

Fundamental limits on the accuracy of demographic inference based on the sample frequency spectrum

Jonathan Terhorst^a and Yun S. Song^{a,b,c,1}

^aDepartment of Statistics, University of California, Berkeley, CA 94720; ^bComputer Science Division, University of California, Berkeley, CA 94720; and ^cDepartment of Integrative Biology, University of California, Berkeley, CA 94720

Edited by Michael S. Waterman, University of Southern California, Los Angeles, CA, and approved May 18, 2015 (received for review February 23, 2015)

The sample frequency spectrum (SFS) of DNA sequences from a collection of individuals is a summary statistic that is commonly used for parametric inference in population genetics. Despite the popularity of SFS-based inference methods, little is currently known about the information theoretic limit on the estimation accuracy as a function of sample size. Here, we show that using the SFS to estimate the size history of a population has a minimax error of at least $O(1/\log s)$, where s is the number of independent segregating sites used in the analysis. This rate is exponentially worse than known convergence rates for many classical estimation problems in statistics. Another surprising aspect of our theoretical bound is that it does not depend on the dimension of the SFS, which is related to the number of sampled individuals. This means that, for a fixed number s of segregating sites considered, using more individuals does not help to reduce the minimax error bound. Our result pertains to populations that have experienced a bottleneck, and we argue that it can be expected to apply to many populations in nature.

minimax rate | population genetics | demographic inference

The past decade has seen a revolution in our ability to interrogate the genome at the molecular level. Fueled by technological advances in DNA sequencing, studies now routinely query thousands or tens of thousands of individuals [refs. 1–4 and UK10K Project (www.uk10k.org) and Exome Aggregation Consortium (exac.broadinstitute.org)] to better understand disease susceptibility, heritability, population history, and other phenomena. In most cases, the conclusions of these studies come in the form of statistical estimates obtained from models that relate the effect of interest to mutation patterns arising in sampled DNA sequences. As genetic sample sizes explode, it is natural to wonder how additional data improve the quality of these estimates. While this general question has received intense focus in theoretical statistics, certain aspects of the genetics setting (for example, non-Gaussianity and lack of independence among samples) complicate efforts to study such models using classical techniques. New methods are needed to theoretically characterize some common models in statistical genetics.

Here, we address this need for a specific estimation problem in population genetics known as demographic inference. As we explain in further detail below, the aim of this problem is to reconstruct the sequence of historical events—including population size changes, migration, and admixture—that gave rise to present-day populations, using DNA samples obtained from those populations. We focus on the simplest problem of estimating the size history of a single population backward in time.

A summary statistic known as the (SFS; defined below) is often used in empirical studies (2, 5–11), but there have been fewer attempts to understand SFS-based estimation from a theoretical perspective. The main result of this paper is to show that, for a common class of estimators that analyze the SFS, there is a fundamental limit on their accuracy as a function of the sample size. More precisely, we show that, under a standard statistical error metric known as minimax error, the rate at which these estimators converge to the truth for certain populations is at best inversely logarithmic in the number of independent segregating sites analyzed, and does not depend at all on the number of individuals sampled. Compared with other types of statistical estimation

problems (for example, linear regression), this is an extremely slow rate of convergence. Our proof is information theoretic in nature and applies to any estimator that operates solely on the SFS. This is the first result we are aware of that characterizes the convergence rate of demographic history estimates as a function of sample size.

The remainder of this paper is organized as follows. In *Preliminaries*, we formally define our notation and model. In *Main Results*, we state our main theoretical results, followed by a discussion of their practical implications in *Discussion*. To streamline our exposition, all mathematical proofs are deferred until *Proofs*.

Preliminaries

The stochastic process underlying the inference procedure we consider is Kingman's coalescent (12–14), which evolves backward in time and describes the genealogy of a collection of chromosomes randomly sampled from a population. The population size is assumed to change deterministically over time and is described by a function $\eta: [0, \infty) \rightarrow (0, \infty)$, with $\eta(t)$ being the population size at time t in the past. The instantaneous rate of coalescence between any pair of lineages at time t is $1/\eta(t)$.

As in the standard infinite sites model of mutation (15), we assume that every dimorphic site (i.e., a site with exactly two observed allelic types) has experienced mutation exactly once in the evolutionary history of the sample. Further, for each such site, we assume that it is known which allele is the ancestral type versus the mutant type. In what follows, we use the terms “dimorphic” and “segregating” interchangeably.

A population size function $\eta(t)$ induces a probability distribution on the number of derived alleles found at a particular segregating site. Specifically, for a sample of $n \geq 2$ randomly sampled individuals, let $\xi_{n,b}^{(\eta)}$, for $1 \leq b \leq n-1$, denote the probability that a segregating site contains b mutant alleles in a sample of n individuals under model η . The vector $\xi_n^{(\eta)} \stackrel{\text{def}}{=} (\xi_{n,1}^{(\eta)}, \dots, \xi_{n,n-1}^{(\eta)})$

Significance

Numerous empirical studies in population genetics have used a summary statistic called the sample frequency spectrum (SFS), which summarizes the information in a sample of DNA sequences. Despite their popularity, the accuracy of inference methods based on the SFS is difficult to characterize theoretically, and it is currently unknown how the estimation accuracy improves as more sites in the genome are used. Here, we establish information theoretic limits on the accuracy of all estimators that use the SFS to infer population size histories. We study the rate of convergence to the true answer as the amount of data increases, and obtain the surprising result that it is exponentially worse than known convergence rates for many classical estimation problems in statistics.

Author contributions: J.T. and Y.S.S. designed research, performed research, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. Email: yss@berkeley.edu.

is called the expected SFS. In the coalescent setting, a general expression for $\xi_{n,b}^{(\eta)}$ is given by (16)

$$\xi_{n,b}^{(\eta)} \propto \sum_{k=2}^{n-b+1} \frac{\binom{n-b-1}{k-2}}{\binom{n-1}{k-1}} \cdot k \cdot \mathbb{E}T_{n,k}^{(\eta)},$$

where $\mathbb{E}T_{n,k}^{(\eta)}$ denotes the amount of time (in coalescent units) during which the genealogy of the sample contained k lineages under model η . The expected waiting time $\mathbb{E}T_{m,m}^{(\eta)}$ to the first coalescence in a sample of m individuals is given by

$$c_m^{(\eta)} \stackrel{\text{def}}{=} \mathbb{E}T_{m,m}^{(\eta)} = \int_0^\infty t \frac{a_m}{\eta(t)} \exp\{-a_m R_\eta(t)\} dt, \quad [1]$$

where $a_m \stackrel{\text{def}}{=} \binom{m}{2}$ and $R_\eta(t) \stackrel{\text{def}}{=} \int_0^t \frac{1}{\eta(s)} ds$ is the cumulative rate of coalescence up to time t . It turns out (17) that there is an invertible linear transformation that relates $(\mathbb{E}T_{n,2}^{(\eta)}, \mathbb{E}T_{n,3}^{(\eta)}, \dots, \mathbb{E}T_{n,n}^{(\eta)})$ to $\mathbf{c}^{(\eta)} \stackrel{\text{def}}{=} (c_2^{(\eta)}, c_3^{(\eta)}, \dots, c_n^{(\eta)})$. Using this relation, the quantity $\xi_{n,b}^{(\eta)}$ can be written as (18)

$$\xi_{n,b}^{(\eta)} = \frac{\langle \mathbf{c}^{(\eta)}, \mathbf{W}_{n,b} \rangle}{\langle \mathbf{c}^{(\eta)}, \mathbf{V}_n \rangle}, \quad [2]$$

where $\mathbf{W}_{n,b} = (W_{n,b,2}, \dots, W_{n,b,n})$ and $\mathbf{V}_n = (V_{n,2}, \dots, V_{n,n})$ are vectors of universal constants that do not depend on the population size function η , and $\langle \cdot, \cdot \rangle$ denotes the l_2 inner product. Under model η , the quantity $\langle \mathbf{c}^{(\eta)}, \mathbf{W}_{n,b} \rangle$ is the total expected length of edges subtending b out of n individuals sampled at time 0, while the quantity $\langle \mathbf{c}^{(\eta)}, \mathbf{V}_n \rangle$ is the total expected tree length for a sample of size n . Both quantities are positive for all population size functions η . For an arbitrary population size function η , we have $\sum_{b=1}^{n-1} W_{n,b,m} = V_{n,m}$ for all $2 \leq m \leq n$, which implies

$$\sum_{b=1}^{n-1} \langle \mathbf{c}^{(\eta)}, \mathbf{W}_{n,b} \rangle = \langle \mathbf{c}^{(\eta)}, \mathbf{V}_n \rangle. \quad [3]$$

For a constant function $\eta(t) \equiv N$,

$$c_m^{(\eta)} = \frac{N}{a_m},$$

$$\langle \mathbf{c}^{(\eta)}, \mathbf{W}_{n,b} \rangle = \frac{2}{b} N, \quad [4]$$

$$\langle \mathbf{c}^{(\eta)}, \mathbf{V}_n \rangle = 2N H_{n-1}, \quad [5]$$

where $H_{n-1} \stackrel{\text{def}}{=} \sum_{b=1}^{n-1} \frac{1}{b}$.

To formulate the problem, we use the following notation. We suppose that a sample of $n \geq 2$ randomly sampled individuals has been typed at s independent segregating sites. These data are used to form the empirical sample frequency spectrum, which is an $(n-1)$ -tuple $(\hat{\xi}_{n,1}, \dots, \hat{\xi}_{n,n-1})$, where $\hat{\xi}_{n,b}$ denotes the proportion of segregating sites with b copies of the mutant allele and $n-b$ copies of the ancestral allele. A frequency-based estimator is any statistic $\hat{\eta}$ that maps an empirical SFS to a population size history.

Main Results

Here, we establish a minimax lower bound on the ability of any estimator $\hat{\eta}$ to accurately reconstruct population size functions.

A General Bound on the Kullback–Leibler Divergence Between Two SFS Distributions. Abusing notation, we use $D(\eta||\eta')$ to denote the Kullback–Leibler (KL) divergence between the probability distributions $\xi_n^{(\eta)}$ and $\xi_n^{(\eta')}$. In *Proofs*, we prove the following general upper bound on the KL divergence between two SFS distributions:

Theorem 1. Let \mathcal{M} denote a general space of population size functions and suppose $\eta, \eta' \in \mathcal{M}$ satisfy $\eta(t) = \eta'(t)$ for all $0 \leq t \leq t_c$ and $\max_{t>t_c} \eta(t) \leq \min_{t>t_c} \eta'(t)$. Then,

$$D(\eta||\eta') \leq \frac{\langle \mathbf{c}^{(\eta')} - \mathbf{c}^{(\eta)}, \mathbf{V}_n \rangle}{\langle \mathbf{c}^{(\eta)}, \mathbf{V}_n \rangle}. \quad [6]$$

Bounds for a Family of Piecewise Constant Models. We now focus on a particular class of population size functions that are easier to analyze and are popular in the literature (11, 19, 20). For a fixed positive integer $K > 1$, let $\mathcal{M}_K \subset \mathcal{M}$ denote the space of piecewise constant size functions with exactly K pieces. A population size function η is a member of \mathcal{M}_K if and only if there exist positive real numbers $t_1 < \dots < t_{K-1}$ and N_1, N_2, \dots, N_K such that

$$\eta(t) = \sum_{k=1}^K N_k \mathbf{1}\{t_{k-1} \leq t < t_k\}, \quad [7]$$

where, by convention, we define $t_0 = 0$ and $t_K = \infty$. For such an η , define

$$S_k^{(\eta)} \stackrel{\text{def}}{=} \sum_{j=1}^k \frac{t_j - t_{j-1}}{N_j}. \quad [8]$$

For $\eta \in \mathcal{M}_K$, the expected waiting time $c_m^{(\eta)}$ defined in Eq. 1 is given by

$$c_m^{(\eta)} = \frac{1}{a_m} \sum_{k=1}^K N_k \left(e^{-a_m S_{k-1}^{(\eta)}} - e^{-a_m S_k^{(\eta)}} \right). \quad [9]$$

Note that since $t_K = \infty$,

$$e^{-a_m S_K^{(\eta)}} \equiv 0, \text{ for all } \eta \in \mathcal{M}_K. \quad [10]$$

To formulate our result, we let I, J denote positive integers that satisfy $I+J=K$, and introduce a subfamily $\mathcal{F}_{IJ} \subset \mathcal{M}_K$ of piecewise constant functions defined as follows. See Fig. 1 for illustration. We assume that all change points $t_1 < \dots < t_{I+J-1}$ are fixed and that the sizes N_1, \dots, N_I of the first I epochs are also fixed, with N_I being the smallest size. So, all functions in \mathcal{F}_{IJ} are identical to each other for the first I epochs, and there is a population bottleneck in the last epoch. Then, for $t \geq t_I$, every function $\eta \in \mathcal{F}_{IJ}$ undergoes jumps according to the following rules:

1. For the interval $t_I \leq t < t_{I+1}$, $\eta(t)$ takes a constant value of either h or $h + \delta$, where $h > N_I$ and $\delta > 0$.
2. At later change points $\{t_{I+1}, \dots, t_{I+J-1}\}$, η either stays the same or jumps upward by δ .

Hence, \mathcal{F}_{IJ} consists of 2^J distinct piecewise constant functions that are nondecreasing functions of t for $t \geq t_I$. Note that $\min_t \eta(t) = N_I$ for all $\eta \in \mathcal{F}_{IJ}$. For ease of notation, we use $\varepsilon \stackrel{\text{def}}{=} N_I$ to denote the bottleneck size and $\tau_B \stackrel{\text{def}}{=} t_I - t_{I-1}$ to denote the bottleneck duration. To facilitate analysis later, we fix $t_{I+j} - t_{I+j-1}$ to some positive constant τ_A for all $j = 1, \dots, J-1$.

For any two models in \mathcal{F}_{IJ} , we obtain the following bound on the difference of their waiting times to the first coalescence:

Lemma 2. For all $\eta, \eta' \in \mathcal{F}_{I,J}$,

$$\left| c_m^{(\eta)} - c_m^{(\eta')} \right| \leq J \frac{\delta}{a_m} e^{-a_m \tau_B / \varepsilon}. \quad [11]$$

Together with Theorem 1, this lemma can be used to show

Theorem 3. Let $\eta, \eta' \in \mathcal{F}_{I,J}$ that satisfy $\max_{t \geq t_1} \eta(t) \leq \min_{t \geq t_1} \eta'(t)$. Then,

$$D(\eta || \eta') \leq J \frac{\delta}{\varepsilon} e^{-\tau_B / \varepsilon}. \quad [12]$$

Proofs of these results are deferred to *Proofs*. It is interesting that the above bound does not depend on the number n of sampled individuals.

Minimax Lower Bounds. Before using the above results to obtain a minimax lower bound, we first note a subtle fact. Given any population size function η , consider a function ζ that satisfies $\zeta(t) = \kappa \cdot \eta(t/\kappa)$ for all $t \in [0, \infty)$, where κ is some positive constant. Such functions are equivalent, as it turns out that $\xi_{n,b}^{(\zeta)} = \xi_{n,b}^{(\eta)}$ for all $n \geq 2$ and $1 \leq b \leq n-1$. To mod out by this equivalence, we assume that every $\eta \in \mathcal{M}$ satisfies $\eta(0) = N_{\text{fix}}$, where N_{fix} is some fixed positive constant.

Let $\|\cdot\|_*$ denote a generic norm (specific examples will be given later) and let $\mathbb{E}_\eta(\cdot)$ denote expectation with respect to the SFS distribution $\xi_n^{(\eta)} = (\xi_{n,1}^{(\eta)}, \dots, \xi_{n,n-1}^{(\eta)})$ induced by population size function η . Then, note that

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{M}} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_* \geq \inf_{\hat{\eta}} \sup_{\eta \in \mathcal{M}_K} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_* \geq \inf_{\hat{\eta}} \sup_{\eta \in \mathcal{F}_{I,J}} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_*.$$

In what follows, we will put a lower bound on the last quantity. We first fix a sensible distance metric on \mathcal{M} . An intuitive way to measure distance between two population size functions is their L_1 distance, $\|\eta_a - \eta_b\|_1 = \int_0^\infty |\eta_a(t) - \eta_b(t)| dt$, but this is unreasonably stringent in that $\|\eta_a - \eta_b\|_1 = \infty$ if η_a and η_b do not agree infinitely far back into the past. Instead we will focus on the following truncated L_1 distance: $\|\eta_a - \eta_b\|_{1,T} \stackrel{\text{def}}{=} \int_0^T |\eta_a(t) - \eta_b(t)| dt$, which measures the discrepancy between η_a and η_b back to some fixed time T in the past.

Henceforth, let $\hat{\eta}$ be any estimator of the population size function that operates on a sample of s independent segregating sites obtained from a sample of n randomly sampled individuals. In *Proofs*, we prove the following main results of our paper:

Theorem 4. Consider the subfamily $\mathcal{F}_{I,J}$ of models described above, and suppose $J > 8$ and $T \geq t_{I+J-1} + \tau_A$. Then,

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{F}_{I,J}} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_{1,T} \geq C \tau_A \frac{(J-8)^2}{J} \frac{\varepsilon}{s} e^{\tau_B / \varepsilon}, \quad [13]$$

where C is a positive constant.




Fig. 1. A family $\mathcal{F}_{I,J}$ of piecewise-constant population size models with $K = I + J$ epochs.




Fig. 2. A family \mathcal{G}_J of population size models with exponential growth in the recent past. This family consists of size histories that are piecewise constant before the bottleneck, and then jump to some level $\gamma\varepsilon$ and undergo (identical) exponential growth from time t_1 to present.

The above theorem applies to all models in $\mathcal{F}_{I,J}$. We now consider the subset $\mathcal{F}_{I,J}^M = \{\eta \in \mathcal{F}_{I,J} : \|\eta\|_\infty < M\}$, which is the set of all models in $\mathcal{F}_{I,J}$ that are bounded by some constant M . For this family of bounded population size functions, a sharper asymptotic lower bound can be obtained as follows.

Theorem 5. Suppose $J > 8$ and $T \geq t_{I+J-1} + \tau_A$. Then,

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{F}_{I,J}^M} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_{1,T} \geq C' \frac{(J-8)^2}{J} \frac{\tau_B \tau_A}{\log s}, \quad [14]$$

where C' is a positive constant.

By specializing $\mathcal{F}_{I,J}^M$, a simplified version of Theorem 5 can be obtained:

Corollary 6. Suppose $T \geq t_{I+J-1} + \tau_A$ and let $\mathcal{F}_{I,*}^M = \cup_{J \geq 1} \mathcal{F}_{I,J}^M$. Then,

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{F}_{I,*}^M} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_{1,T} \geq C'' (T - t_1) \frac{\tau_B}{\log s}, \quad [15]$$

where C'' is a positive constant.

Note that the above lower bounds do not depend on the dimension of the SFS (which is equal to $n-1$). Hence, for a fixed number s of segregating sites considered, using more individuals does not diminish the error bounds.

Bottleneck Followed by Exponential Growth. In the results presented above, we dropped smaller terms to obtain the dominant contribution to our lower bound. Here, we provide a more detailed analysis to study how the model in the recent past (i.e., the period $0 \leq t \leq t_{I-1}$) affects the lower bound. A slight modification of the above results permits us to analyze the following model class, which is of interest in, for example, human genetics (2, 3, 7): Let \mathcal{G}_J be the family of models illustrated in Fig. 2 with exponential growth in the recent past. Specifically, $\eta(t) = \eta_0 e^{-\beta(\eta_0)t}$ for the period $0 \leq t \leq t_1$. The rate of growth $\beta(\eta_0) = \log(\eta_0/\gamma\varepsilon)/t_1$ is defined so that $\eta(t_1) = \gamma\varepsilon$ for all $\eta \in \mathcal{G}_J$, where $\gamma \geq 1$. The part for $t > t_1$ is the same as that for $t > t_{I-1}$ in $\mathcal{F}_{I,J}$ (Fig. 1). We obtain the following result for the subfamily \mathcal{G}_J :

Theorem 7. Consider the subfamily \mathcal{G}_J of models described above, and suppose $J > 8$ and $T \geq t_J + \tau_A$. Then,

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{G}_J} \mathbb{E}_\eta \|\hat{\eta} - \eta\|_{1,T} \geq C \tau_A \frac{(J-8)^2}{J} \frac{\varepsilon}{s} \exp \left[\frac{\tau_B}{\varepsilon} + t_1 \frac{1}{\log(\eta_0) - \log(\gamma\varepsilon)} \right]. \quad [16]$$

Theorem 4 is a measure of how (a lower bound on) estimation error depends on growth following a bottleneck. The two extremes

$\eta_0 \rightarrow \infty$ and $\eta_0 \rightarrow \gamma\varepsilon$ have intuitive interpretations. For large η_0 , the bound in Eq. 16 tends to the corresponding bound given by Theorem 4, as expected since coalescences become increasingly less likely in the first time period. Small η_0 has the effect of “prolonging” the bottleneck, thus increasing the minimax lower bound. In particular, if $\gamma = 1$ then $t_1[(1/\gamma\varepsilon) - (1/\eta_0)]/[\log(\eta_0) - \log(\gamma\varepsilon)] \rightarrow (t_1/\varepsilon)$ as $\eta_0 \rightarrow \gamma\varepsilon$, so that the effect of low population growth on the minimax lower bound is to simply prolong the bottleneck effect by an additional t_1 time periods.

Discussion

In this paper, we have theoretically characterized fundamental limits on the accuracy of demographic inference from data. We have shown that the minimax error rate for estimating the piecewise-constant demography of a single population is at least $O(1/\log s)$, where s is the number of independent segregating sites analyzed. In contrast, the minimax error for many classical estimation problems in statistics (for example, non-parametric regression or density estimation) decays inverse polynomially in the sample size (21). Compared with these problems, exponentially more samples would be required to estimate a population size history function to within a similar magnitude of error. The paper that most closely relates to the present work is by Kim et al. (22), who obtain lower bounds on the amount of exact coalescence time data necessary to distinguish between size histories in a hypothesis testing framework. Since coalescence times are never observed and must be estimated from data, these bounds place a limit on the accuracy with which a population size function can be inferred. The authors also describe an estimator that uses coalescence times (again observed without noise) to accurately recover the underlying population size function with high probability, at a rate that roughly matches the lower bound.

Another line of work centers around the identifiability of the parameter $\eta(t)$ using the SFS. Roughly speaking, a family of statistical models $\{P_{\theta}\}_{\theta \in \Theta}$ defined over a parameter space Θ is identifiable if, for any $\theta_1, \theta_2 \in \Theta$ with $\theta_1 \neq \theta_2$, the sampling distributions induced by P_{θ_1} and P_{θ_2} are different. In our context, this simply says that, for all n , $\xi_n^{(\eta_1)} \neq \xi_n^{(\eta_2)}$ unless $\eta_1 = \eta_2$ almost everywhere. Standard desiderata for statistical estimators (e.g., consistency or unbiasedness) are impossible without identifiability, so it is the weakest possible regularity condition one can impose on a useful family of models.

Perhaps surprisingly, it turns out that, in general, a population size function is not identifiable from the SFS (23). Indeed, for any given $\eta(t)$, it has been shown that an infinite number of smooth functions $F(t)$ exist such that $\xi_n^{(\eta)} = \xi_n^{(\eta+F)}$. Moreover, explicit examples can be constructed that demonstrate this phenomenon (23). On the other hand, these counterexamples consist of functions that exhibit an unbounded frequency of oscillatory behavior near the present time, which is perhaps unrealistic when modeling naturally occurring populations. More recently, it has been shown (19) that identifiability holds for many classes of population size functions used by practitioners (including piecewise constant, piecewise exponential, and piecewise generalized exponential). Furthermore, the number n of sampled individuals sufficient for identifiability can be explicitly given and is a function of the complexity of the underlying class of models being studied (19).

Identifiability asserts that, given an infinite amount of data (specifically, taking the number of segregating sites $s \rightarrow \infty$), the model parameter $\eta(t)$ can be uniquely recovered. In practice, s is finite, and only a perturbed version of the expected frequency spectrum, say $\hat{\xi}_n^{(\eta)}$, is observed. From a practical standpoint, it is important to understand how these perturbations ultimately affect the parameter estimate $\hat{\eta}(t)$. It is this question that forms the starting point for the present work.

A single population evolving under a piecewise-constant demography is a special case of many richer classes of demographic models. For example, it is a (limiting) member of the family of exponential growth models, seen by taking each exponential growth parameter to zero. In the multispecies coalescent setting

(10, 24), multiple population size histories must be estimated, and the error of that estimate must necessarily be lower bounded by that of estimating a single such history. Thus, our result can be expected to apply to a broader class of models than the one we have studied here.

As detailed in *Proofs*, the result in Theorem 5 follows from setting $\varepsilon = \tau_B/\log s$ and $\delta \propto \varepsilon^2 \exp(\tau_B/\varepsilon)$ in the subfamily \mathcal{F}_{JJ}^M . The size $\tau_B/\log s$ is in coalescent units. In terms of the number of individuals, it is proportional to $g_B/\log s$, where g_B is the number of generations corresponding to duration τ_B in the coalescent limit. Intuitively, as the severity of the bottleneck increases, the population is increasingly likely to find its most recent common ancestor (MRCA) during that time; farther back in time than the MRCA, no information is conveyed concerning the demographic events experienced by the population.

One might object to considering models with a bottleneck size that scales inversely with the number s of segregating sites in the data, and it is indeed possible that a better convergence rate may be achievable for populations that are known not to contain a bottleneck. On the other hand, we note that $1/\log s$ decreases sufficiently slowly with s that our result can be expected to apply to many real-world examples. For example, for $s \approx 10^8$, which is a conservative upper bound for most organisms, $g_B/\log s \approx 0.054g_B$. This implies that for populations that have experienced roughly an order-of-magnitude increase in effective population size during their history, accurate estimation of demographic events that occurred before this expansion is difficult using SFS-based methods. Additionally, an interesting aspect of our work is that our minimax lower bounds do not depend on the number n of sampled individuals; increasing n is not enough to overcome the information barrier imposed by the presence of a bottleneck. This is intuitively plausible since, as n increases, the $(n + 1)$ th sampled lineage becomes more likely to coalesce early on.

An interesting question that we have not attempted to analyze is whether the $O(1/\log s)$ rate is optimal, i.e., whether there exists some estimator $\hat{\eta}(t)$ that achieves the minimax lower bound established here. In practice, from Eqs. 2, 8, and 9, it can be seen that naively maximizing the likelihood of the observed SFS with respect to $\eta(t)$ requires solving a nonconvex optimization problem, so that convergence to the global maximum is not even guaranteed. Computational issues aside, finding such an estimator remains an open theoretical challenge.

In closing, we stress that our result is specific to SFS-based estimators, which analyze only independent sites. The main allure of these estimators is their mathematical tractability, rather than their realism. In fact, a rich source of additional information exists in the correlation structure found among linked sites in the genome. Methods that seek to exploit this structure by modeling the action of recombination pose greater mathematical and computational difficulties, but there has been recent progress in this area (20, 25–29). Our result serves to underscore the importance of pursuing more realistic models of genomic evolution, challenging though they may be.

Proofs

Proof of Theorem 1. To simplify the notation, we write $\mathbf{c} = \mathbf{c}^{(\eta)}$ and $\mathbf{c}' = \mathbf{c}^{(\eta')}$. Then, using Eq. 2, we can write

$$D(\eta||\eta') = \sum_{b=1}^{n-1} \xi_{n,b}^{(\eta)} \log \frac{\xi_{n,b}^{(\eta)}}{\xi_{n,b}^{(\eta')}} = \sum_{b=1}^{n-1} \xi_{n,b}^{(\eta)} \left[\log \left(\frac{\langle \mathbf{c}, \mathbf{W}_{n,b} \rangle}{\langle \mathbf{c}', \mathbf{W}_{n,b} \rangle} \right) + \log \left(\frac{\langle \mathbf{c}', \mathbf{V}_n \rangle}{\langle \mathbf{c}, \mathbf{V}_n \rangle} \right) \right].$$

The assumption $\min_{t>t_c} \eta'(t) \geq \max_{t>t_c} \eta(t)$ implies that, for all times $t, t' > t_c$, the instantaneous rate of coalescence at time t in model η is greater than or equal to \geq the instantaneous rate of coalescence at time t' in model η' . Hence, this assumption together with $\eta(t) = \eta'(t)$ for all $0 \leq t \leq t_c$ implies $\langle \mathbf{c} - \mathbf{c}', \mathbf{W}_{n,b} \rangle \leq 0$ for

all $1 \leq b \leq n-1$; equivalently, $\log(\langle \mathbf{c}, \mathbf{W}_{n,b} \rangle / \langle \mathbf{c}', \mathbf{W}_{n,b} \rangle) < 0$. Additionally, $\langle \mathbf{c}' - \mathbf{c}, \mathbf{V}_n \rangle / \langle \mathbf{c}, \mathbf{V}_n \rangle > -1$ and $\log(1+x) \leq x$ for all $x \geq -1$. Combining these facts, we obtain

$$D(\eta \| \eta') \leq \sum_{b=1}^{n-1} \xi_{n,b}^{(\eta)} \log \left(\frac{\langle \mathbf{c}', \mathbf{V}_n \rangle}{\langle \mathbf{c}, \mathbf{V}_n \rangle} \right) \leq \sum_{b=1}^{n-1} \xi_{n,b}^{(\eta)} \frac{\langle \mathbf{c}' - \mathbf{c}, \mathbf{V}_n \rangle}{\langle \mathbf{c}, \mathbf{V}_n \rangle} = \frac{\langle \mathbf{c}' - \mathbf{c}, \mathbf{V}_n \rangle}{\langle \mathbf{c}, \mathbf{V}_n \rangle},$$

where we have used $\sum_{b=1}^{n-1} \xi_{n,b}^{(\eta)} = 1$ in the final equality.

Proof of Lemma 2. We distinguish two particular models, $\eta_\ell, \eta_u \in \mathcal{F}_{IJ}$, which are the lower and the upper envelopes of \mathcal{F}_{IJ} . The function η_ℓ stays constant at h for all $t \geq t_I$, while η_u jumps upward by δ at every change point t_I, \dots, t_{I+J-1} . Hence, $\eta_\ell \leq \eta \leq \eta_u$ pointwise for all $\eta \in \mathcal{F}_{IJ}$. The two enveloping functions will form the basis of subsequent analysis.

Fix $\eta, \eta' \in \mathcal{F}_{IJ}$ and note that, by the definition of \mathcal{F}_{IJ} , one of these functions must pointwise dominate the other. Therefore, assume without loss of generality that $\eta(t) \leq \eta'(t)$ for all t . Then, for all t ,

$$\eta_\ell(t) \leq \eta(t) \leq \eta'(t) \leq \eta_u(t),$$

which implies

$$c_m^{(\eta_u)} \leq c_m^{(\eta)} \leq c_m^{(\eta')} \leq c_m^{(\eta_\ell)},$$

for all $m=2, \dots, n$. Using these inequalities, we conclude

$$c_m^{(\eta')} - c_m^{(\eta)} \leq c_m^{(\eta_u)} - c_m^{(\eta_\ell)},$$

so it suffices to demonstrate Eq. 11 for $c_m^{(\eta_u)} - c_m^{(\eta_\ell)}$. Now, by Eq. 9 and the definition of η_ℓ ,

$$a_m c_m^{(\eta_\ell)} = \sum_{i=1}^I N_i \left[e^{-a_m S_{i-1}^{(\eta_\ell)}} - e^{-a_m S_i^{(\eta_\ell)}} \right] + \sum_{j=1}^J h \left[e^{-a_m S_{I+j-1}^{(\eta_\ell)}} - e^{-a_m S_{I+j}^{(\eta_\ell)}} \right] = \sum_{i=1}^I N_i \left[e^{-a_m S_{i-1}^{(\eta_\ell)}} - e^{-a_m S_i^{(\eta_\ell)}} \right] + h e^{-a_m S_I^{(\eta_\ell)}},$$

where we have used Eq. 10. Similarly,

$$a_m c_m^{(\eta_u)} = \sum_{i=1}^I N_i \left[e^{-a_m S_{i-1}^{(\eta_u)}} - e^{-a_m S_i^{(\eta_u)}} \right] + \sum_{j=1}^J (h + j\delta) \left[e^{-a_m S_{I+j-1}^{(\eta_u)}} - e^{-a_m S_{I+j}^{(\eta_u)}} \right] = \sum_{i=1}^I N_i \left[e^{-a_m S_{i-1}^{(\eta_u)}} - e^{-a_m S_i^{(\eta_u)}} \right] + h e^{-a_m S_I^{(\eta_u)}} + \sum_{j=1}^J j\delta \left[e^{-a_m S_{I+j-1}^{(\eta_u)}} - e^{-a_m S_{I+j}^{(\eta_u)}} \right].$$

Now, using the fact that η_ℓ and η_u agree on the first I epochs, we obtain

$$a_m \left[c_m^{(\eta_u)} - c_m^{(\eta_\ell)} \right] = \sum_{j=1}^J j\delta \left[e^{-a_m S_{I+j-1}^{(\eta_u)}} - e^{-a_m S_{I+j}^{(\eta_u)}} \right] = \delta \sum_{j=1}^J e^{-a_m S_{I+j-1}^{(\eta_u)}} \leq J\delta e^{-a_m \tau_B / \varepsilon}, \quad [17]$$

where the second line follows from telescoping and the fact that $S_{I+j}^{(\eta_u)} = \infty$, while the last line follows from the fact that $\frac{\tau_B}{\varepsilon} \leq S_{I+j-1}^{(\eta_u)}$ for all $j=1, \dots, J$.

Proof of Theorem 3. For ease of notation, define $\mathbf{c} = \mathbf{c}^{(\eta)}$ and $\mathbf{c}' = \mathbf{c}^{(\eta')}$. By Lemma 2,

$$\langle \mathbf{c}' - \mathbf{c}, \mathbf{V}_n \rangle = \sum_{m=2}^n (c'_m - c_m) V_{n,m} \leq J\delta \sum_{m=2}^n \frac{V_{n,m}}{a_m} e^{-a_m \tau_B / \varepsilon} \leq J\delta e^{-\tau_B / \varepsilon} \sum_{m=2}^n \frac{V_{n,m}}{a_m},$$

where the second inequality follows from $e^{-a_m \tau_B / \varepsilon} \leq e^{-\tau_B / \varepsilon}$ for all $m=2, \dots, n$. Now, noting that $\sum_{m=2}^n (V_{n,m}/a_m)$ corresponds to the total tree length for the constant population size function $\eta \equiv 1$ and using Eq. 5, we obtain

$$\langle \mathbf{c}' - \mathbf{c}, \mathbf{V}_n \rangle \leq J\delta e^{-\tau_B / \varepsilon} 2H_{n-1}. \quad [18]$$

To finish the proof, recall that $\langle \mathbf{c}, \mathbf{V}_n \rangle$ is the total expected branch length of the coalescent tree under model η . Since $\min_t \eta(t) = \varepsilon$, we have that $\langle \mathbf{c}, \mathbf{V}_n \rangle$ is at least as large as the corresponding quantity under a model with constant population size ε . By Eq. 5, the total expected tree length under the latter model equals $2\varepsilon H_{n-1}$. Thus, $\langle \mathbf{c}, \mathbf{V}_n \rangle \geq 2\varepsilon H_{n-1}$, and combining this result with Eq. 18 gives

$$\frac{\langle \mathbf{c}' - \mathbf{c}, \mathbf{V}_n \rangle}{\langle \mathbf{c}, \mathbf{V}_n \rangle} \leq J \frac{\delta}{\varepsilon} e^{-\tau_B / \varepsilon}.$$

Finally, Eq. 12 follows from this inequality and Theorem 1.

Proof of Theorem 4. Our proof uses a generalized form of Fano's inequality (30). Adapted to our setting and notation, the method reads as follows.

Theorem 8 (Fano's method). *Consider a space \mathcal{M} of population size models. Let $r \geq 2$ be an integer, and let $S_n^r = \{\eta_1, \eta_2, \dots, \eta_r\} \subset \mathcal{M}$ contain r population size functions such that for all $a \neq b$, $\|\eta_a - \eta_b\|_* \geq \alpha_r$ and $D(\xi_n^{(\eta_a)} \| \xi_n^{(\eta_b)}) \leq \beta_r$. Let $\hat{\eta}^{(n,s)} = \hat{\eta}^{(n,s)}(X_1, \dots, X_s)$ be an estimator of η based on the SFS data X_1, \dots, X_s sampled independently from $\xi_n^{(\eta)}$; i.e., X_1, \dots, X_s are SFS data for n individuals at s independent segregating sites. Then,*

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{M}} \mathbb{E}_{\eta} \|\hat{\eta}^{(n,s)} - \eta\|_* \geq \frac{\alpha_r}{2} \left(1 - \frac{s \cdot \beta_r + \log 2}{\log r} \right). \quad [19]$$

This theorem places a lower bound on the minimax rate of convergence of a population size history estimator based on the SFS.

For $\eta \in \mathcal{F}_{IJ}$, let w_j denote the variable $\in \{0,1\}$ indicating whether η jumps by δ at change point t_{I+j} . Let $\mathcal{Y} = \{\mathbf{w} = (w_0, \dots, w_{J-1}) \mid w_i \in \{0,1\}\}$, where $J \geq 8$. By the Varshamov-Gilbert lemma (see ref. 31, Lemma 4.7), there exist $\mathcal{X} = \{\mathbf{w}^0, \dots, \mathbf{w}^M\} \subset \mathcal{Y}$ such that (i) $\mathbf{w}^0 = (0, \dots, 0)$, (ii) $M \geq 2^{J/8}$, and (iii) $H(\mathbf{w}^i, \mathbf{w}^j) \geq J/8$, where $H(\cdot, \cdot)$ denotes the Hamming distance.

Let $\mathcal{F}_{IJ}^{\mathcal{X}}$ denote the subset of $2^{J/8} + 1$ functions in \mathcal{F}_{IJ} with the indicator variable for δ jumps at t_I, \dots, t_{I+J-1} given by $\mathbf{w} \in \mathcal{X}$. Then, for any two $\eta_a \neq \eta_b \in \mathcal{F}_{IJ}^{\mathcal{X}}$, we have

$$\|\eta_a - \eta_b\|_{1,T} \geq \frac{J}{8} \cdot \tau_A \cdot \delta. \quad [20]$$

Using Theorem 8 via Eq. 20 and Theorem 3, we obtain

$$\inf_{\hat{\eta}} \sup_{\eta \in \mathcal{F}_{IJ}^{\mathcal{X}}} \mathbb{E}_{\eta} \|\hat{\eta}^{(n,s)} - \eta\|_{1,T} \geq \frac{J \cdot \tau_A \cdot \delta}{16} \left[1 - \frac{sJ \frac{\delta}{\varepsilon} e^{-\tau_B / \varepsilon} + \log 2}{\log(2^{J/8} + 1)} \right] \geq \frac{J \cdot \tau_A \cdot \delta}{16} \left[1 - \frac{sJ \frac{\delta}{\varepsilon} e^{-\tau_B / \varepsilon} + \log 2}{\frac{J}{8} \log 2} \right]. \quad [21]$$

We now optimize the bound with respect to δ . A straightforward calculation shows that the maximum is attained at

$$\delta^* = \frac{(J-8)\log 2}{16J} \left(\frac{\varepsilon}{s}\right) e^{\tau_B/\varepsilon}, \quad [22]$$

and setting $\delta = \delta^*$ in Eq. 21 yields the result.

Proof of Theorem 5. The result is obtained by scaling ε with the number of segregating sites s . Denote this scaling by $\varepsilon(s)$; we will determine $\varepsilon(s)$ that produces the largest possible lower bound. Starting from Eq. 22 in the proof of Theorem 4, note that δ^* scales as $(\varepsilon/s)e^{\tau_B/\varepsilon} =: f(\varepsilon)$. To satisfy the constraint that $\|\eta\|_\infty < M$ for all $\eta \in \mathcal{F}_{IJ}^M$ and s , the condition

$$\limsup_{s \rightarrow \infty} \max \left\{ \frac{\varepsilon(s)}{s} e^{\tau_B/\varepsilon(s)}, \varepsilon(s) \right\} < \infty \quad [23]$$

must therefore hold. This implies that $\varepsilon(s)s^p \rightarrow \infty$ as $s \rightarrow \infty$ for all $p > 0$. Suppose that $q = \liminf_{s \rightarrow \infty} [(\varepsilon(s)\log s)/\tau_B] < 1$; note that $\varepsilon(s) > 0$ implies $q > 0$. Then there exists a diverging sequence $s_1, s_2, \dots \rightarrow \infty$ with $\log(s_i) < [(1+q)/2][\tau_B/(\varepsilon(s_i))]$ for all i , whence

$$\begin{aligned} \limsup_{s \rightarrow \infty} \frac{\varepsilon(s)}{s} e^{\tau_B/\varepsilon(s)} &\geq \limsup_{i \rightarrow \infty} \frac{\varepsilon(s_i)}{s_i} e^{\frac{2}{1+q} \log(s_i)} \\ &= \limsup_{i \rightarrow \infty} \varepsilon(s_i) s_i^{\frac{1-q}{1+q}} = \infty. \end{aligned}$$

From this, it follows that $\varepsilon(s) \geq \tau_B/\log s$ for sufficiently large s . Now, on the interval $(0, \infty)$, the function $f(\varepsilon)$ is convex with a unique minimum at $\varepsilon = \tau_B$. Let ε' be a point where $f(\varepsilon') > f(\tau_B/\log s) = \tau_B/\log s$. Then $\varepsilon' \notin [\tau_B/\log s, \tau_B]$. If $\varepsilon' > \tau_B$, then $f(\varepsilon') < (\varepsilon'/s)e^1$. Since $\frac{\tau_B}{\log s} < f(\varepsilon')$, we then conclude $\varepsilon' > s\tau_B/(e^1 \log s)$, which is not bounded as $s \rightarrow \infty$.

In summary, we see that the largest possible lower bound that obeys Eq. 23 must have $f(\varepsilon)$ asymptotically $\leq \tau_B/\log s$, and that this bound is achieved by setting $\varepsilon(s) = \tau_B/\log s$. Plugging this in to Eq. 19 yields the claim.

Proof of Corollary 6. For $c \in (0, 1)$, choose J large enough so that $(J-8)/J > c$, and fix τ_A so that $T = t_1 + J\tau_A$. Then

$(J-8)\tau_A \geq cJ\tau_A = c(T-t_1)$. Substituting the above inequalities into Eq. 14 and letting $C'' = C'c^2$ yields the desired result.

Proof of Theorem 7. The theorem is obtained by suitably modifying the preceding results to account for the effect of exponential growth in the first period. Let η_u, η_t be the analogously defined upper and lower envelope functions for \mathcal{G}_J . Then

$$\begin{aligned} \int_0^{t_{j+1}} \frac{ds}{\eta_u(s)} &= \frac{e^{\beta(\eta_0)t_1} - 1}{\eta_0\beta(\eta_0)} + \frac{\tau_B}{\varepsilon} + \sum_{i=3}^{j+1} \frac{t_i - t_{i-1}}{N_i} \\ &= t_1 \frac{1}{\gamma\varepsilon} \frac{1}{\eta_0} + \frac{\tau_B}{\varepsilon} + \sum_{i=3}^{j+1} \frac{t_i - t_{i-1}}{N_i}, \end{aligned}$$

where we have used the definition of $\beta(\eta_0)$ in the second equality. Since all size histories in \mathcal{G}_J are equal up to period t_2 , the steps of Lemma 2 all go through unchanged. Starting from Eq. 17, we obtain the modified bound

$$a_m \left[c_m^{(\eta_u)} - c_m^{(\eta_t)} \right] \leq J\delta \exp \left\{ -a_m t_1 \frac{1}{\gamma\varepsilon} \frac{1}{\eta_0} \right\} e^{-a_m \tau_B/\varepsilon}. \quad [24]$$

Propagating the modified bound (Eq. 24) through Theorems 3 and 4 ultimately yields the claim.

ACKNOWLEDGMENTS. We thank Anand Bhaskar for helpful comments on a draft of this paper and for suggesting Corollary 6 to simplify the presentation of the main result. We also thank Jack Kamm and Jeff Spence for useful feedback. This research is supported in part by a Citadel Fellowship (to J.T.), National Institutes of Health Grant R01-GM109454 (to Y.S.S.), a Packard Fellowship for Science and Engineering (to Y.S.S.), and a Miller Research Professorship (to Y.S.S.).

1. Abecasis GR, et al.; 1000 Genomes Project Consortium (2010) A map of human genome variation from population-scale sequencing. *Nature* 467(7319):1061–1073.
2. Nelson MR, et al. (2012) An abundance of rare functional variants in 202 drug target genes sequenced in 14,002 people. *Science* 337(6090):100–104.
3. Tennessen JA, et al.; Broad GO; Seattle GO; NHLBI Exome Sequencing Project (2012) Evolution and functional impact of rare coding variation from deep sequencing of human exomes. *Science* 337(6090):64–69.
4. Fu W, et al.; NHLBI Exome Sequencing Project (2013) Analysis of 6,515 exomes reveals the recent origin of most human protein-coding variants. *Nature* 493(7431):216–220.
5. Nielsen R (2000) Estimation of population parameters and recombination rates from single nucleotide polymorphisms. *Genetics* 154(2):931–942.
6. Gutenkunst RN, Hernandez RD, Williamson SH, Bustamante CD (2009) Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genet* 5(10):e1000695.
7. Coventry A, et al. (2010) Deep resequencing reveals excess rare recent variants consistent with explosive population growth. *Nat Commun* 1:131.
8. Gazave E, et al. (2014) Neutral genomic regions refine models of recent rapid human population growth. *Proc Natl Acad Sci USA* 111(2):757–762.
9. Gravel S, et al.; 1000 Genomes Project (2011) Demographic history and rare allele sharing among human populations. *Proc Natl Acad Sci USA* 108(29):11983–11988.
10. Excoffier L, Dupanloup I, Huerta-Sánchez E, Sousa VC, Foll M (2013) Robust demographic inference from genomic and SNP data. *PLoS Genet* 9(10):e1003905.
11. Bhaskar A, Wang YXR, Song YS (2015) Efficient inference of population size histories and locus-specific mutation rates from large-sample genomic variation data. *Genome Res* 25(2):268–279.
12. Kingman JFC (1982) The coalescent. *Stochastic Process Appl* 13(3):235–248.
13. Kingman JFC (1982) On the genealogy of large populations. *J Appl Probab* 19A:27–43.
14. Kingman JFC (1982) *Exchangeability in Probability and Statistics*, eds Koch G, Spizzichino F (North-Holland, Amsterdam), pp 97–112.
15. Kimura M (1969) The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations. *Genetics* 61(4):893–903.
16. Griffiths R, Tavaré S (1998) The age of a mutation in a general coalescent tree. *Commun Stat Stochastic Models* 14(1-2):273–295.
17. Polanski A, Bobrowski A, Kimmel M (2003) A note on distributions of times to coalescence, under time-dependent population size. *Theor Popul Biol* 63(1):33–40.
18. Polanski A, Kimmel M (2003) New explicit expressions for relative frequencies of single-nucleotide polymorphisms with application to statistical inference on population growth. *Genetics* 165(1):427–436.
19. Bhaskar A, Song YS (2014) Descartes’ rule of signs and the identifiability of population demographic models from genomic variation data. *Ann Stat* 42(6):2469–2493.
20. Li H, Durbin R (2011) Inference of human population history from individual whole-genome sequences. *Nature* 475(7357):493–496.
21. Tsybakov AB (2009) *Introduction to Nonparametric Estimation* (Springer, New York).
22. Kim J, Mossel E, Rácz MZ, Ross N (2014) Can one hear the shape of a population history? *Theor Popul Biol* 100:26–38.
23. Myers S, Fefferman C, Patterson N (2008) Can one learn history from the allelic spectrum? *Theor Popul Biol* 73(3):342–348.
24. Chen H (2012) The joint allele frequency spectrum of multiple populations: A coalescent theory approach. *Theor Popul Biol* 81(2):179–195.
25. Paul JS, Steinrücken M, Song YS (2011) An accurate sequentially Markov conditional sampling distribution for the coalescent with recombination. *Genetics* 187(4):1115–1128.
26. Sheehan S, Harris K, Song YS (2013) Estimating variable effective population sizes from multiple genomes: A sequentially markov conditional sampling distribution approach. *Genetics* 194(3):647–662.
27. Steinrücken M, Paul JS, Song YS (2013) A sequentially Markov conditional sampling distribution for structured populations with migration and recombination. *Theor Popul Biol* 87:51–61.
28. Rasmussen MD, Hubisz MJ, Gronau I, Siepel A (2014) Genome-wide inference of ancestral recombination graphs. *PLoS Genet* 10(5):e1004342.
29. Schiffels S, Durbin R (2014) Inferring human population size and separation history from multiple genome sequences. *Nat Genet* 46(8):919–925.
30. Yu B (1997) Assouad, Fano, and Le Cam. *Festschrift for Lucien Le Cam*, ed Pollard D, Torgersen E, Yang GL (Springer, New York), pp 423–435.
31. Massart P (2007) *Concentration Inequalities and Model Selection* (Springer, Berlin).