



Unbiased estimation of sampling variance for Simpson's diversity index

Andreas Tiffeau-Mayer 

Division of Infection and Immunity & Institute for the Physics of Living Systems, University College London, London WC1E 6BT, United Kingdom

 (Received 6 October 2023; revised 25 February 2024; accepted 3 June 2024; published 21 June 2024)

Quantification of measurement uncertainty is crucial for robust scientific inference, yet accurate estimates of this uncertainty remain elusive for ecological measures of diversity. Here, we address this longstanding challenge by deriving a closed-form unbiased estimator for the sampling variance of Simpson's diversity index. In numerical tests the estimator consistently outperforms existing approaches, particularly for applications in which species richness exceeds sample size. We apply the estimator to quantify biodiversity loss in marine ecosystems and to demonstrate ligand-dependent contributions of T-cell-receptor chains to specificity, illustrating its versatility across fields. The novel estimator provides researchers with a reliable method for comparing diversity between samples, essential for quantifying biodiversity trends and making informed conservation decisions.

DOI: [10.1103/PhysRevE.109.064411](https://doi.org/10.1103/PhysRevE.109.064411)

Living systems are characterized by immense diversity across multiple scales from molecules to ecosystems [1–5]. Quantitatively understanding how this diversity is produced and supports biological function has been a central question in the physics of living systems: On the ecosystem scale, statistical physics approaches have, for instance, identified conditions under which diverse interacting species can be stably maintained [1–3], while on the molecular and cellular scale, probabilistic modeling has shed light on how antibody and T-cell-receptor diversity in the adaptive immune system are generated [4–7].

To compare predictions from ecological theory to experimentally measured diversities requires quantification of measurement uncertainty; similarly comparing changes in biodiversity in response to habitat loss or climate change requires determining whether diversity differs more than expected by sampling chance alone [8–11]. A number of estimators have been proposed to quantify sampling variance in diversity estimation [12–17]. However, none of the current estimators are unbiased outside of asymptotically large samples, and numerical tests suggest that their finite sample bias can be severe. This is a practically important limitation, as overestimation of sampling uncertainty diminishes the ability to detect true trends in biodiversity. Reversely, underestimating sampling variance can inflate the apparent significance of observed changes in diversity, potentially leading to spurious conclusions.

Here, we address this gap by deriving an unbiased estimator of the sampling variance of Simpson's index. We show the superior performance of the estimator compared to previous approaches on simulated and real data and apply it to estab-

lish ligand-dependent differences in T-cell-receptor diversity. Simpson's index is used widely from ecology [12,18] and microbiology [14,19] to economics [20] as a measure of species diversity. The index is defined as the probability of species coincidence of a random pair of individuals,

$$p_C = \sum_{i=1}^S p_i^2, \quad (1)$$

where p_i is the frequency of species i in the population $i = 1, \dots, S$ and S is the number of distinct species. This index can be converted into an effective number of species $D = 1/p_C$ [18]. Importantly, the unbiased estimator introduced here provides interval estimates for diversity that do not depend systematically on sample size. It thus addresses a recently highlighted gap in methods [10] to compare diversity estimates between samples of different sizes without rarefaction.

To illustrate the problem setting, we reanalyze data from Albano *et al.* [11] (Fig. 1), who considered the following question: Has climate change led to a biodiversity loss in the Mediterranean Sea? Figure 1(a) shows the raw data, the number n_i of counted molluscs belonging to species i in a patch on the sea floor (top panel) and corresponding counts for empty mollusc shells found at the same site in death assemblages (bottom panel). A diversity index turns these species counts into scalar measures of current and past biodiversity [points in Fig. 1(b)]. However, less than 500 molluscs were counted across 65 distinct species, so some sampling variability in the diversity estimates is expected. The method we introduce allows robust quantification of this sampling variance [error bars in Fig. 1(b)]. Whereas the primary study assessed the statistical significance of biodiversity loss by comparisons across multiple independent sampling sites [11], adding confidence intervals to the biodiversity estimates shows the strong statistical support for biodiversity collapse that exists even at the level of an individual site.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

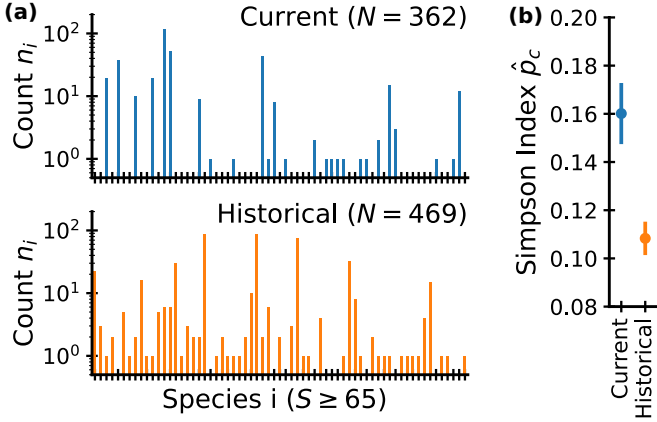


FIG. 1. Quantifying collapse of marine biodiversity in the Eastern Mediterranean Sea. Illustration of the problem setting of uncertainty quantification for finite sample diversity estimates. (a) Distributions of sampled mollusc species currently alive (top panel) and found in surficial dead assemblages (bottom panel). (b) Simpson Diversity index with error bars calculated using the proposed method [Eq. (4)]. Error bars show $\hat{p}_C \pm \sqrt{\text{Var}(\hat{p}_C)}$ and demonstrate statistical significance of the difference in diversity. Data: Noninvasive mollusc species found on rocky subtidal ground at 12-m depth off the coast of Ashqelon [11].

I. THE UNBIASED ESTIMATOR

In 1949 Simpson published a short letter in Nature that proposed the index that now bears his name [12]. In the same publication Simpson also showed that

$$\hat{p}_C = \frac{\sum_{i=1}^S n_i(n_i - 1)}{N(N - 1)} \quad (2)$$

provides an unbiased estimate of underlying population diversity from a finite sample of size N . Here n_i with $\sum_{i=1}^S n_i = N$ is the number of counts of the i th species in the sample, which follows a multinomial distribution,

$$P(n_1, \dots, n_S) = \frac{N!}{\prod_{i=1}^S n_i!} \prod_{i=1}^S p_i^{n_i}, \quad (3)$$

under the commonly used assumption that each individual from the population is sampled with equal probability [12,14,15].

Here we propose that the variance $\text{Var}(\hat{p}_C)$ of the point estimate [Eq. (2)] can be calculated without bias using the following closed-form estimator:

$$\hat{\text{Var}}(\hat{p}_C) = \frac{a}{1-b} \hat{p}_T - \frac{b}{1-b} \hat{p}_C^2 + \frac{c}{1-b} \hat{p}_C, \quad (4)$$

where

$$\hat{p}_T = \frac{\sum_{i=1}^S n_i(n_i - 1)(n_i - 2)}{N(N - 1)(N - 2)} \quad (5)$$

and

$$a = \frac{4(N - 2)}{N(N - 1)}, \quad b = \frac{2(2N - 3)}{N(N - 1)}, \quad c = \frac{2}{N(N - 1)}. \quad (6)$$

II. BACKGROUND AND DERIVATION

For the following derivation, it is instructive to recall the insight leading to the unbiased point estimator given in Eq. (8). To estimate p_C from a sample, it is tempting to simply replace the population frequencies p_i in Eq. (1) with the sampled frequencies $f_i = n_i/N$. However, one is well advised to resist this temptation as such plug-in estimators are known to be severely biased in small samples [12,13,15,21]. Instead Eq. (2) should be used, which is the probability of coincidence when drawing pairs of items from the sample without replacement. This estimator of p_C is unbiased, i.e., $\langle \hat{p}_C \rangle = p_C$, where $\langle \cdot \rangle$ is an average over repeated samples of a fixed size. Evaluating the expectation requires calculating the factorial moment $\langle n_i(n_i - 1) \rangle$, which can be calculated most conveniently from the probability generating function $G(z_1, \dots, z_S) = \langle \prod_{i=1}^S z_i^{n_i} \rangle$. From the definition of G it follows that $\frac{\partial}{\partial z_i^2} G(1, \dots, 1) = \langle n_i(n_i - 1) \rangle$. For the multinomial distribution, $G(z_1, \dots, z_S) = (\sum_{i=1}^S p_i z_i)^N$, and thus

$$\langle n_i(n_i - 1) \rangle = N(N - 1)p_i^2, \quad (7)$$

which can be plugged into Eq. (2) to complete the proof that \hat{p}_C is unbiased.

We now turn to reviewing the state of the art for variance estimation for Simpson's index. We first recall the formula for the sampling variance of the point estimator was again already given by Simpson [12] (derived in Appendix A for completeness),

$$\text{Var}(\hat{p}_C) = ap_T - bp_C^2 + cp_C. \quad (8)$$

This estimator is a linear combination of the triplet coincidence probability,

$$p_T = \sum_{i=1}^S p_i^3, \quad (9)$$

the square of the coincidence probability p_C^2 and p_C with sample-size-dependent parameters a , b , and c given in Eq. (6). Based on this formula, Grundmann *et al.* [14] proposed estimating variance by plugging the empirical frequencies $f_i = n_i/N$ into an asymptotic expansion of Eq. (8) for $N \rightarrow \infty$,

$$\hat{\text{Var}}_{\text{Grundmann}}(\hat{p}_C) = \frac{4}{N} \left[\sum_{i=1}^S f_i^3 - \left(\sum_{i=1}^S f_i^2 \right)^2 \right]. \quad (10)$$

However, we will find that this plug-in estimator is substantially biased for small N , similarly to plug-in estimators of diversity indices themselves. We also find substantial biases for the nonasymptotic plug-in estimator

$$\hat{\text{Var}}_{\text{plug-in}} = a \sum_{i=1}^S f_i^3 - b \left(\sum_{i=1}^S f_i^2 \right)^2 + c \sum_{i=1}^S f_i^2. \quad (11)$$

In addition to these plug-in estimators, another popular approach is due to Chao *et al.* [15]. This method is widely used in the field due to its implementation in the R package INEXT [22]. Chao *et al.*'s method estimates variances by bootstrapping from a population constructed from the sample by coverage-reweighting observed species frequencies and by augmenting the sample with an estimated number of rare




FIG. 2. Benchmarking of variance estimators on simulated data. (a) Frequency-rank plot of species abundances. Probabilities of $S = 1000$ species were drawn from the steady-state Dirichlet distribution of a neutral model with immigration and stochastic drift (parameter $\alpha = 1$). (b) Bias and (c) variance as a function of sample size N . Expectation values were calculated over 1000 repeated draws at each sample size from the population distribution. Bias and variance are expressed as fractions, i.e., divided by the true value or its square, respectively. The shaded areas differentiate sampling regimes: blue $N < \sqrt{2S}$, orange $\sqrt{2S} < N < S$, and green $N > S$. The number of bootstrap samples for Chao *et al.*'s method was set to 200 as recommended [15].

unseen species. Surprisingly, despite its widespread use, we find that this estimator has the largest bias and variance of tested methods.

Can we generalize the coincidence counting approach underlying the unbiased point estimate to the problem of interval estimation? To derive a better estimator for the variance of Simpson's index we exploit the linearity of Eq. (8) and decompose the problem into the unbiased estimation of each of the three terms. For the first term, the analogy to \hat{p}_C suggests to estimate the triplet probability using Eq. (5). To show that this estimator is indeed unbiased requires calculating the third factorial moment $\langle n_i(n_i - 1)(n_i - 2) \rangle$. This factorial moment can again be calculated by taking derivatives of the probability generating function G of the multinomial distribution,

$$\langle n_i(n_i - 1)(n_i - 2) \rangle = \frac{\partial^3 G(1, \dots, 1)}{\partial z_i^3} \quad (12)$$

$$= N(N - 1)(N - 2)p_i^3, \quad (13)$$

which can be plugged into the definition of \hat{p}_T to demonstrate its absence of bias. For the second term, we re-express the squared coincidence probability as

$$p_C^2 = \langle \hat{p}_C^2 \rangle - \text{Var}(\hat{p}_C). \quad (14)$$

where we have used the variance decomposition formula, $\text{Var}(\hat{p}_C) = \langle \hat{p}_C^2 \rangle - \langle \hat{p}_C \rangle^2$, and the unbiasedness of Simpson's estimator, $\langle \hat{p}_C \rangle = p_C$. Plugging this expression into Eq. (8) and solving for $\text{Var}(\hat{p}_C)$ yields

$$\text{Var}(\hat{p}_C) = \frac{a}{1-b} p_T - \frac{b}{1-b} \langle \hat{p}_C^2 \rangle + \frac{c}{1-b} p_C. \quad (15)$$

Finally, the third term can be estimated using Eq. (2). Combining these results proves our central finding, the unbiasedness of the estimator proposed in Eq. (4).

III. BENCHMARKING ON SIMULATED DATA

To compare the empirical performance of the different estimators we turned to numerical experiments, applying estimators to samples from a range of ecologically relevant

species abundance distributions. We simulated drawing samples of different sizes N from the population ranging from $N = 10$ to $N = 10000$. Repeated sampling at a given sample size allows evaluation of how empirical estimates deviate from the ground truth value, $\text{Var}(\hat{p}_C)$, computed using Eq. (8) from the species abundances. Given an estimator \hat{x} of a parameter with true value x a natural measure of its quality is the mean squared error,

$$\text{MSE}(\hat{x}) = \langle (\hat{x} - x)^2 \rangle, \quad (16)$$

which can be decomposed into a bias and variance term,

$$\text{MSE}(\hat{x}) = \text{Bias}(\hat{x})^2 + \text{Var}(\hat{x}), \quad (17)$$

where $\text{Bias}(\hat{x}) = \langle \hat{x} - x \rangle$ and $\text{Var}(\hat{x}) = \langle (\hat{x} - \langle \hat{x} \rangle)^2 \rangle$. To make values more readily interpretable we normalize Bias and Variance by the true value x to the fractional values, Bias/x and Var/x^2 . We display bias [Fig. 2(b)] and variance [Fig. 2(c)] separately to investigate any potential trade-offs between bias and variance [23].

In our first experiments, we drew relative species abundances from Dirichlet distributions, $\rho(p_i) \propto p_i^{\alpha-1}$. These distributions arise in ecology as the steady state of neutral birth, death, and immigration dynamics in a Wright-Fisher diffusion limit [7,24]. We sampled a species abundance distribution of support $S = 1000$ by sampling p_i uniformly from the probability simplex, this is from a Dirichlet distribution with $\alpha = 1$ [Fig. 2(a)]. The numerical results demonstrate that the proposed estimator is not only unbiased [Fig. 2(b)] but also has lower variance than all other estimators [Fig. 2(c)]. We also generated a more uniform and more peaked distribution of species abundances, corresponding respectively to high or low relative immigration rates, by using Dirichlet parameters $\alpha = 4$ [Fig. 5(a)] and $\alpha = 0.25$ [Fig. 5(d)], respectively. We find that the unbiased estimator consistently performs best regardless of the choice of α [Figs. 5(b), 5(c), 5(e), and 5(f)].

Different sample-size regimes, which we analyze in detail in the next section, predict the performance of the different estimators. For sample sizes $N > S$ (green shading) all estimators have relatively low variance, but Chao *et al.*'s

estimator and the full plug-in estimator are substantially biased unless $N \gg S$. All estimators are highly variable for very small sample sizes $N < \sqrt{2S}$ (blue shading), which corresponds to the sample size at which the expected number of coincidences for a uniform distribution of S species is one.

To test the generality of these findings beyond neutral models, we repeated the numerical experiment with alternative species abundance distributions. In theoretical ecology, chaotic competition in generalized Lotka-Volterra models [3] and stochastic environmental fluctuations [25] have been shown to lead to the emergence of heavy-tailed species abundance distributions. Empirically, longstanding evidence shows that species abundance distributions in many complex ecosystems are well fit by lognormal [26] or power-law distributions [27]. We thus tested our estimator on samples from species abundance distributions following a lognormal form (Fig. 6) and a power-law form (Fig. 7). The results again show that other methods have substantial bias in small samples that is removed by the proposed estimator.

IV. INSIGHTS INTO SAMPLE-SIZE SCALING

To gain intuition into how estimator performance depends on N we consider limits of the variance formula [Eq. (8)]. For large N , the variance is asymptotically equal to

$$\text{Var}(\hat{p}_c) = \frac{4}{N} \left[\sum_i p_i^3 - \left(\sum_i p_i^2 \right)^2 \right]. \quad (18)$$

This shows that in large samples the variance of the estimator scales with the familiar $1/N$ scaling of an arithmetic average. Note further that the expression in brackets can be interpreted as the variance of species probabilities, $\text{Var}(p) = \langle p^2 \rangle - \langle p \rangle^2 = p_T - p_C^2$. When this variance is estimated by plugging in empirical frequencies, there is an additional sampling variance contribution, which explains why plug-in estimators are positively biased.

Conversely, when the number of species S is increased at fixed N the third term in Eq. (8) asymptotically dominates as $p_T \sim 1/S^2$ and $p_C^2 \sim 1/S^2$ while $p_C \sim 1/S$, and thus

$$\text{Var}(\hat{p}_c) = \frac{2}{N(N-1)} p_C. \quad (19)$$

Interestingly, the variance scales as $1/N^2$ in this limit, which explains the sharp rise in estimator variances as $N \sim \sqrt{2S}$. As coincidences are rare, they occur roughly independently across the $N(N-1)/2$ possible pairs and the distribution of the total number of coincidences n_C is approximately Poissonian [28] with mean $\langle n_C \rangle = N(N-1)p_C/2$. The variance of a Poisson distribution equals its mean, and as $\hat{p}_c = \frac{2n_C}{N(N-1)}$, we have $\text{Var}(\hat{p}_c) = \frac{4}{N^2(N-1)^2} \langle n_C \rangle = \frac{2}{N(N-1)} p_C$. A Poisson approximation thus recovers Eq. (19), identifying counting noise as the dominant source of variance in small samples.

The scaling analysis suggests that in terms of mean squared error it might be preferable to only estimate the Poisson term in the very smallest sample to reduce variance stemming from the estimation of the first two terms in Eq. (8). Benchmarking

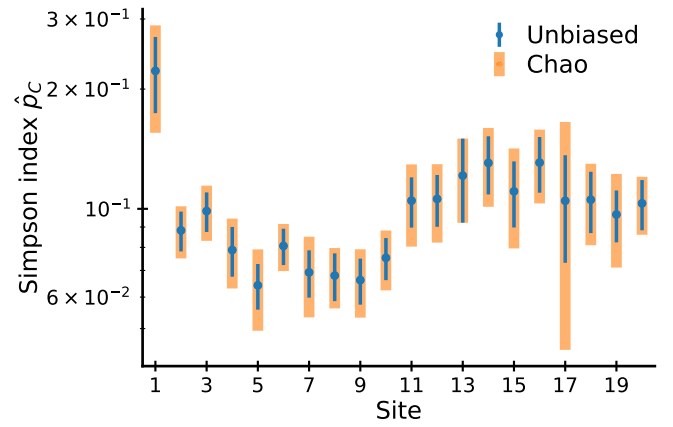


FIG. 3. Comparison of interval estimation on empirical data. The dataset presented in Jongman *et al.* [30] contains the abundances of 30 plant species at 20 different sampling sites on the Dutch island of Terschelling. Interval estimates for the biodiversity at each site were obtained from the unbiased estimator (blue) and the Chao estimator (orange) as $\hat{p}_c \pm \sqrt{\hat{\text{Var}}_{\text{method}}(\hat{p}_c)}$.

of the Poisson estimator,

$$\hat{\text{Var}}_{\text{Poisson}}(\hat{p}_c) = \frac{2}{N(N-1)} \hat{p}_c, \quad (20)$$

confirms this intuition: The Poisson estimator greatly reduces variance in small samples at the expense of moderate negative bias (Fig. 8). In practice, we propose using the maximum value of the Poisson and unbiased estimator to increase robustness in the smallest samples. Future work might more formally address the problem of combining the two estimators to minimize overall mean squared error using the statistical framework of shrinkage estimators [29].

V. COMPARISONS ON EMPIRICAL DATA

To demonstrate the practical importance of unbiased estimation we next compared estimators on empirical data. In a textbook dataset on how vegetation patterns depend on the management of dune meadows [30] we find that the unbiased estimator produces consistently lower estimates of variance than Chao *et al.*'s estimator (Fig. 3). In this dataset N varies from 15 to 48, which is of the same order of magnitude as the $S = 30$ species which are distinguished in this dataset. The wider interval estimates of Chao *et al.*'s method are compatible with the previously demonstrated bias of this method for $N \sim S$. Overestimation of sampling variance decreases the percentage of pairs of sites with nonoverlapping error bars from 44% for the unbiased method to 29% for Chao *et al.*'s method, demonstrating the potential gain in statistical power using the proposed estimator.

We next compared estimator performance for the problem of quantifying T-cell-receptor (TCR) diversity. Stochastic genetic recombination creates hypervariable TCRs, which are the molecular basis for how our adaptive immune system responds to diverse pathogens [6] [Fig. 4(a)]. TCR diversity can be quantified by considering receptors as “species” with associated probabilities corresponding to the likelihood of clonal lineages encoding the same receptor. The number S of

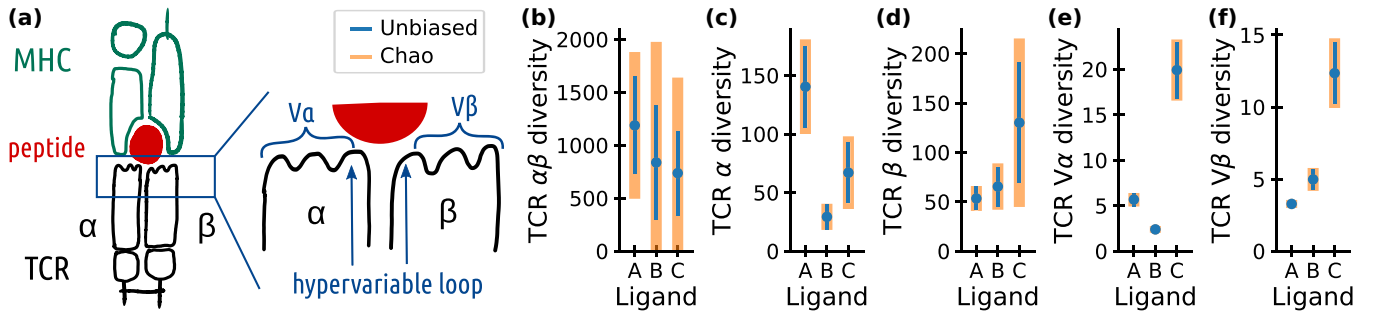


FIG. 4. Quantifying the diversity of ligand-specific T-cell receptors. (a) Schematic diagram of the interaction between TCRs and their ligands and peptides bound to major histocompatibility complex (MHC). The $\alpha\beta$ TCR is a heterodimer composed of an α and β chain, each containing variable loops that together determine TCR specificity to its ligand binding partners. The sequence of the variable loops are determined during genetic recombination which involves the choice of gene segments, called V genes, and additional diversification within the hypervariable complementary determining region 3. (b)–(f) Interval estimates for the diversity of pMHC-specific receptors were obtained from the unbiased estimator (blue) and the Chao estimator (orange) for three viral epitopes. Diversity was assessed on the level of (b) the full receptor, (c) the α chain, and (d) β chain, as well as (e) the $V\alpha$ gene choice and (f) the $V\beta$ gene choice. Diversities are shown as effective number equivalents of Simpson’s index, $D = 1/\hat{p}_C$, with error bars calculated by error propagation $\sqrt{\hat{V}\text{ar}(\hat{p}_C)}/\hat{p}_C^2$. Data: Dash *et al.* [34]. Ligand A: Influenza virus peptide M1₅₈. Ligand B: Epstein-Barr virus peptide BMLF1₂₈₀. Ligand C: Human cytomegalovirus peptide pp65₄₉₅. The three peptides are presented by a common MHC, the human leukocyte antigen (HLA) A*02:01.

receptors that can be created by recombination is immense with estimates as large as $S \sim 10^{39}$ for the TCR β chain [31]. Therefore, this problem illustrates a practical use case for diversity estimation methods in the $N \ll S$ regime.

To test the variance estimators we constructed a metarepertoire of 3×10^7 T-cell clonotypes by combining samples from multiple healthy donors from a cohort study [32] and then split this repertoire into nonoverlapping pools of different sizes. The unbiased estimator outperforms all other estimators in terms of bias [Fig. 9(b)] and variance [Fig. 9(c)]. (Chao *et al.*’s estimator was excluded from this comparison due to its slow computational speed at tested sample sizes.) Our results show that using the unbiased estimator only $> 10\,000$ sequences are needed to estimate $\text{Var}(\hat{p}_C)$ with a coefficient of variation $\ll 1$. Such sampling depths are now readily obtainable via bulk [32] or single-cell TCR sequencing [33], allowing this estimator to be applied to quantify differences in TCR diversity across individual samples.

VI. APPLICATION: HOW MANY T-CELL RECEPTORS BIND TO A GIVEN LIGAND?

Having established the good empirical performance of the unbiased estimator, we sought to exploit our statistical advance to provide a quantitative answer to an important open question in immunology (Fig. 4): How degenerate is the mapping between antigen receptors and their ligands? The diversity of antigen receptors binding to a ligand determines the breadth and polyclonality of the adaptive immune response and quantification of this diversity is thus of central interest in the field [34–36]. Recent experimental advances in single-cell sequencing of T cells sorted for specificity to multimerized ligands allow experimental probing of this diversity [34,37,38]. By quantifying the diversity of ligand-specific TCRs at multiple levels, we demonstrate the statistical significance of ligand-dependent differences in the contribution

of the two chains of the heterodimeric receptor to binding specificity.

The dataset we consider consists of the sequences of 415 $\alpha\beta$ TCRs specific to three ligands [34], which are important viral epitopes (see Fig. 4 caption). We label them A, B, and C in the text for conciseness. Each experiment involved sorting T cells from multiple individuals, but for simplicity we determined overall TCR diversity regardless of donor origin, a limitation which can be relaxed as dataset sizes increase. Using our method we determined the diversity of the full receptor [Fig. 4(b)] and its parts [Figs. 4(c)–4(f)] along with their associated sampling uncertainty. We find that the effective diversity of TCRs binding each ligand is on the order of 1000 receptors [Fig. 4(b)] albeit with a large associated sampling uncertainty. Interestingly, when zooming in on the diversity of the component parts of the heterodimeric receptor, we find strong statistical support for hypothesized differences [34,38] in α - and β -chain diversity among ligand-specific receptors [Figs. 4(c)–4(f)]. For instance, TCRs specific to ligand A have significantly more diverse α chains than β chains, while the reverse is true for ligand B. The diversity of V gene segments found in specific TCRs also varies significantly between ligands and is largest among TCRs specific to ligand C.

The unbiased interval estimates (blue) are equivalent to Chao *et al.*’s estimates (orange) for V-gene diversity [Figs. 4(e) and 4(f)], but are substantially tighter for the diversity estimates on the full receptor level [Fig. 4(b)] and receptor chain level [Figs. 4(c) and 4(d)], again highlighting the upward bias of alternative estimators. Importantly, the quantification of TCR diversity at different levels leads to hypotheses about the structural basis of recognition for each ligand. For example, TCRs specific to ligand C are expected to make fewer contacts on average between the V-gene encoded CDR1 and CDR2 loops and the ligand. Similarly, diversity restriction among the two chains in ligands A and B might be reflective of how many contacts each chain is making with

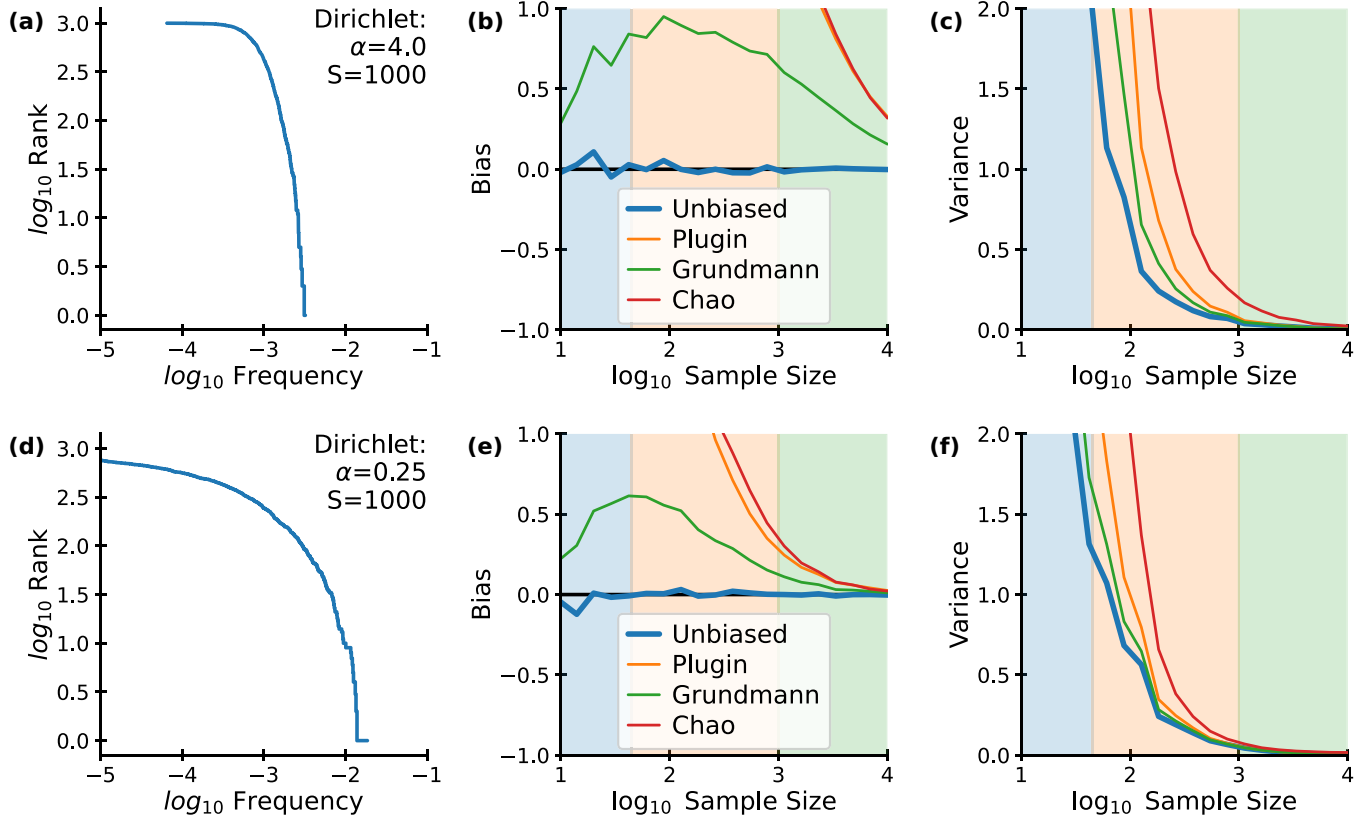


FIG. 5. Performance of variance estimators for species abundances drawn from a Dirichlet distribution with other choices of the parameter α . Top panels: $\alpha = 4.0$. Bottom panels: $\alpha = 0.25$.

the ligand. These hypotheses will soon become testable as more structures of TCRs in complex with their cognate ligands are solved [35,39].

VII. CONCLUSION AND DISCUSSION

This work introduced a method to estimate the sample variance of Simpson's diversity index without bias for arbitrary species abundance distributions. This unbiased estimator does not seem to be widely known despite its superior statistical properties compared to existing methods. Additionally the unbiased estimator has a closed-form analytical expression and is thus fast to calculate even for large samples, in contrast to bootstrapping approaches.

As the estimator is unbiased it has the practically important property of producing estimates that do not vary systematically with sample size. This is an important advantage for practical applications in which sample sizes vary between ecological communities. Such variation has become an increasingly important concern in ecology, as the field has moved to apply techniques developed for field studies with well-controlled sampling effort to the assessment of microbiome [8,10] or immune repertoire [16,36,40] diversity from high-throughput sequencing experiments. By using the unbiased estimator introduced here ecologists can avoid the loss of information inherent in the common practice of subsampling larger samples down to the smallest sample size, known as rarefaction. Our method thus fills a previously identified gap in

the ecological literature to overcome the need for rarefaction by bias-corrected interval estimators [10].

An extension of our work could revisit methods for interval estimation for other diversity metrics such as Shannon entropy. For these metrics past work has focused on reducing bias in the point estimates themselves given the absence of an unbiased estimator [13,21]. Our work might be generalized to address the variance estimation problem for these bias-corrected estimators for other diversity metrics. Another direction for future work is to compare the performance of the estimators on samples with overdispersion [41], which goes beyond the multinomial sampling assumption that underlies all tested estimators.

We note that the negative logarithm of Simpson's index, $-\log p_C$, is the Renyi entropy (of order 2) [9,42]. The Renyi entropy in turn lower-bounds the Shannon entropy $-\sum_i p_i \log p_i$, a relation that has been exploited to estimate entropy rates of dynamical systems [13,43] and neural spike trains [44]. Thus we expect that our estimator will also be of use outside of ecology in the many other areas that use the concept of entropy. Interestingly, the estimator we have introduced can determine sampling variances even when the total number of species S exceeds the sample size N . This shows that the surprising ability to infer entropies way before the distribution is fully sampled, known in statistical physics as Ma's square-root regime of entropy estimation [13,43,45] and in probability theory as the birthday paradox [46,47], generalizes from point to interval estimation.




FIG. 6. Performance of variance estimators for lognormally distributed species abundances with $\log p_i \propto N(0, \sigma)$ with $\sigma = 1$.

Application of our estimator experimentally identified ligand-specific T-cell receptors, showed that their effective diversity is on the order of ~ 1000 receptors, and demonstrated ligand-dependent restriction of TCR chain diversity. The effective number of receptors is very small compared to the multiple trillions of $\alpha\beta$ receptors that can be produced by recombination, demonstrating the stringent selection of antigen-specific TCRs in these experiments [36]. Knowing how many TCRs on average bind a given ligand is important in experimental design for TCR screens as it can help guide the breadth and depth of sampling strategies. Quantification of variability in the effective number of TCRs binding to different ligands using the method introduced in this paper could yield insights into the mechanistic basis of immunodominance hierarchies and help quantify how much the effective diversity of specific TCRs depends on cutoffs on TCR avidity imposed by different experimental assays, prior exposure, or age [35,38,48]. Estimates of the diversity of different features of the TCR relative to null expectations can furthermore bound which features are most informative about receptor specificity [49]. Taken together, quantification of ligand-specific TCR diversity might help explain variation in performance of machine learning models for different ligands [50,51].

We hope that by improving both the trustworthiness and practicality of interval estimates for biodiversity measurements our work will lead to a wider adoption of sampling uncertainty quantification in ecological research. To aid

adoption of the method we have made a reference implementation of the estimator available as an open source PYTHON package [52]. We are particularly hopeful that our method for interval estimation of diversity will enable focusing of sampling efforts for monitoring biodiversity loss in our changing world.

ACKNOWLEDGMENTS

The author thanks Curtis G. Callan and Yuta Nagano for critical reading of the manuscript. The work was supported in part by the Royal Free Charity and was concluded at the Aspen Center for Physics, which is supported by the National Science Foundation under Grant No. PHY-2210452.

APPENDIX A: DERIVATION OF THE VARIANCE EXPRESSION

To derive the variance of Simpson's estimator we need to calculate various (cross-)moments of the multinomial distribution. Taking derivatives of the probability generating function demonstrates that the factorial moments are equal to

$$\left\langle \prod_i^S n_i^{(a_i)} \right\rangle = N^{(\sum_i a_i)} \prod_i^S p_i^{a_i}, \quad (\text{A1})$$

where $x^{(n)} = x(x-1)\dots(x-n+1)$ denotes the falling factorial. To calculate the raw moments of the distribution we




FIG. 7. Performance of variance estimators for Zipf-distributed species abundances with $p_i \sim 1/i$.




FIG. 8. Comparison of the Poisson and unbiased estimators. (a) Frequency-rank plot as in Fig. 2. (b) Bias and (c) variance as a function of sample size N for the unbiased estimator [Eq. (4)] and the Poisson estimator [Eq. (20)]. While the Poisson estimator is substantially negatively biased in large samples, it has only modest bias in small samples and variance lower than that of the unbiased estimator.

make use of the moment generating function,

$$M(t_1, \dots, t_S) = G(e^{t_1}, \dots, e^{t_S}) = \langle e^{\sum_i t_i n_i} \rangle, \quad (\text{A2})$$

which for the multinomial distribution is equal to

$$M(t_1, \dots, t_S) = \left(\sum_i p_i e^{t_i} \right)^N. \quad (\text{A3})$$

Calculating the partial derivatives of the moment generating function at $t_1 = \dots = t_S = 0$, we obtain

$$\langle n_i \rangle = N p_i, \quad (\text{A4})$$

$$\langle n_i^2 \rangle = N^{(2)} p_i^2 + N p_i, \quad (\text{A5})$$

$$\langle n_i^3 \rangle = N^{(3)} p_i^3 + 3N^{(2)} p_i^2 + N p_i, \quad (\text{A6})$$

$$\langle n_i^4 \rangle = N^{(4)} p_i^4 + 6N^{(3)} p_i^3 + 7N^{(2)} p_i^2 + N p_i. \quad (\text{A7})$$

The variance of \hat{p}_c can be expressed as

$$\text{Var}(\hat{p}_c) = \frac{\langle [\sum_i n_i (n_i - 1)]^2 \rangle}{[N(N-1)]^2} - p_c^2. \quad (\text{A8})$$

The key calculation concerns the numerator of the first term, which is equal to

$$\sum_i \langle (n_i^{(2)})^2 \rangle + \sum_i \sum_{j \neq i} \langle n_i^{(2)} n_j^{(2)} \rangle. \quad (\text{A9})$$

Expanding the first term and evaluating the second average using Eq. (A1) yields

$$\sum_i (\langle n_i^4 \rangle - 2\langle n_i^3 \rangle + \langle n_i^2 \rangle) + \sum_i \sum_{j \neq i} N^{(4)} p_i^2 p_j^2. \quad (\text{A10})$$

Using the expression for the moments Eqs. (A5)–(A7), and noting that $\sum_{j \neq i} p_j^2 = \sum_j p_j^2 - p_i^2 = p_c - p_i^2$, we obtain

$$4N^{(3)} \sum_i p_i^3 + 2N^{(2)} p_c + N^{(4)} p_c^2. \quad (\text{A11})$$




FIG. 9. Benchmarking of variance estimators on T-cell-receptor repertoire data. (a) Frequency-rank plot of TCR multiplicities among 3×10^7 clonotypes from a metarepertoire. (b) Bias and (c) variance as a function of sample size N . Expectation values were calculated at each sample size by splitting the total sequence pool into nonoverlapping subsets. Bias and variance are expressed as fractions, i.e., divided by the true value or its square, respectively. The shaded areas differentiate sampling regimes: blue $N < \sqrt{2D}$, orange $\sqrt{2D} < N$ with effective diversity $D = 1/\hat{p}_c \sim 1.5 \times 10^6$ [estimated by applying Eq. (2) to the complete dataset]. Note that in panel (b) a symlog scale is used to account for the large bias of the plug-in method, i.e., bias values are transformed by $\text{sign}(x) \log_{10}(|x| + 1)$. Data: A metarepertoire of 3×10^7 complementary determining region 3 (CDR3) sequences of the TCR β -chain constructed by random selection from the combined productive clonotypes of 200 healthy human donors from the Emerson *et al.* [32] study.

Plugging this numerator into Eq. (A8), we obtain after some algebra

$$\text{Var}(\hat{p}_C) = \frac{4N^{(3)} \sum_i p_i^3 - 2N^{(2)}(2N - 3)p_C^2 + 2N^{(2)}p_C}{[N(N - 1)]^2}, \quad (\text{A12})$$

the expression first published by Simpson [12].

APPENDIX B: SUPPLEMENTAL FIGURES

Figures 5–7 are supplemental to Fig. 2. Figure 8 shows a comparison of the Poisson and unbiased estimators, and Fig. 9 shows benchmarking of the variance estimators on bulk TCR repertoire data.

-
- [1] S. Allesina and S. Tang, *Nature (London)* **483**, 205 (2012).
- [2] A. Posfai, T. Taillefumier, and N. S. Wingreen, *Phys. Rev. Lett.* **118**, 028103 (2017).
- [3] M. T. Pearce, A. Agarwala, and D. S. Fisher, *Proc. Natl. Acad. Sci. USA* **117**, 14572 (2020).
- [4] Q. Marcou, T. Mora, and A. M. Walczak, *Nat. Commun.* **9**, 561 (2018).
- [5] Z. Sethna, Y. Elhanati, C. G. Callan, A. Walczak, and T. Mora, *Bioinformatics* **35**, 2974 (2019).
- [6] A. Mayer, V. Balasubramanian, T. Mora, and A. M. Walczak, *Proc. Natl. Acad. Sci. USA* **112**, 5950 (2015).
- [7] A. Mayer, V. Balasubramanian, A. M. Walczak, and T. Mora, *Proc. Natl. Acad. Sci. USA* **116**, 8815 (2019).
- [8] B. Haegeman, J. Hamelin, J. Moriarty, P. Neal, J. Dushoff, and J. S. Weitz, *ISME J.* **7**, 1092 (2013).
- [9] M. Roswell, J. Dushoff, and R. Winfree, *Oikos* **130**, 321 (2021).
- [10] A. D. Willis, *Front. Microbiology* **10**, 2407 (2019).
- [11] P. G. Albano, J. Steger, M. Bošnjak, B. Dunne, Z. Guifarro, E. Turapova, Q. Hua, D. S. Kaufman, G. Rilov, and M. Zuschin, *Proc. R. Soc. B* **288**, 20202469 (2021).
- [12] E. H. Simpson, *Nature (London)* **163**, 688 (1949).
- [13] I. Nemenman, F. Shafee, and W. Bialek, *Proceedings of the 14th International Conference on Neural Information Processing Systems (NIPS'01)*, edited by T. Dietterich, S. Becker, and Z. Ghahramani (MIT Press, 2001).
- [14] H. Grundmann, S. Hori, and G. Tanner, *J. Clin. Microbiol.* **39**, 4190 (2001).
- [15] A. Chao, N. J. Gotelli, T. C. Hsieh, E. L. Sander, K. H. Ma, R. K. Colwell, and A. M. Ellison, *Ecol. Monogr.* **84**, 45 (2014).
- [16] J. Kaplinsky and R. Arnaout, *Nat. Commun.* **7**, 11881 (2016).
- [17] M. Grabchak, E. Marcon, G. Lang, and Z. Zhang, *PLoS One* **12**, e0173305 (2017).
- [18] L. Jost, *Oikos* **113**, 363 (2006).
- [19] P. R. Hunter and M. A. Gaston, *J. Clin. Microbiol.* **26**, 2465 (1988).
- [20] J. P. Gibbs and W. T. Martin, *Am. Soc. Rev.* **27**, 667 (1962).
- [21] P. Grassberger, [arXiv:physics/0307138v2](https://arxiv.org/abs/physics/0307138v2).
- [22] T. C. Hsieh, K. H. Ma, and A. Chao, *Methods Ecol. Evol.* **7**, 1451 (2016).
- [23] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, *Phys. Rep.* **810**, 1 (2019).
- [24] A. Etheridge, *Some Mathematical Models from Population Genetics* (Springer-Verlag, Berlin, 2012).
- [25] M. U. Gaimann, M. Nguyen, J. Desponds, and A. Mayer, *eLife* **9**, e61639 (2020).
- [26] F. W. Preston, *Ecology* **29**, 254 (1948).
- [27] G. U. Yule, *Philos. Trans. R. Soc. B* **213**, 21 (1924).
- [28] D. Aldous, *Probability Approximations via the Poisson Clumping Heuristic*, Applied Mathematical Sciences Vol. 77 (Springer, New York, 1989).
- [29] J. Hausser and K. Strimmer, *J. Mach. Learn. Res.* **10**, 1469 (2009).
- [30] R. H. Jongman, C. J. F. ter Braak, and O. F. R. Van Tongeren, Editors, *Data Analysis in Community and Landscape Ecology* (Cambridge University, Cambridge, England, 1995).
- [31] T. Mora and A. Walczak, in *Systems Immunology*, edited by J. Das and C. Jayaprakash (CRC Press, Boca Raton, FL, 2019), Chap. 11.
- [32] R. O. Emerson, W. S. DeWitt, M. Vignali, J. Gravley, J. K. Hu, E. J. Osborne, C. Desmarais, M. Klinger, C. S. Carlson, J. A. Hansen *et al.*, *Nat. Genet.* **49**, 659 (2017).
- [33] R. G. H. Lindeboom, K. B. Worlock, L. M. Dratva, M. Yoshida, D. Scobie, H. R. Wagstaffe, L. Richardson, A. Wilbrey-Clark, J. L. Barnes, K. Polanski *et al.*, *Nature* (2024), doi:10.1038/s41586-024-07575-x.
- [34] P. Dash, A. J. Fiore-Gartland, T. Hertz, G. C. Wang, S. Sharma, A. Souquette, J. C. Crawford, E. B. Clemens, T. H. Nguyen, K. Kedzierska *et al.*, *Nature (London)* **547**, 89 (2017).
- [35] G. Chen, X. Yang, A. Ko, X. Sun, M. Gao, Y. Zhang, A. Shi, R. A. Mariuzza, and N.-p. Weng, *Cell Rep.* **19**, 569 (2017).
- [36] A. Mayer and C. G. Callan, Jr., *Proc. Natl. Acad. Sci. USA* **120**, e2213264120 (2023).
- [37] J. Glanville, H. Huang, A. Nau, O. Hatton, L. E. Wagar, F. Rubelt, X. Ji, A. Han, S. M. Krams, C. Pettus *et al.*, *Nature (London)* **547**, 94 (2017).
- [38] A. A. Minervina, M. V. Pogorelyy, A. M. Kirk, J. C. Crawford, E. K. Allen, C. H. Chou, R. C. Mettelman, K. J. Allison, C. Y. Lin, D. C. Brice *et al.*, *Nat. Immunol.* **23**, 781 (2022).
- [39] I. Song, A. Gil, R. Mishra, D. Ghersi, L. K. Selin, and L. J. Stern, *Nat. Struct. Mol. Biol.* **24**, 395 (2017).
- [40] D. J. Laydon, C. R. Bangham, and B. Asquith, *Philos. Trans. R. Soc. B* **370**, 20140291 (2015).
- [41] R. Scherer, F. Schaarschmidt, S. Prescher, and K. U. Priesnitz, *Biom. J.* **55**, 246 (2013).
- [42] S. Xu, L. Böttcher, and T. Chou, *Phys. Biol.* **17**, 031001 (2020).
- [43] S.-k. Ma, *J. Stat. Phys.* **26**, 221 (1981).
- [44] S. P. Strong, R. Koberle, R. de R. van Steveninck, and W. Bialek, *Phys. Rev. Lett.* **80**, 197 (1998).
- [45] D. G. Hernández, A. Roman, and I. Nemenman, *Phys. Rev. E* **108**, 014101 (2023).
- [46] T. S. Nunnikhoven, *Am. Stat.* **46**, 270 (1992).
- [47] L. Paninski, *IEEE Trans. Inf. Theory* **54**, 4750 (2008).

- [48] C. E. van de Sandt, T. H. O. Nguyen, N. A. Gherardin, J. C. Crawford, J. Samir, A. A. Minervina, M. V. Pogorelyy, S. Rizzetto, C. Szeto, J. Kaur *et al.*, *Nat. Immunol.* **24**, 1890 (2023).
- [49] J. Henderson, Y. Nagano, M. Milighetti, and A. Tiffeau-Mayer, [arXiv:2404.12565](https://arxiv.org/abs/2404.12565).
- [50] A. Montemurro, L. E. Jessen, and M. Nielsen, *Front. Immunol.* **13**, 1055151 (2022).
- [51] G. Croce, S. Bobisse, D. L. Moreno, J. Schmidt, P. Guillame, A. Harari, and D. Gfeller, *Nat. Commun.* **15**, 3211 (2024).
- [52] <https://github.com/andim/pydiver>.