at least one $m_{z}$ is not uniquely achieved by $y \in z$. When $P \in \mathcal{P}_{N}$, a more explicit characterization of the distribution
of $Z$ seems very elusive, and in particular, is not Gaussian. To see why, observe that [2, Corollary 5] implies that $\min _{x \in C_{z}} Z_{z}^{(x)}$ has probability density function (p.d.f.):

$$
f_{z}(t)=\frac{c_{z}}{\sigma_{z}} \cdot \varphi\left(\frac{t}{\sigma_{z}}\right) \cdot \Phi_{c_{z}-1}\left(\frac{t \sqrt{m_{z}}}{\sigma_{z}^{2}}, \ldots, \frac{t \sqrt{m_{z}}}{\sigma_{z}^{2}} ; \rho_{x} \mathbb{I}_{c_{z}-1}+\frac{\Sigma_{x}}{\sigma_{z}^{2}}\right)
$$

where $\varphi(\cdot)$ is the p.d.f. of a standard Normal random variable, $\Phi_{k}(\cdot ; \Sigma)$ is the cumulative distribution function (c.d.f.) of a zero-mean $k$-dimensional multivariate normal distribution with variance-covariance matrix $\Sigma$, and $\mathbb{I}_{k}$ is the $k$-dimensional identity matrix.

Unfortunately, the above probability densities are not enough to describe the distribution of $Z$ due to the correlation between the minima in equation (4.1). Furthermore, Theorem 4.3 does not hold when $P \in \mathcal{P}_{N}$, because whenever $c_{x}>1$, $\min _{x \in C_{z}}\left(Z_{z}^{(x)}+Y_{z}^{(x)}\right)-\min _{x \in C_{z}} Y_{z}^{(x)} \neq \min _{x \in C_{z}} Z_{z}^{(x)}$ with probability 1. As a result, the weak convergence in the last line of equation (4.8) is to a version of the distribution of $Z$, but with some Gaussian perturbation.

Luckily, however, the bootstrap estimator of $Z$ when $P \in \mathcal{P}_{U}$ can be made consistent by choosing a resample size $n_{0}$ of order $o(n)$. In this case, we would redefine $\delta^{*}:=\sqrt{n_{0}}\left(\hat{\lambda}_{n}^{*}-\hat{\lambda}\right)$ and $\Delta^{*}:=\sqrt{n_{0}}\left(\sum_{z \in\left[\mathcal{X}^{d}\right]}|z| \cdot\left(\hat{M}_{z}^{*}-\hat{M}_{z}\right)\right)$, so that

$$
\begin{aligned}
\hat{\psi} & =\mathbb{P}\left(\delta^{*} \leq t \mid \mathbb{X}\right)=\mathbb{P}\left(\Delta^{*} \leq t \mid \mathbb{X}\right)+o_{p}(1) \\
& =\psi\left(\left(\sqrt{n_{0}}\left(\hat{p}_{x}-p_{x}\right)_{z \in\left[\mathcal{X}^{d}\right], x \in C_{z}}\right)\right)+o_{p}(1) \\
& \xrightarrow{p} \psi(0),
\end{aligned}
$$

because $\sqrt{n_{0}}\left(\hat{p}_{x}-p_{x}\right) \xrightarrow{p} 0$.
Based on simulations, we have found that it is usually more accurate to use a full size- $n$ resample for the purposes of correcting the bias of $\hat{\lambda}_{n}$, even at moderate sample sizes and when $P \in \mathcal{P}_{N}$. Additionally, Monte Carlo bootstrap estimates of $V(\hat{\lambda})$ tend to be more accurate than the asymptotic formula given in derived from Corollary 4.2. With very large samples, users may wish to try using a size $n_{0}:=2 \sqrt{n}$ resample. Users might also explore ad hoc methods for combined estimators of $\min _{y \in[x]} p_{y}$ when there is strong reason to believe that $P \in \mathcal{P}_{N}$. In what follows in this manuscript, bias and variance of $\hat{\lambda}$ are approximated from a Monte Carlo estimate of the bootstrap distribution $\sqrt{n}\left(\hat{\lambda}_{n}^{*}-\hat{\lambda}\right) \mid \mathbb{X} \stackrel{d}{\approx} \sqrt{n}(\hat{\lambda}-\lambda)$.

An explicit Berry-Esseen type bound on the error using the limiting normal distribution to approximate the sampling distribution of $\hat{\lambda}_{n}$ remains elusive, therefore we recommend selecting sample size by simulation. The largest difficulty in estimating $\hat{\lambda}$ is controlling negative bias, especially when $P \in \mathcal{P}_{N}$. Therefore we recommend to simulate data from several test sources $T \in \mathcal{E}$ to approximate a worst-case sampling distribution for a given sample size. In particular, it is suitable to choose $T$ with $T\left(\{a\}^{d}\right)=0$ for each $a \in \mathcal{X}$, and uniform mass elsewhere, because each of the outcomes $\{a\}^{d}$ belongs to a singleton equivalence class. Any observation of these outcomes can only increase the estimate of $\lambda(T)$. In our experience on a variety of sample spaces $\mathcal{X}^{d}$, estimation of $\lambda(T)$ has the largest bias and standard deviation of any source in $\mathcal{P}\left(\mathcal{X}^{d}\right)$.

We suggest the following heuristic: first, select several candidate sample sizes $\left(n_{i}\right)_{i=1,2, \ldots}$. For each candidate sample size, repeatedly simulate $n_{i}$ outcomes from the test source $T$ described above to get an empirical estimate of standard error and


Fig. 1: Diagram of a typical WGBS experiment. The blue rectangle represents a segment of ssDNA, with the location of CpGs on that strand marked by black vertical bars. Black horizontal line segments represent reads mapped to a reference strand, open and closed circles represent the partially-observed joint methylation status several CpGs.
bias of $\lambda(\hat{T})$ (the test source above satisfies $\lambda(T)=1$ ). After selecting a sample size for which standard deviation and bias of $\hat{\lambda}_{n_{i}}(T)$ appear acceptably small, collect samples of this size for every source in a coarse grid over $\mathcal{P}\left(\mathcal{X}^{d}\right)$. This is to ensure that even on the most pathological sources we can obtain acceptable estimates of $\lambda$.
5. DNA methylation analysis. When a DNA sequence contains a cytosine residue ( C ) followed by a guanine residue ( G ) in the $5^{\prime}$-to- $3^{\prime}$ sense, this dimer is referred to as a CpG. The cytosine in a CpG may or may not have methyl group bonded to it in the $5^{\prime}$ position of its pyrimidine ring. This methylation is regulated by reversible enzymatic processes and is known to modulate gene expression; increased methylation in gene promoters is associated with transcriptional silencing [19], and specific DNA methylation patterns have been linked to human disease [30]. In particular, certain aberrant methylation patterns are a hallmark of some cancers [18], and as such considerable effort has been expended to determine regions of DNA that have differential methylation under different cellular conditions.

One popular modern assay to assess DNA methylation is Whole-Genome Bisulfite Sequencing (WGBS) [23], a procedure in which unmethylated cytosines are chemically transformed into thymine ( T ) through treatment with a bisulfite catalyst. When bisulfite-treated DNA is then sequenced by high-throughput shotgun technology, methylated CpGs can be distinguished from unmethylated ones by the observation of a "CG" dimer versus a "TG" dimer, as depicted in Figure 1.

When attempting to describe DNA methylation, it is routine to use a sliding window approach $[15,1]$ wherein all observations of methylated and unmethylated CpGs are counted in the window, typically 1 Kb in length, to summarize local methylation. It is common to compare methylation between two different biological samples using, e.g., Fisher's exact test [1], concluding that a window is differentially methylated if the null hypothesis of equal distribution can be rejected. This approach assumes that in a single window the methylation status of each CpG contributes identically in its biological effect. For example, if we denote an unmethylated CpG as a ' 0 ' and a methylated CpG as a ' 1 ' and consider ten consecutive CpGs in different tissues, with
the first tissue always producing the configuration ' 1111100000 ' and the second always producing the configuration ' 0000011111 ', the above approach would be unable to identify this locus as differentially methylated. So current approaches for differential methylation implicitly assume that binary sequences representing methylation inside each window have an exchangeable distribution.

To evaluate this assumption, we examined WGBS data from 121 experimental replicates representing 77 unique biological samples, publicly available from ENCODE [3, 7]. These replicates include clinical tissue samples, cell lines, and primary cells. We selected replicates using single-end reads which were not flagged by ENCODE as having low coverage or insufficient read length. The list of the sample identifiers (ENCODE_IDs.xlsx) and processed datasets can be found on GitHub (https://github.com/antonypearson/OnContaminationofSymbolicDatasets).

Each replicate is associated with a BAM file generated by mapping reads to GRCh38 using Bismark [21]. We used the MethPipe methylation software suite [31] to convert BAM files into MethPipe format and generate epiread files, an efficient format reporting the genomic index and methylation status of each CpG contained in a read.

To investigate the exchangeability of local DNA methylation we focused on sets of 3 consecutive CpGs, which we call "triplets." Due to Theorem 3.4, a genomic region containing a highly unexchangeable triplet must have highly unexchangeable methylation overall.

For each replicate, we used the epireads file generated to extract data from "wellcovered" triplets - i.e. those where all three CpGs are jointly covered by at least 100 reads, discarding reads which report a CpG with ambiguous methylation status. In all datasets, we observed 637,612 well-covered triplets, representing 72,815 unique loci. For each autosome in each sample we estimated the exchangeable weight of each well-covered triplet and corrected for estimator bias using a sample mean of $N=1000$ full bootstrap resamples. Although each estimate of a triplet's exchangeable weight $\hat{\lambda}$ lies in $[0,1]$, the bias-adjusted estimate $\left(\hat{\lambda}-\overline{{ }_{\lambda}^{*}}\right)$ may be larger than 1 or smaller than 0 . Therefore we truncate these estimates to $[0,1]$. Available online are Numpy files containing processed triplets corresponding to each BAM file ID. Each row corresponds to a well-covered triplet, with columns corresponding to 1) chromosome number, 2) index of the triplet on the chromosome, 3) an estimate of the total variation distance to the class of exchangeable distributions, 4) an estimate of the exchangeable weight of the triplet (bias-corrected), 5) a bootstrap estimate of the standard deviation of $\hat{\lambda}, 6-13$ ) the counts of each of the 8 possible triplet configurations (ordered lexicographically, i.e. ' 000 ', ' 001 ', etc.), and 14-21) an estimate of the largest exchangeable component.

Estimates of triplet exchangeable weight are depicted in Figure 2. As seen in the figure, in some chromosomal regions, particularly e.g. on chromosomes 6 and 13, there are triplets whose exchangeable weights are very small. In fact, some appear completely unexchangeable.

As seem in Table 1, triplet exchangeability does not appear strongly correlated with the genomic distance to the nearest promoter. Further, as seen in Figure 3, within each chromosome, and within each dataset, the correlation between triplet exchangeability and distance from a promoter is usually small. There is a noticeable trend, however, that triplet exchangeabile weight is more likely to be negatively correlated with distance from a promoter. Indeed, both a two-sided Wald test and a Spearman rank-order test of the null hypothesis that TSS proximity and estimated


Fig. 2: Estimated exchangeable weights of well-covered triplets by chromosome. Dashed red lines denote the mean, and green plots the histograms associated with these weights.
triplet exchangeable weight are uncorrelated give very small p-values $\left(p \ll 10^{-10}\right)$. That is, despite the small magnitude of the effect, we can detect that triplets close to promoters tend to have more-exchangeable methylation configurations.

Table 1: Correlation per chromosome between estimated exchangeable weight of wellcovered triplets, and distance between their center and the nearest transcription start site (TSS).

| Chromosome | Correlation |  | Chromosome | Correlation |
| :---: | :---: | :---: | :---: | :---: |
| 1 | $-1.200 \times 10^{-1}$ |  | 12 | $+2.676 \times 10^{-2}$ |
| 2 | $+1.236 \times 10^{-1}$ |  | 13 | $-2.406 \times 10^{-1}$ |
| 3 | $+4.738 \times 10^{-2}$ |  | 14 | $-2.977 \times 10^{-1}$ |
| 4 | $-1.742 \times 10^{-1}$ |  | 15 | $-1.999 \times 10^{-1}$ |
| 5 | $-1.584 \times 10^{-1}$ | 16 | $+4.356 \times 10^{-2}$ |  |
| 6 | $+7.900 \times 10^{-2}$ | 17 | $+9.088 \times 10^{-2}$ |  |
| 7 | $+7.423 \times 10^{-2}$ | 18 | $+8.553 \times 10^{-2}$ |  |
| 8 | $+2.869 \times 10^{-2}$ | 19 | $-6.973 \times 10^{-2}$ |  |
| 9 | $+2.586 \times 10^{-2}$ | 20 | $-8.210 \times 10^{-2}$ |  |
| 10 | $-3.123 \times 10^{-2}$ | 21 | $+6.563 \times 10^{-2}$ |  |
| 11 | $+5.987 \times 10^{-2}$ | 22 | $-2.552 \times 10^{-1}$ |  |




By chromosome


By ENCODE dataset

Fig. 3: Top, plot of estimated exchangeable weight for each well-covered triplet versus its distance from the nearest TSS. Bottom, correlation between distance to the nearest TSS and exchangeable weight of each triplet, per chromosome (left) and per dataset (right).

To confirm that these highly unexchangeable loci are not due to uncertainty in estimation of the exchangeable weight, we simulated data from the uniform distribution over binary triplets except ' 000 ' or ' 111 '. Based on empirical study, this source is the worst case for estimating the exchangeable weight in terms of bias and standard error. Nevertheless, as seen in Figure 4, the empirical distribution of estimated exchangeable weights of all triplets gives much greater probability mass near 0 than the corresponding sampling distribution of the synthetic data. That is, uncertainty from statistical estimation does not account for the apparent phenomenon of highly unexchangeable loci.

In general, it is impossible to disentangle contamination which is caused by, e.g., sequencing errors or incomplete enzymatic conversion of unmethylated cytosines, from biological processes discriminating specific configurations of methylation. Under the assumption that contamination of the former kind is small, i.e. that we have a truly accurate picture of how methylation is configured in cells, we would expect triplets to have exchangeable weights close to one if overall methylation levels govern biological function. This might mean that in some cell types methylation far away from


Fig. 4: Histogram of estimated exchangeable weights of all triplets from aggregated WGBS data (green) and synthetic samples of size $n=100$ from the uniform distribution on $\{(0,0,1),(0,1,0),(0,1,1),(1,0,0),(1,0,1),(1,1,0)\}$ (red). The averages of estimated exchangeable weights from real and synthetic data are given by the dashed green and red lines, respectively.
promoters (and likely far from CpG islands) is "locked in," and specific patterns of methylation rather than overall methylation levels modulate biological function.

Identifying the biological reason for highly unexchangeable loci remains an open question, which may not have a universal answer. We conclude that there are some loci which are far from exchangeable - that is, some configurations of methylation are discriminated at these triplets. The identification of these loci opens opportunities for more high-resolution understanding of methylation patterns. In particular, these loci represent regions where very specific configurations of methylation may regulate function.

## REFERENCES

[1] A. Akalin, M. Kormaksson, S. Li, F. E. Garrett-Bakelman, M. E. Figueroa, A. MelNICK, AND C. E. MASON, methylKit: a comprehensive $R$ package for the analysis of genome-wide DNA methylation profiles, Genome Biology, 13 (2012), p. R87.
[2] R. B. Arellano-Valle and M. G. Genton, On the exact distribution of the maximum of absolutely continuous dependent random variables, Statistics \& Probability Letters, 78 (2008), pp. 27-35.
[3] B. Bernstein, E. Birney, I. Dunham, E. Green, C. Gunter, M. Snyder, EnCODE Project Consortium, and T. Hubbard, An integrated encyclopedia of DNA elements in the human genome, NATURE, 489 (2012), pp. 57-74.
[4] P. J. Bickel and D. A. Freedman, Some asymptotic theory for the bootstrap, Ann. Statist., 9 (1981), pp. 1196-1217.
[5] S. Chestnut and M. E. Lladser, Occupancy distributions via Doeblin's ergodicity coefficient, in Discrete Mathematics and Theoretical Computer Science Proceedings, vol. AM, 2010, pp. 79-92.
[6] L. J. Core, J. J. Waterfall, and J. T. Lis, Nascent RNA sequencing reveals widespread
pausing and divergent initiation at human promoters, Science, 322 (2008), pp. 1845-1848.
[7] C. A. Davis, B. C. Hitz, C. A. Sloan, E. T. Chan, J. M. Davidson, I. Gabdank, J. A. Hilton, K. Jain, U. K. Baymuradov, A. K. Narayanan, K. C. Onate, K. Graham, S. R. Miyasato, T. R. Dreszer, J. S. Strattan, O. Jolanki, F. Y. Tanaka, and J. M. Cherry, The encyclopedia of DNA elements (ENCODE): data portal update, Nucleic Acids Research, 46 (2017), pp. D794-D801.
[8] B. de Finetti, La prévision: ses lois logiques, ses sources subjectives, Annales de l'institut Henri Poincaré, 7 (1937), pp. 1-68.
[9] P. Diaconis, Finite forms of de Finetti's theorem on exchangeability, Synthese, 36 (1977), pp. 271-281.
[10] P. Diaconis and D. Freedman, Finite exchangeable sequences, Ann. Probab., 8 (1980), pp. 745-764.
[11] A. V. Gnedin, A class of exchangeable sequences, Statistics \& Probability Letters, 28 (1996), pp. $159-164$.
[12] P. I. Good, Extensions of the concept of exchangeability and their applications, Journal of Modern Applied Statistical Methods, 1 (2002), p. 34.
[13] P. Hall, W. Härdle, and L. Simar, On the inconsistency of bootstrap distribution estimators, CORE Discussion Papers RP 1062, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 1993, https://EconPapers.repec.org/RePEc: cor:louvrp:1062.
[14] J. Hampton and M. E. Lladser, Estimation of Distribution Overlap of Urn Models, PLoS ONE, 7 (2012), p. e42368.
[15] K. D. Hansen, B. Langmead, and R. A. Irizarry, BSmooth: from whole genome bisulfite sequencing reads to differentially methylated regions, Genome Biology, 13 (2012), p. R83.
[16] P. J. Huber, Robust estimation of a location parameter, Ann. Math. Statist., 35 (1964), pp. 73101.
[17] P. J. Huber, A robust version of the probability ratio test, Ann. Math. Statist., 36 (1965), pp. 1753-1758.
[18] P. A. Jones, Functions of DNA methylation: islands, start sites, gene bodies and beyond, Nature Reviews Genetics, 13 (2012), pp. 484-492.
[19] P. L. Jones, G. C. J. Veenstra, P. A. Wade, D. Vermaak, S. U. Kass, N. Landsberger, J. Strouboulis, And A. P. Wolffe, Methylated dna and mecp2 recruit histone deacetylase to repress transcription, Nature genetics, 19 (1998), p. 187.
[20] J. F. C. Kingman, On the genealogy of large populations, Journal of Applied Probability, 19 (1982), p. 27-43.
[21] F. Krueger and S. R. Andrews, Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications, Bioinformatics, 27 (2011), pp. 1571-1572.
[22] T. Lindvall, Lectures on the Coupling Method, Wiley Series in Probability and Statistics Applied Probability and Statistics Section, Wiley, 1992.
[23] R. Lister, R. C. O’Malley, J. Tonti-Filippini, B. D. Gregory, C. C. Berry, A. H. MilLAR, AND J. R. Ecker, Highly integrated single-base resolution maps of the epigenome in arabidopsis, Cell, 133 (2008), pp. 523-536.
[24] M. E. Lladser, J. G. Azofeifa, M. A. Allen, and R. D. Dowell, RNA Pol II transcription model and interpretation of GRO-seq data, Journal of Mathematical Biology, 74 (2017), pp. 77-97.
[25] M. E. Lladser and S. Chestnut, Approximation of sojourn-times via maximal couplings: Motif frequency distributions, Journal of mathematical biology, 69 (2013).
[26] M. E. Lladser, R. Goevet, and J. Reeder, Extrapolation of Urn Models via Poissonization: Accurate Measurements of the Microbial Unknown, PLoS One, 6 (2011).
[27] S. Newcomb, A generalized theory of the combination of observations so as to obtain the best result, American Journal of Mathematics, 8 (1886), pp. 343-366.
[28] P. J. PARk, Chip-seq: advantages and challenges of a maturing technology, Nature reviews genetics, 10 (2009), p. 669.
[29] A. Punzo and P. D. McNicholas, Parsimonious mixtures of multivariate contaminated normal distributions, Biometrical Journal, 58 (2016), pp. 1506-1537.
[30] K. D. Robertson, DNA methylation and human disease, Nature Reviews Genetics, 6 (2005), p. 597.
[31] Q. Song, B. Decato, E. E. Hong, M. Zhou, F. Fang, J. Qu, T. Garvin, M. Kessler, J. Zhou, and A. D. Smith, A reference methylome database and analysis pipeline to facilitate integrative and comparative epigenomics, PLoS ONE, 8 (2013), p. e81148.
[32] M. M. Suzuki and A. Bird, DNA methylation landscapes: provocative insights from epigenomics, Nature Reviews Genetics, 9 (2008), p. 465.
[33] J. W. Tukey, A Survey of Sampling From Contaminated Distributions, Contributions to Probability and Statistics (in: I. Olkin et al., eds), (1960), pp. 448-485.
[34] A. W. van der Vaart, Asymptotic statistics, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 1998.
[35] Z. Wang, M. Gerstein, and S. M., RNA-Seq: a revolutionary tool for transcriptomics, Nat Rev Genet, 10 (2009), pp. 57-63.


[^0]:    *Submitted to the editors DATE.
    Funding: This work was partially supported by the NSF Graduate Research Fellowship Program grant No. 2016198773 (Pearson), and the NSF IGERT grant No. 1144807.
    ${ }^{\dagger}$ Department of Applied Mathematics, University of Colorado, Boulder, CO.
    ${ }^{\ddagger}$ Department of Applied Mathematics, University of Colorado, Boulder, CO (emailmanuel.lladser@colorado.edu).

