General Triallelic Frequency Spectrum Under Demographic Models with Variable Population Size

Paul A. Jenkins,* Jonas W. Mueller,[†] and Yun S. Song^{*,§,1}

*Department of Statistics, University of Warwick, Coventry CV4 7AL, United Kingdom, [†]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, and [§]Department of Statistics and

[§]Computer Science Division, University of California, Berkeley, California 94720

ABSTRACT It is becoming routine to obtain data sets on DNA sequence variation across several thousands of chromosomes, providing unprecedented opportunity to infer the underlying biological and demographic forces. Such data make it vital to study summary statistics that offer enough compression to be tractable, while preserving a great deal of information. One well-studied summary is the site frequency spectrum—the empirical distribution, across segregating sites, of the sample frequency of the derived allele. However, most previous theoretical work has assumed that each site has experienced at most one mutation event in its genealogical history, which becomes less tenable for very large sample sizes. In this work we obtain, in closed form, the predicted frequency spectrum of a site that has experienced at most two mutation events, under very general assumptions about the distribution of branch lengths in the underlying coalescent tree. Among other applications, we obtain the frequency spectrum of a triallelic site in a model of historically varying population size. We demonstrate the utility of our formulas in two settings: First, we show that triallelic sites are more sensitive to the parameters of a population that has experienced historical growth, suggesting that they will have use if they can be incorporated into demographic inference. Second, we investigate a recently proposed alternative mechanism of mutation in which the two derived alleles of a triallelic site are created *simultaneously* within a single individual, and we develop a test to determine whether it is responsible for the excess of triallelic sites in the human genome.

BECAUSE of the recent advances in DNA sequencing technologies, it has become feasible to obtain data on sequence variation across tens of thousands of chromosomes (e.g., Coventry *et al.* 2010; Keinan and Clark 2012; Nelson *et al.* 2012; Tennessen *et al.* 2012) and hence to study the impact of variants of very low population frequency. Classical models underlying population genetic studies have typically assumed that each site is affected by at most one mutation event in the genealogical history relating a sample, but for very large samples this assumption is less tenable. One must then account for sites experiencing repeat mutations, which skew the site frequency spectrum and can generate triallelic and even quadra-allelic sites. Triallelic sites are therefore becoming increasingly common, appearing as

a few percent of segregating sites in large-scale resequencing studies, particularly as the threshold on masking sites below a given minor allele frequency is being reduced. There are now examples of studies that have found an association between a triallelic single-nucleotide polymorphism (SNP) and a disease phenotype, including coronary heart disease (Crawford *et al.* 2006) and inflammatory bowel disease (discussed in Hübner *et al.* 2007).

Triallelic sites also have potential use in inference, using frequency spectrum data. The observed frequency spectrum of *diallelic* sites is well recognized as an important summary of genomic data, maintaining a great deal of the information encapsulated by the full data while being relatively simple to interpret. It is therefore well studied: The effects of a host of modeling assumptions on the frequency spectrum have been investigated and many theoretical predictions have been made, typically using either coalescent-based or diffusion-based models. For example, one can obtain analytic results incorporating the effects of a population of varying size (Griffiths and Tavaré 1998; Wooding and Rogers 2002; Polanski and Kimmel 2003; Polanski *et al.* 2003), selection (Griffiths 2003), and population subdivision with instantaneous migration events (Chen 2012).

Copyright © 2014 by the Genetics Society of America
doi: 10.1534/genetics.113.158584

Manuscript received October 10, 2013; accepted for publication November 6, 2013;
published Early Online November 7, 2013.

Supporting information is available online at <http://www.genetics.org/lookup/suppl/doi:10.1534/genetics.113.158584/-/DC1>.

¹Corresponding author: Department of Electrical Engineering and Computer Sciences, University of California, 683 Soda Hall, No. 1776, Berkeley, CA 94720-1776.

E-mail: yss@cs.berkeley.edu

The Poisson random field framework of Sawyer and Hartl (1992) is attractive in this respect because of its amenability to the incorporation of natural selection (Sawyer and Hartl 1992; Bustamante *et al.* 2001). This and other diffusion-based approaches can also be extended to obtain numerical solutions for more complicated underlying population demographic histories, including a single population of variable size (Williamson *et al.* 2005; Evans *et al.* 2007; Boyko *et al.* 2008) or a hierarchy of splitting subpopulations with restricted migration between them (Gutenkunst *et al.* 2009; Lukić *et al.* 2011; Lukić and Hey 2013). Essentially, one writes down the Kolmogorov forward equation for the underlying diffusion approximation and then obtains a numerical solution using finite differences (Williamson *et al.* 2005; Evans *et al.* 2007; Gutenkunst *et al.* 2009) or spectral methods (Lukić *et al.* 2011; Lukić and Hey 2013). Examples of inference using the frequency spectrum such as these are important because they can help us learn about recent human population history, estimate the strength of natural selection, and calibrate our expectations prior to a disease association study. However, none of these approaches make use of the information from triallelic sites since they rely on an infinite-sites assumption in which triallelic sites are never observed (although see Desai and Plotkin 2008; Song and Steinrücken 2012; Steinrücken *et al.* 2013). There have been some extensions to incorporate recurrent mutations into the theory of the frequency spectrum (Sargsyan 2006; Hobolth and Wiuf 2009; Jenkins and Song 2011; Bhaskar *et al.* 2012), but with the exception of Sargsyan (2006) these all assume a simple demography of a stationary, panmictic population of constant size.

In this article, we obtain a closed-form expression for the sample frequency spectrum of a site that has experienced two mutation events, under an extension of the standard coalescent model that allows for very general assumptions about the distribution of times between coalescence events. This allows us to obtain predictions for the shape of the frequency spectrum, allowing for both recurrent mutations and varying historical population size.

To emphasize the usefulness of our results, we consider two applications. First, we investigate the sensitivity of the triallelic frequency spectrum to the assumed demographic history. In particular, our interest is in the following question: How much power to distinguish between different demographic models do we gain by looking at a triallelic, rather than diallelic, site? In a manner quantified further below, we show that although triallelic sites are far less abundant than diallelic sites, they have rather greater value *per site* in capturing the effects of demographic history.

This application relies on a frequency spectrum in which the two mutation events arose independently during the genealogical history of the site. Recently, however, Hodgkinson and Eyre-Walker (2010) noted that there are approximately twice as many triallelic sites in the human genome as would be expected by chance. They explored a number of potential explanations and ultimately favored the idea of

Figure 1 A possible mechanism of simultaneous mutation suggested by Hodgkinson and Eyre-Walker (2010). Each nucleotide of a single base pair mutates due to instability of the mismatch. Replication of the resulting base pair, using each of the parental strands, results in two new derived alleles. Here, the G allele has given rise to a C and a T (boxed).

a new mutational mechanism: namely, the simultaneous generation of two new alleles due to mutation within a single individual. Although the precise mechanism is unknown, they suggest the instability of base mismatches as a plausible explanation. For example, a mutation of a G = C base pair to an unstable G = A mismatch could give rise to a further mutation to C = A. DNA replication of this mismatch means the ancestral G = C has given rise to both a derived A = T and a derived C = G base pairing (Figure 1). Another possibility is that both strands of the DNA duplex mutate simultaneously due to a chemical or a radiation event. As a second application of our results, we design and implement a frequency spectrum-based test for the hypothesis that a subset of triallelic sites was generated by a *simultaneous* mutation event within a single individual, giving rise to the two derived alleles. The test allows us to account for variable historical population size explicitly, and when we do so we do not find evidence in favor of the existence of such a mechanism (although, as we discuss below, it is likely that this is further confounded by population subdivision in the samples used).

This article is organized as follows. In the following section we first introduce some notation and summarize some previous results. We then obtain closed-form formulas for a triallelic frequency spectrum under a general model of coalescence times distributions. In the sections thereafter we consider two applications. First, we perform an extensive simulation study to discern between the sensitivities of diallelic and triallelic frequency spectra to the underlying demographic model. Second, we obtain the triallelic frequency spectrum under the proposed simultaneous-mutation mechanism of Hodgkinson and Eyre-Walker (2010), and develop a likelihood ratio test to compare it with the null triallelic frequency spectrum under independently-occurring mutations. The test is applied to sequence data taken from the Environmental Genome Project (NIEHS SNPs 2011) and the SeattleSNPs project (SeattleSNPs 2011) to examine whether some fraction of the triallelic sites in these datasets are in fact the product of simultaneous mutations.

Notation and Previous Results

In this section we introduce our notation and summarize some existing results. Denote by N_0 the diploid effective population size in the present generation and by u the probability of

a mutation event at a given locus per meiosis. For simplicity we assume throughout that the “locus” is a single site, although we note that the theory extends easily to other loci that may be of interest. Let $\theta = 4N_0u$ be the population-scaled mutation rate, which we take to be fixed in the usual diffusion limit as $N_0 \rightarrow \infty$. In this limit and on a timescale of $2N_0$ generations, we denote by N_t the effective population size at time t back in the past, which we take to be a nonrandom function of time such that $N_t \gg 1$ so that a coalescent limit exists for all times (see Slatkin and Hudson 1991 and Griffiths and Tavaré 1994 for details). We assume a general K -allele mutation model with mutation transition matrix $\mathbf{P} = (P_{ij})$, so that P_{ij} is the probability forward in time of a mutation taking allele i to allele j , given that a type i mutated. It is usual to treat $\mathbf{P}$ (and K) as fixed and known. We further denote by $a \in \{1, \dots, K\}$ the ancestral allele at the site of interest and by $\mathbf{n} = (n_1, n_2, \dots, n_K)$ the unordered sample configuration taken from that site, with total sample size $n = \sum_{i=1}^K n_i$. A unit K vector whose k th entry is 1 and all other entries are 0 is denoted by $\mathbf{e}_k$. Finally, let E_s denote the event that there were precisely s mutation events at the site in the genealogical history relating the sample.

The sample frequency spectrum can be obtained first by finding the probability of the observed sample configuration under the assumptions of an appropriate coalescent model. This may be partitioned according to the number of mutation events in the genealogical history relating the sample. However, for humans the average per-generation mutation rate for SNPs is small; recent studies estimate that on average $u \approx 1.2 \times 10^{-8}$ (Campbell *et al.* 2012; Kong *et al.* 2012). Classical population genetics results on the frequency spectrum can be obtained formally by conditioning on precisely one mutation event in the history of the site and then letting $\theta \rightarrow 0$. Denoting the ancestral and derived alleles in a diallelic model respectively by a and b , it is well known (Watterson 1975; Fu 1995; Griffiths and Tavaré 1998) that for a constant population size, $N_t \equiv N_0$, and a sample configuration of the form $\mathbf{n} = (n_a, n_b)$,

$$\phi(i) := \lim_{\theta \rightarrow 0} \mathbb{P}[(n_a, n_b) = (n - i, i) | E_1] = \frac{i^{-1}}{\sum_{j=1}^{n-1} j^{-1}}, \quad (1)$$

since $\mathbb{P}[(n_a, n_b) = (n - i, i) | E_1] = \theta i^{-1} + O(\theta^2)$. We refer to the quantity $\phi(i)$ as the *sample* frequency spectrum [as distinguished from the density of the expected number of mutations at each frequency $x \in (0, 1)$ in a population of genomes comprising many polymorphic sites, which is also referred to as the (site) frequency spectrum]. Throughout this work we obtain the sample frequency spectrum in a finite-alleles model and in the limit as $\theta \rightarrow 0$, after conditioning on the required number of mutation events (for triallelic sites, at least two mutation events are of course necessary). In fact, the result (1) is usually obtained by positing a model of *infinitely many sites* of mutation and then finding the distribution of the number of copies

of the mutant allele at any random position at which a mutation occurred. Because we condition on looking at a mutant site, this distribution is equivalent to that of a finite-alleles model at a *fixed* site and conditioned on one mutation event, with the implicit assumption that $P_{aa} = 0$ so that the overall rate of mutation in the two models is the same.

There are two extensions to the above result that are relevant to the present work. The first is to general coalescent trees in which the collection of intercoalescence times, $\mathbf{T} = (T_n, T_{n-1}, \dots, T_2)$ is not necessarily given by the standard sequence of independent, exponentially distributed random variables. In a standard coalescent model we have that the time T_k during which there exist k distinct ancestors to the sample satisfies $T_k \sim \text{Exp}(\frac{k}{2})$ on the coalescent timescale. However, certain extensions to this model yield a more complicated distribution for $\mathbf{T}$ but leave the topological structure of the tree otherwise unchanged. In this setting, Griffiths and Tavaré (1998) have obtained the following result: Under a coalescent model with general intercoalescence times $\mathbf{T}$ and conditional on precisely one mutation event at a given site, the sample frequency spectrum is given by

$$\phi(i) = \frac{\sum_{k=2}^n \alpha_k^{(n-i,i)} \mathbb{E}(T_k)}{\sum_{k=2}^n \beta_k \mathbb{E}(T_k)}, \quad (2)$$

where $\alpha_k^{(n-i,i)} = ((n-i-1)!(i-1)!/(n-1)!)k(k-1) \binom{n-k}{i-1}$ and $\beta_k = k$. One application of this result is to a coalescent model with a nonconstant population size N_t . The distribution for $\mathbf{T}$ does not have a simple form, but an expression for $\mathbb{E}(T_k)$ is given by Griffiths and Tavaré (1998), and an expression for the marginal density of T_k is given by Wooding and Rogers (2002), Polanski *et al.* (2003), and Polanski and Kimmel (2003). In the *Appendix* we provide a new proof of (2), to illustrate our general strategy. For now we merely remark that the topological structure of any polymorphic site having experienced precisely one mutation event in its genealogical history must be of the form shown in Figure 2. Coalescent trees of this form are studied in detail by Wiuf and Donnelly (1999).

A second extension of (1) is to allow for two mutation events at a polymorphic site. In this case we must consider the exact form of the mutation transition matrix $\mathbf{P}$. In particular, it may allow for the second mutation to revert a derived allele to its ancestral state (a *back* mutation) or for the second mutation to create a second independent copy of the extant derived allele (a *parallel* mutation). Such mutations do not give rise to triallelic sites, whereas in practice we typically identify sites having experienced two mutations only when three alleles are actually observed. Thus, in extending the definition of the sample frequency spectrum to triallelic sites, we condition on *observing* three alleles, an event we denote O_3 , rather than E_2 . Jenkins and Song (2011) have obtained the following result: Under a standard coalescent model with $N_t \equiv N_0$ the triallelic sample frequency spectrum is given by

Figure 2 A coalescent tree with one mutation. The allele of each leaf is annotated. Also annotated is the variable l_a determining the number of lineages ancestral to allele a at the time of the sole mutation event; in this example, $l_a = 3$.

$$\begin{aligned} \phi(n_a, n_b, n_c) &:= \lim_{\theta \rightarrow 0} \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c | O_3), \\ &= \frac{1}{C} \left[P_{ab} P_{bc} d(n_a, n_b, n_c) + P_{ac} P_{cb} d(n_a, n_c, n_b) \right. \\ &\quad \left. + P_{ab} P_{ac} \left(\frac{1}{n_b n_c} - d(n_a, n_b, n_c) - d(n_a, n_c, n_b) \right) \right], \end{aligned} \quad (3)$$

where

$$\begin{aligned} C &= \left[\sum_{x \neq a} \sum_{y \neq a, x} P_{ax} P_{xy} \right] (H_n + \frac{1}{n} - 2) \\ &\quad + \left[\sum_{x \neq a} \sum_{y \neq a, x} P_{ax} P_{ay} \right] \left(\frac{(H_{n-1})^2}{2} - \frac{H_{n-1}^{(2)}}{2} - H_n - \frac{1}{n} + 2 \right), \end{aligned}$$

and

$$\begin{aligned} d(n_a, n_b, n_c) &= \frac{1}{(n_a + n_b)(n_a + n_b - 1)} \\ &\quad \times \left[1 + \frac{n}{n_c} - \frac{2n(H_n - H_{n-1})}{n_a + n_b + 1} \right], \\ H_m &= \sum_{j=1}^m \frac{1}{j}, \\ H_m^{(2)} &= \sum_{j=1}^m \frac{1}{j^2}. \end{aligned}$$

In the above expression a , b , and c are distinct alleles with $n_a + n_b + n_c = n$; a is the ancestral allele and b and c are derived alleles. We set $H_0 := 0$ by convention. If the diagonal of $\mathbf{P}$ is zero, then the sums involving $\mathbf{P}$ in C respectively simplify to $[1 - (\mathbf{P}^2)_{aa}]$ and $[1 - (\mathbf{P}\mathbf{P}^T)_{aa}]$. (That this simplification requires the diagonal of $\mathbf{P}$ to be 0 was inadvertently omitted from Jenkins and Song 2011, Corollary 6.1.)

General Triallelic Frequency Spectrum

In this section we obtain a closed-form expression for the sample frequency spectrum of a triallelic site under a general coalescent model with variable population

Figure 3 Coalescent trees with two mutations. (A) Two nested mutations. (B) Two nonnested mutations. The allele of each leaf is annotated. Also annotated are variables determining the number of each type at the times of the mutation events; for example, in A we have $m = 1$, $l_y = 3$, and $l_o = 3$.

size, N_t . This generalizes (2) to the case of two mutation events at a single site and generalizes (3) to the case of a variable population size. Our arguments and notation are similar to those in Jenkins and Song (2011), and a brief proof is deferred to an Appendix. In that article, a key observation is that the event E_2 can be partitioned as follows:

E_{2N} : The two mutation events are genealogically nested.

E_{2NN} : The two mutation events are genealogically non-nested and at least one of them does not reside on the basal (adjacent to the root) branches of the tree.

$E_{2\mathcal{B}}$: The two mutation events reside on the two different basal (adjacent to the root) branches of the tree.

E_{2S} : The two mutation events reside on the same branch of the tree.

Only the first two cases can lead to a triallelic site, and so we do not consider the last two any further. It is straightforward to obtain analogous generalizations for the last two cases, although we omit them. The events E_{2N} and E_{2NN} are illustrated in Figure 3. In these examples, the older of the two mutation events gives rise to the allele b and the younger gives rise to the allele c , and the subsets of E_{2N} and E_{2NN} satisfying these constraints are denoted by $E_{2N}^{(b,c)}$ and $E_{2NN}^{(b,c)}$, respectively. To find the sample frequency spectrum we first consider the joint probability of observing our triallelic sample with each of these events.

Lemma 1. Let $n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c$ denote a triallelic sample as in (3). Under a coalescent model with time-dependent population size N_t , the joint probability of such a sample together with the way the two mutations are placed on the coalescent tree satisfies

$$\begin{aligned} \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c, E_{2N}^{(b,c)}) \\ = \frac{\theta^2}{4} P_{ab} P_{bc} \sum_{k=3}^{n_a + n_b + 1} \sum_{j=2}^{k-1} C_{j,k}^{(n_a, n_b)} \mathbb{E}[T_j T_k] + O(\theta^3), \end{aligned} \quad (4)$$

$$\begin{aligned} \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c, E_{2\mathcal{N}\mathcal{N}}^{(b,c)}) \\ = \frac{\theta^2}{4} P_{ab} P_{ac} \sum_{k=3}^{n_a+n_b+1} \sum_{j=2}^k F_{j,k}^{(n_a, n_b)} \mathbb{E}[T_j T_k] + O(\theta^3), \end{aligned} \quad (5)$$

as $\theta \rightarrow 0$. Furthermore,

$$\mathbb{P}(E_{2\mathcal{N}}^{(b,c)}) = \frac{\theta^2}{4} P_{ab} P_{bc} \sum_{k=3}^n \sum_{j=2}^{k-1} D_{j,k} \mathbb{E}[T_j T_k] + O(\theta^3), \quad (6)$$

$$\mathbb{P}(E_{2\mathcal{N}\mathcal{N}}^{(b,c)}) = \frac{\theta^2}{4} P_{ab} P_{ac} \sum_{k=3}^n \sum_{j=2}^k G_{j,k} \mathbb{E}[T_j T_k] + O(\theta^3). \quad (7)$$

The coefficients in the above expressions are

$$\begin{aligned} C_{j,k}^{(n_a, n_b)} &= \sum_{l=j-1}^{k-2} \binom{n_a-1}{l-1} \binom{n_b-1}{k-l-2} \binom{k-j}{k-1-l} \binom{n-1}{k-1}^{-1} \\ &\quad \times \binom{k-1}{k-l}^{-1} j(j-1), \end{aligned}$$

$$D_{j,k} = j \left[k \binom{k-2}{j-1} - (j-1) \binom{k-1}{j} \right] \binom{k-1}{j-1}^{-1},$$

$$\begin{aligned} F_{j,k}^{(n_a, n_b)} &= \sum_{l=(j-2) \vee 1}^{k-2} \binom{n_a-1}{l-1} \binom{n_b-1}{k-l-2} \binom{k-j}{k-2-l} \\ &\quad \times \binom{n-1}{k-1}^{-1} \binom{k-1}{l+1}^{-1} \frac{j(j-1)}{1+\delta_{j,k}}, \end{aligned}$$

$$G_{j,k} = \left(k(j-1) - \frac{2\delta_{j,2}}{k-1} \right) \frac{1}{1+\delta_{j,k}},$$

where $\delta_{j,k}$ denotes the Kronecker delta.

Proof. See the Appendix. $\square$

From Lemma 1, we can obtain our main result in a straightforward manner.

Theorem 1. Let $n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c$ denote a triallelic sample as above: a , b , and c are distinct alleles with $n_a + n_b + n_c = n$; a is the ancestral allele and b and c are derived alleles. Under a coalescent model with time-dependent population size N_t and in which mutation events occur independently in the tree, the sample frequency spectrum is

$$\begin{aligned} \phi_0(n_a, n_b, n_c) &:= \lim_{\theta \rightarrow 0} \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c | O_3), \\ &= \frac{\sum_{k=3}^n \sum_{j=2}^k \gamma_{j,k}^{(n_a, n_b, n_c)} \mathbb{E}[T_j T_k]}{\sum_{k=3}^n \sum_{j=2}^k \kappa_{j,k} \mathbb{E}[T_j T_k]}, \end{aligned} \quad (8)$$

where

Figure 4 Statistical error in the frequency spectrum computation due to approximating first- and second-order moments of intercoalescence times, using simulations. For a constant population size, the plots show mean (base 10 log) KL divergence (over 10 repetitions) of the frequency spectrum approximated using simulations of intercoalescence times, from the true frequency spectrum computed by using the exact expected intercoalescence times. The diallelic frequency spectrum is shown in blue and the triallelic frequency spectrum is shown in red, for samples of size (A) 10 and (B) 100.

$$\begin{aligned} \gamma_{j,k}^{(n_a, n_b, n_c)} &= \left(P_{ab} P_{bc} C_{j,k}^{(n_a, n_b)} + P_{ab} P_{ac} F_{j,k}^{(n_a, n_b)} \right) \mathbb{I} \{k \leq n_a + n_b + 1\} \\ &\quad + \left(P_{ac} P_{cb} C_{j,k}^{(n_a, n_c)} + P_{ab} P_{ac} F_{j,k}^{(n_a, n_c)} \right) \mathbb{I} \{k \leq n_a + n_c + 1\}, \end{aligned}$$

$$\kappa_{j,k} = \left[\sum_{x \neq a} \sum_{y \neq a, x} P_{ax} P_{xy} \right] D_{j,k} + \left[\sum_{x \neq a} \sum_{y \neq a, x} P_{ax} P_{ay} \right] G_{j,k},$$

and $\mathbb{I}\{\cdot\}$ denotes the indicator function.

Proof. As in Jenkins and Song (2011, Theorem 6.2), this follows from

$$\begin{aligned} \mathbb{P}(\mathbf{n} | O_3) &= \frac{\mathbb{P}(\mathbf{n}, O_3, E_2)}{\mathbb{P}(O_3, E_2)} + O(\theta), \\ &= \frac{\mathbb{P}(\mathbf{n}, E_{2\mathcal{N}}^{(b,c)}) + \mathbb{P}(\mathbf{n}, E_{2\mathcal{N}\mathcal{N}}^{(c,b)}) + \mathbb{P}(\mathbf{n}, E_{2\mathcal{N}\mathcal{N}}^{(b,c)}) + \mathbb{P}(\mathbf{n}, E_{2\mathcal{N}\mathcal{N}}^{(c,b)})}{\left[\sum_{x \neq a} \sum_{y \neq a, x} \mathbb{P}(E_{2\mathcal{N}}^{(x,y)}) \right] + \left[\sum_{x \neq a} \sum_{y \neq a, x} \mathbb{P}(E_{2\mathcal{N}\mathcal{N}}^{(x,y)}) \right]} \\ &\quad + O(\theta). \end{aligned}$$

Now substitute for each term on the right-hand side using Lemma 1 and let $\theta \rightarrow 0$. $\square$

Thus, while the frequency spectrum for a site experiencing one mutation event depends only on the first moments of the elements of $\mathbf{T}$ [see (2)], the frequency spectrum for a site experiencing two mutation events depends only on the second moments of the elements of $\mathbf{T}$ (Theorem 1). These moments are considered in further detail by Polanski *et al.* (2003) and Živković and Wiehe (2008). Under a suitable choice of historical population size function, N_t , the frequency spectrum given by (8) serves as our null model for triallelic sites. As a check on (8), we can fix the population size, $N_t \equiv N_0$, so that

$$\mathbb{E}[T_j T_k] = (1 + \delta_{j,k}) \binom{j}{2}^{-1} \binom{k}{2}^{-1}.$$

Inserting this expression into (8) leads to (3), after extensive simplification.

While this article was under review we learned of related work by Sargsyan (2006), who also obtains an expression for the frequency spectrum of a site experiencing two mutation events under an arbitrary distribution on $\mathbf{T}$ (Sargsyan 2006, Lemma 34). Our work strengthens his result, which relies on higher-order and exponential moments of the elements of $\mathbf{T}$. Our work also allows for a more general model of mutation and disentangles the relative contributions of nested and nonnested mutations.

Application I: Sensitivity to Demography

In this section, we compare the frequency spectrum of a diallelic site with that of a triallelic site. Given a sample taken from a population whose recent history is described by a demographic model $\mathcal{M}_1$, we can measure the information that is lost if one erroneously applies the frequency spectrum according to another model $\mathcal{M}_0$. To quantify this difference in information, we employ the Kullback–Leibler (KL) divergence, a measure defined for this task (Kullback and Leibler 1951; Burnham and Anderson 2002). We define the KL divergence from $\mathcal{M}_1$ to $\mathcal{M}_0$ by

$$D(\mathcal{M}_1 \parallel \mathcal{M}_0) = \mathbb{E}_{\mathcal{M}_1} \left(\log \frac{\phi_{\mathcal{M}_1}(\mathbf{n})}{\phi_{\mathcal{M}_0}(\mathbf{n})} \right),$$

where $\phi_{\mathcal{M}_i}$ is the appropriate sampling distribution under model $\mathcal{M}_i$ and $\mathbb{E}_{\mathcal{M}}$ denotes expectation with respect to random samples $\mathbf{n}$ drawn under model $\mathcal{M}$. Thus, KL divergence is the expected likelihood ratio when testing an alternative $\mathcal{M}_1$ against a null $\mathcal{M}_0$ and the alternative is true. Although KL divergence properly refers to distributions under these models rather than to the models themselves, when we refer to the divergence between two models, it should be clear from the context that we are referring to either their diallelic or their triallelic sample frequency spectrum.

We focus on the divergence from one model of population growth to another. The KL divergence (amount of information

Figure 5 Seven models of exponential population growth examined in our analysis. Models 1–7 correspond to growth curves with the onset of growth occurring from $t_1 = 100$ generations ago to $t_1 = 700$ generations ago, respectively. The growth rate per generation, r_1 , is chosen so that $N_0 = 10^{6.5}$ in each model.

loss) can thus be compared for samples from a diallelic site vs. samples from a triallelic site. A larger value of KL divergence for triallelic sites would suggest that such sites are potentially very informative for demographic inference. Throughout this section, a symmetric mutation matrix is used in the frequency spectra calculations (*i.e.*, it is assumed that all transitions between alleles are equally likely). To illustrate how divergences vary at different scales, we focus our analysis on frequency spectra for samples of 10 and 100 individuals [a typical magnitude of sample sizes in demographic inference studies (Williamson *et al.* 2005; Gutenkunst *et al.* 2009)].

The effect of sample size in simulations

To compute frequency spectra under general models of historical population size we need first- and second-order moments of $\mathbf{T}$ (*cf.* Equation 2 and Theorem 1). We precompute these by simulating coalescent trees, using *ms* (Hudson 2002). To investigate the effect of sample size in simulations and of the differing dimensions of the two frequency spectra, we first consider the case of a constant population size (*i.e.*, $N_t \equiv N_0$), for which the expected frequency spectra are known in closed form (see Equations 1 and 3). Specifically, we compute the KL divergence from the expected frequency spectrum computed exactly to the expected frequency spectrum obtained by simulation of N_{trees} to approximate the first- and second-order moments of $\mathbf{T}$. The results are shown in Figure 4. Although it might be considered unfair to compare KL divergences in diallelic frequency spectra (one-dimensional distributions) with those between triallelic spectra (two-dimensional distributions), Figure 4 clearly shows that these divergences exhibit extremely similar behavior as we increase N_{trees} . The degree to which the true spectrum is approximated by its Monte Carlo counterpart is almost exactly the same in the di- and triallelic cases, regardless of the choice of N_{trees} . We also note that for both di- and triallelic spectra, the KL divergences have similar magnitude between sample sizes 10 and 100 for all choices of

Figure 6 KL divergence from one model of exponential population growth to another. A and B depict the (base 10 log) KL divergence $D(\mathcal{G}_i \parallel \mathcal{G}_0)$ of the sample frequency spectrum taken from a population under growth model $\mathcal{G}_i$ from the sample frequency spectrum taken from a population of fixed size ($\mathcal{G}_0$). C and D depict the (base 10 log) KL divergence $D(\mathcal{G}_i \parallel \mathcal{G}_{i-1})$ of the sample frequency spectrum taken from a population under growth model $\mathcal{G}_i$ from the sample frequency spectrum taken from a population under growth model $\mathcal{G}_{i-1}$. Growth models $\mathcal{G}_i$, $i = 1, \dots, 7$ are defined in Figure 5; results are shown for the diallelic (blue) and the triallelic (red) sample frequency spectrum.

N_{trees} . Thus, KL divergence appears to be a good measure of the difference between two frequency spectra that is relatively invariant to the differences in dimensionality and sample size in our study. Because Figure 4 illustrates that using $N_{\text{trees}} = 10^6$ results in negligibly small divergences (on the order of 10^{-8}) from the true (closed-form) frequency spectra, we fix $N_{\text{trees}} = 10^6$ in the remainder of this section.

Exponential growth

Next, we examine how sampling from a population with historical exponential growth affects the resulting frequency spectrum. As specified in Figure 5, we investigate seven models of exponential growth, $\mathcal{G}_i$, $i = 1, \dots, 7$ (with $\mathcal{G}_0$ representing a population of fixed size). To compute their respective moments of intercoalescence times, we first simulate 10^6 trees from populations according to each model. In Figure 6, A and B, we compute the KL divergence $D(\mathcal{G}_i \parallel \mathcal{G}_0)$ of the sample frequency spectrum under $\mathcal{G}_0$ from the sample frequency spectrum under $\mathcal{G}_i$, for $i = 1, \dots, 7$. To investigate the potential benefit of triallelic spectra in fine-tuning between two population growth models with different degrees of exponential growth, we also compute KL divergences $D(\mathcal{G}_i \parallel \mathcal{G}_{i-1})$ of sample frequency spectra under growth model $\mathcal{G}_{i-1}$ from growth model $\mathcal{G}_i$, for $i = 1, \dots, 7$ (Figure 6, C and D).

Figure 6 demonstrates clear superiority of using triallelic spectra to distinguish between demographic models with varying degrees of exponential population growth. The mean KL

divergence from exponential growth model i to $i - 1$ is increased by 87% when triallelic spectra are used in place of diallelic spectra for samples of size 10 (and the divergence is increased by 99% for sample size 100). This indicates that triallelic sites contain information that may significantly increase our ability to discern between competing exponential growth models with similar parameters.

Instantaneous growth

We next investigate in further detail the effect of sample size on KL divergence from a growth model to a model of fixed population size. For the growth model we assumed a function of historical human population growth as inferred by Williamson *et al.* (2005), who assumed an instantaneous expansion of the population from an ancestral size νN_0 to a modern size N_0 a time τ ago. Using data from the Environmental Genome Project and working within the framework of the Poisson random field model (Sawyer and Hartl 1992), Williamson *et al.* (2005) inferred maximum-likelihood estimates (MLEs) of $\nu = 0.160$ and $\hat{\tau} = 0.00885$, the latter in units of $2N_0$ generations. (The authors estimated $N_0 \approx 51,340$ directly by comparing polymorphism and divergence data, in which case the latter estimate is calibrated as $\hat{\tau} = 908$ generations or, further assuming a 20-year generation time, $\hat{\tau} \approx 18,200$ years.) We denote this model by $\mathcal{G}_W$. Thus, given samples that actually stem from a population such as the one described by Williamson *et al.* (2005), the KL divergence $D(\mathcal{G}_W \parallel \mathcal{G}_0)$ quantifies the ability to distinguish that these samples do not come from

a fixed-size population. This analysis therefore studies the effect of sample size on the ability to perform inference of population growth parameters under realistic settings in a problem of great interest, for both di- and triallelic sites.

Figure 7 illustrates that the KL divergence using triallelic spectra for the two models is much greater, for any sample size, than the divergence using the corresponding diallelic spectra. Furthermore, we see that the increase in KL divergence that results from the presence of the third allele at a triallelic site also grows with increasing sample size, at least up to sample sizes of 25, before leveling off beyond 25 and providing a consistent $\sim 97\%$ increase in KL divergence.

We further address the effect of the parameters of a model of instantaneous population size change, as follows. Adopting $\mathcal{G}_W$ as a reference point, we examine variations in the two parameters ν and τ . First, the amount of instantaneous growth is varied while keeping the time of the size change fixed to same value as $\mathcal{G}_W$ (see models 8–20 in Table 1), and subsequently, different times of size-change occurrence are examined while the amount of instantaneous growth is fixed so that the prechange population size is 15% of the postchange size (models 21–29 in Table 1).

Repeating the steps of our analysis of the exponential growth models, we computed KL divergences $D(\mathcal{G}_i \parallel \mathcal{G}_0)$ for $i = 8, \dots, 29$ (Figure 8, A and B, and Figure 9, A and B), and to investigate the potential benefit of triallelic spectra in the more subtle problem of distinguishing between two instantaneous population growth models with different degrees of growth, we computed $D(\mathcal{G}_i \parallel \mathcal{G}_{i-1})$ for each $i = 9, \dots, 20$ and $i = 22, \dots, 29$ (Figure 8, C and D, and Figure 9, C and D).

From Figure 8 and Figure 9, we again find that KL divergence from the spectrum under one growth model to another is larger when we use a triallelic, rather than diallelic, sample frequency spectrum, although this advantage fades as the time of the sudden size change is moved extremely far into the past while the amount of instantaneous growth is kept constant (as in models 21 and 22 in Figure 9). The advantage of the triallelic sample frequency spectrum is even more pronounced for sample size 100 than for sample size 10. The mean KL divergence $D(\mathcal{G}_i \parallel \mathcal{G}_{i-1})$ (over $i = 8, \dots, 20$; see Figure 8, C and D) is increased by 79% when triallelic spectra are used in place of diallelic spectra for samples of size 10 (and this divergence is increased by 97% for sample size 100). For models 21–29 (Figure 9, C and D, the mean KL divergence $D(\mathcal{G}_i \parallel \mathcal{G}_{i-1})$ is increased by 56% through the inclusion of the third allele for samples of size 10 (and the divergence is increased by 83% for sample size 100). Thus, the inclusion of the third allele in triallelic spectra contains information that may considerably increase our power to discern between competing instantaneous growth models with similar parameters.

Application II: Simultaneous Mutation Model

Theory

As discussed above, Hodgkinson and Eyre-Walker (2010) propose that there exists another mechanism of mutation respon-

Figure 7 KL divergence $D(\mathcal{G}_W \parallel \mathcal{G}_0)$ from a population model with the growth function of Williamson *et al.* (2005) to a population model of fixed size, for various sample sizes and using di- (blue) and tri- (red) allelic frequency spectra. The percentages between the curves denote the percentage of increase in KL divergence that results from computing divergence using triallelic spectra rather than diallelic spectra.

sible for the observed excess of triallelic sites in samples of human genomes: the simultaneous generation of two new alleles within a single individual. It is estimated that this mechanism is responsible for the generation of $\sim 3\%$ of all human SNPs. In support of this hypothesis they developed a phylogenetic statistic to test whether the two minor alleles of a triallelic site are closer to each other on a reconstructed phylogenetic tree than would be expected by chance. Using this test, they find significant evidence at the 5% level for proximity of the minor alleles when probabilities are combined across all triallelic sites in their data, although the null hypothesis of two independent mutation events is rejected at a rate close to the nominal 5% when each of 113 triallelic sites is tested independently. There are, however, some limitations to this test. First, as Hodgkinson and Eyre-Walker (2010) observe, phylogenies are reconstructed using local haplotype information but ignoring the confounding effects of recombination. Second, it uses the mean branch length between leaf nodes subtended by minor alleles as an indirect measure of the branch length between the minor alleles themselves. Finally, it uses the *proximity* of two mutation events on the phylogeny as evidence for what is in fact a stricter hypothesis of *colocation*. In this section we use our results on the triallelic frequency spectrum to develop a complementary test that does not suffer from these issues. Our approach is to find the sample frequency spectrum as predicted by the simultaneous mutation mechanism and then to compare it with the results of Theorem 1 via a likelihood-ratio test. It will be clear that the two spectra give very different predictions, particularly with regard to the expected number of singleton alleles in a sample.

The key observation that enables us to obtain the frequency spectrum under this model is as follows. Suppose there exists a mechanism whereby two new alleles are produced within a single individual, such as that described in the Introduction:

Table 1 Twenty-two models of instantaneous population size increase from νN_0 to N_0 at a fixed historical time point τ

Model	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29
ν	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.25	0.2	0.15	0.1	0.05	0.01									
τ														1.0	0.5	0.1	0.05	0.01	0.005	0.001	0.0005	0.0001

In models 8–20 the time τ is fixed at the same value as in the human population growth model proposed by Williamson *et al.* (2005) ($\tau = 0.0044$ in units of $4N_0$ generations), while in models 21–29 the magnitude of growth is fixed at $\nu = 0.15$.

A single DNA duplex within a diploid cell experiences subsequent mutations of both of the nucleotides within a base pair. The duplex then undergoes replication so that two new alleles are produced (Figure 1). Now, the two alleles are observed in a sample taken from the population in the present day. The individual responsible for the creation of these alleles must have been an ancestor of individuals in the sample carrying *either* of the derived alleles. Moreover, this ancestor was the *most recent common ancestor* of any pair of individuals carrying the two distinct derived alleles. The genealogy relating the sample at this site must be of the form illustrated in Figure 10; in particular, the simultaneous mutation event coincides with the coalescence node uniting the two clades defined by individuals carrying the two derived alleles. We can therefore condition on this coincidence event in deriving the sample frequency spectrum under this model, by choosing uniformly among the $n - 1$ coalescence nodes. As noted by Hodgkinson and Eyre-Walker (2010), the probability that both products of a single human meiosis leave descendants in the following generation is negligible, so the posited simultaneous mutation event is presumed to occur during the mitotic phase of germline development.

In its full generality, the mutations under this model are parameterized by a $K \times \binom{K}{2}$ transition matrix $\mathbf{Q} = (Q_{i,\{j,k\}})$ whose $(i,\{j,k\})$ th entry specifies the probability that a simultaneous mutation affecting allele i gives rise to alleles j and k . For notational simplicity we assume $Q_{i,\{j,k\}} = 0$ if i, j, k are not all distinct—it is straightforward to make the appropriate modifications to relax such an assumption. In this setting we have the following theorem.

Theorem 2. Let $n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c$ denote a triallelic sample, and let E'_s denote the event that there were s instances of the mechanism of simultaneous mutation in the genealogical history relating the sample. Then the sample frequency spectrum is

$$\begin{aligned} \phi_S(n_a, n_b, n_c) &= \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c | O_3, E'_1), \\ &= Q_{a,\{b,c\}} \cdot \frac{2n}{n-2} \cdot \frac{1}{(n-n_a-1)(n-n_a)(n-n_a+1)}. \end{aligned} \quad (9)$$

Proof. See the Appendix. $\square$

We remark that in the above we conditioned on E'_1 ; equivalently one could introduce a rate parameter for the occurrences of simultaneous mutations and then let it go to zero after conditioning on O_3 only, in which case none of $E'_2, E'_3, \dots$ contributes to the frequency spectrum.

Importantly, the frequency spectrum in (9) depends on the distribution of topologies of coalescent trees but not on the

distribution of $\mathbf{T}$. Thus, Theorem 2 continues to apply when we allow a general distribution of intercoalescence times as in the null model, and in particular this includes a model of variable population size, N_t . To summarize: Under a model in which the population size N_t is allowed to vary in time, the sample frequency spectrum of a triallelic site is given by (8) when the two mutation events occur independently and by (9) when they occur simultaneously within a single individual. An example of the two spectra is shown in Figure 11. Clearly, the largest difference occurs in the frequency class of double singletons, $(n_a, n_b, n_c) = (n - 2, 1, 1)$, which contributes 0.37 of the total probability mass under the simultaneous mutation mechanism compared with 0.09 when the two mutations occur independently along the tree. Other nearby configurations in which both derived alleles are at very low frequency are also over-represented according to the simultaneous mutation mechanism by comparison with independent mutations, while configurations in which one or both derived alleles are at moderate frequency are slightly underrepresented. This underrepresentation is greatest for frequencies of the form $(n_a, 1, n_c)$ and $(n_a, n_b, 1)$, *i.e.*, along the axes in Figure 11.

Likelihood-ratio test of independent vs. simultaneous mutation

Our goal now is to test for a relative excess of triallelic sites that conform to the frequency spectrum of the simultaneous mutation mechanism. We take as our null hypothesis that each triallelic site was generated by two independent mutation events, so that the frequency spectrum is given by ϕ_0 (Equation 8). An appropriate alternative is that some fraction, $\lambda > 0$, of triallelic sites arose as the result of a simultaneous mutation event. Under this model, the sample frequency spectrum is given by the mixture

$$\phi_\lambda(n_a, n_b, n_c) = (1 - \lambda)\phi_0(n_a, n_b, n_c) + \lambda\phi_S(n_a, n_b, n_c). \quad (10)$$

Suppose we observe M triallelic sites, and the i th site has configuration $(n_{a_i}^{(i)}, n_{b_i}^{(i)}, n_{c_i}^{(i)})$. If each pair of sites is sufficiently far apart that their genealogical histories are independent, then a likelihood-ratio statistic for these data is

$$\Lambda = \frac{\prod_{i=1}^M \phi_\lambda(n_{a_i}^{(i)}, n_{b_i}^{(i)}, n_{c_i}^{(i)})}{\prod_{i=1}^M \phi_0(n_{a_i}^{(i)}, n_{b_i}^{(i)}, n_{c_i}^{(i)})}, \quad (11)$$

where $\hat{\lambda}$ is a MLE for λ , with $0 \leq \hat{\lambda} \leq 1$. We reject the null hypothesis at level α if Λ lies within the $100(1 - \alpha)$ th percentile tail of its null distribution. Unfortunately, the null

Figure 8 KL divergence from one model of instantaneous population growth to another. A and B depict (base 10 log) KL divergence $D(\mathcal{G}_i \parallel \mathcal{G}_0)$ of the sample frequency spectrum taken from a population under growth model $\mathcal{G}_i$ from the sample frequency spectrum taken from a population of fixed size ($\mathcal{G}_0$), for each $i = 8, \dots, 20$. C and D depict (base 10 log) KL divergence $D(\mathcal{G}_{i-1} \parallel \mathcal{G}_i)$ of the sample frequency spectrum taken from a population under growth model $\mathcal{G}_i$ from the sample frequency spectrum taken from a population under growth model $\mathcal{G}_{i-1}$, for $i = 8, \dots, 20$ (where we define the reference growth model before model 8 to simply be a demography with fixed-population size). Results are shown for the diallelic (blue) and the triallelic (red) sample frequency spectrum.

distribution of $-2 \ln \Lambda$ does not tend toward the usual χ^2_1 -distribution due to the possibility that the mixture parameter λ lies on the boundary of its permissible set (Self and Liang 1987). We thus employ bootstrap estimation to determine the null distribution of the test statistic by simulation (further described in supporting information, File S1).

Our aim is to apply this test to empirical SNP data, so we take $K = 4$ with alleles $\{A, C, G, T\}$. However, to apply the test we must also specify both transition matrices $\mathbf{P}$ and $\mathbf{Q}$. While $\mathbf{P}$ can be estimated by, for example, using the empirical frequencies of each type of mutation event inferred from diallelic SNPs and their corresponding outgroup alleles (Chan *et al.* 2012), there is no guidance on how to choose $\mathbf{Q}$. In fact, we do not expect the test to be greatly influenced by the *identities* of the three alleles at a site compared to the information contained in the sample counts themselves. Therefore we choose to eliminate the appearance of the entries of $\mathbf{Q}$ in Λ by *conditioning* on the identities of the observed alleles at each site. Formally, we replace O_3 in the definitions above with $O_3^{(a,b,c)}$, the event that three alleles are observed and these alleles are a, b , and c , with a ancestral. This leads to the slightly modified test statistic $\tilde{\Lambda}$, obtained by replacing the appearances of ϕ_S and ϕ_0 in each of (10) and (11) with

$$\tilde{\phi}_S(n_a, n_b, n_c) = \mathbb{P}(n = n_a e_a + n_b e_b + n_c e_c | O_3^{(a,b,c)}, E'_1),$$

$$\tilde{\phi}_0(n_a, n_b, n_c) = \lim_{\theta \rightarrow 0} \mathbb{P}(n = n_a e_a + n_b e_b + n_c e_c | O_3^{(a,b,c)}).$$

By repeating the reasoning that led to expressions for ϕ_S and ϕ_0 , we obtain $\tilde{\phi}_S$ from ϕ_S by dividing by $Q_{a,\{b,c\}}$ in (9) and $\tilde{\phi}_0$ from ϕ_0 by replacing $\kappa_{j,k}$ in (8) with

$$\tilde{\kappa}_{j,k} = [P_{ab}P_{bc} + P_{ac}P_{cb}]D_{j,k} + [2P_{ab}P_{ac}]G_{j,k}.$$

The resulting test statistic $\tilde{\Lambda}$ is independent of $\mathbf{Q}$ and is more robust than Λ to the choice of $\mathbf{P}$.

Data

To apply our test to SNP data we first need to specify $\mathbf{P}$, and to do this we followed the procedure described in Chan *et al.* (2012). Their method requires empirical counts of each type of diallelic SNP with the ancestral allele at each SNP specified, as well as the overall abundance of each nucleotide in the genome. To obtain empirical diallelic SNP data we used the Genome Variation Server (v6.01) (GVS 2011). We obtained each diallelic SNP from the GVS database, discarding those for which the orthologous chimpanzee allele was unavailable or did not match any of the human alleles. To restrict our attention to neutral mutation events occurring at a typical genomic rate, we further discarded SNPs at which one of the alleles would produce a CpG dinucleotide; we discarded SNPs residing in coding regions; and to keep our estimate of $\mathbf{P}$ independent of the test data set we discarded sites at which more than two alleles were observed. Assuming that the chimp carries the ancestral allele at each site polymorphic in humans and that each remaining SNP represents a single mutation event from the ancestral to the

Figure 9 (A–D) KL divergence from one model of instantaneous population growth to another, as in Figure 8, but this time for models $i = 21, \dots, 29$.

derived allele, we were left with the following counts of each type of mutation event:

$$\begin{array}{lll}
 N_{A \rightarrow T} = 2299 & N_{A \rightarrow C} = 1886 & N_{A \rightarrow G} = 7226 \\
 N_{T \rightarrow A} = 2238 & N_{T \rightarrow C} = 6956 & N_{T \rightarrow G} = 1960 \\
 N_{C \rightarrow A} = 2819 & N_{C \rightarrow T} = 9940 & N_{C \rightarrow G} = 2395 \\
 N_{G \rightarrow A} = 9931 & N_{G \rightarrow T} = 2870 & N_{G \rightarrow C} = 2394.
 \end{array}$$

These counts can be converted to rates of mutation by comparison with the overall genomic abundance of each type of nucleotide, which were obtained from the University of California, Santa Cruz Genome Browser (hg19) (Kent *et al.* 2002), after excluding both CpG dinucleotides and the Y chromosome and weighting the counts for the X chromosome by 3/4. This left the following counts of each type of nucleotide in the human genome: $N_A = 835,878,173$ (30.1%), $N_T = 836,874,687$ (30.0%), $N_C = 552,795,868$ (19.9%), and $N_G = 553,090,147$ (19.9%). Using these values together with the mutation counts given above, we obtained the following empirical $\mathbf{P}$ matrix, using the method described in Chan *et al.* (2012):

$$\mathbf{P} = \begin{array}{c} \begin{array}{cccc} & \text{A} & \text{T} & \text{C} & \text{G} \\ \text{A} & 0.503 & 0.100 & 0.082 & 0.315 \\ \text{T} & 0.097 & 0.515 & 0.303 & 0.085 \\ \text{C} & 0.186 & 0.655 & 0.002 & 0.158 \\ \text{G} & 0.654 & 0.189 & 0.158 & 0 \end{array} \end{array}.$$

We reanalyzed the triallelic data set of Hodgkinson and Eyre-Walker (2010, their Table S1), which comprised 113

triallelic sites. These were in turn obtained from 896 nuclear genes sequenced as part of the Environmental Genome Project (NIEHS SNPs 2011) and the SeattleSNPs project (SeattleSNPs 2011), which provide high-quality resequencing data, avoiding problems such as ascertainment bias. Only sites of high quality ($Q > 25$), outside CpG dinucleotides, and outside coding regions were included in the data. Orthologous chimpanzee alleles corresponding to each site in the data were kindly provided to us by Alan Hodgkinson, CHU Sainte-Justine, Montréal, Canada; sites for which the chimp allele was unavailable were excluded from our analysis, leaving $M = 96$ triallelic sites. Since these sites originate from different experiments using different population panels, the sample size varied across sites. The minimum, mean, and maximum sample sizes across the 96 triallelic sites were 71, 160, and 190, respectively.

To compute Λ , the required formulas are given in terms of the joint moments $\mathbb{E}[T_j T_k]$, and as before we precompute these numerically. This precomputation step can be reused for different choices of λ and across segregating sites with the same sample size, and so it does not add to the computational burden significantly.

The most striking feature of the historical human population size is its recent rapid growth (Keinan and Clark 2012). To examine this, we used the model $\mathcal{G}_W$ of Williamson *et al.* (2005) as described above. Computing the MLE $\hat{\lambda}$ under this demographic model yielded $\hat{\lambda} = 0$, conflicting with the conclusions of Hodgkinson and Eyre-Walker

Figure 10 A coalescent tree with one simultaneous mutation event. The allele of each leaf is annotated. Also annotated is the variable l (here, $l = 5$), the number of lineages ancestral to the sample just prior to the simultaneous mutation event.

(2010). This illustrates the importance of accounting for demographic changes when using frequency spectrum data; had we assumed a population of constant size we would find $\hat{\lambda} = 0.21$ ($P < 0.001$; see [File S1](#), in which we also examine the robustness of this result to assumptions about $\mathbf{P}$ and to potential sequencing errors), in better agreement with Hodgkinson and Eyre-Walker's (2010) estimate of $\lambda \approx 0.5$. The decrease in the value of our estimated mixture parameter can be explained by the fact that singleton sites—and, similarly, doubly singleton triallelic sites—are relatively more probable under a null model with population growth than one without. Thus, the abundance of doubly singleton sites that previously gave rise to an extreme likelihood-ratio statistic is now explicable under the null model.

A further potential complication of the data used here is population subdivision with migration between subpopulations. Analytic results for the frequency spectrum under complex demography are unavailable, even ignoring the issue of recurrent mutations (but for recent progress on this problem see [Chen 2012](#)), and so in [File S1](#) we further investigate this issue by a simulation study. We find $\hat{\lambda} = 0.16$, although the lack of availability of subpopulation labels with the data renders the result nonsignificant ($P \approx 0.42$; [File S1](#)).

Discussion

In this article we have obtained closed-form expressions for the frequency spectrum of a site experiencing two mutation events, of which triallelic sites are an important example, and have allowed for the possibility of a varying historical population size; the results generalize those of Griffiths and Tavaré (1998) and Jenkins and Song (2011). We applied our formulas to the question of the ability of the frequency spectrum to discern between closely related models of population growth and to the question of the mechanism of mutation that gives rise to triallelic sites.

Demographic inference from SNP data has thus far relied solely on diallelic sites. As sample sizes in sequencing studies grow with the falling cost of the technology, it is likely that an ever-increasing fraction of segregating sites found will be triallelic. In this article, we have illustrated that the triallelic sampling frequency spectrum is more sensitive to historical population size changes under both exponential and instantaneous growth models, and thus it seems likely that improved estimates of population growth parameters may be obtained by incorporating this growing number of triallelic sites in demographic inference analyses. Furthermore, the increased sensitivity of the triallelic spectrum over the diallelic spectrum to discern between different growth models becomes more exaggerated as sample sizes are increased. While triallelic sites remain relatively rare compared to diallelic ones, the above analysis suggests that they are more valuable per site in distinguishing between a variety of demographic models. We hope that the ability of triallelic sites to fine-tune between competing growth models will be especially useful when looking at recent *superexponential* growth (Keinan and Clark 2012), although convenient software for this type of growth is not yet available.

We also found the frequency spectrum under a model in which the two derived alleles of a triallelic site may be generated simultaneously (Hodgkinson and Eyre-Walker 2010), and we have developed a likelihood-ratio test for the existence of such a mechanism. This approach is parameterized by the mixture parameter λ that represents the fraction of triallelic sites having arisen as a result of the simultaneous mutation mechanism. Assuming a simple randomly mating

Figure 11 The triallelic sample frequency spectrum when mutations occur (A) independently and (B) simultaneously, for a sample of size $n = 20$. In this example there are $K = 3$ alleles a , b , and c , with uninformative mutation matrices $P_{ab} = P_{ac} = P_{bc} = P_{cb} = \frac{1}{2}$, and $Q_{a,\{b,c\}} = 1$. The population size is constant in time.

population of constant size, we find a MLE of $\hat{\lambda} = 0.21$, which is significantly nonzero. We show that another explanation for the excess of doubly singleton triallelic sites is rapid recent population growth, but when we posit a realistic demographic model that includes population subdivision and migration as well as recent growth, we find only a minor adjustment to $\hat{\lambda} = 0.16$, supporting the idea that at least some triallelic sites were generated as the result of a simultaneous mutation event. This latter estimate is not, however, significantly different from 0, a state of affairs we can at least partly attribute to a considerable loss of power—the individuals sampled at a substantial fraction of triallelic sites in our data set are lacking subpopulation labels. [In particular, none of the triallelic sites of the form $(n^{(i)} - 2, 1, 1)$ remained triallelic after removing from the sample individuals of unknown origin.]

Because of the lack of power associated with the available data, we treat these results with caution and do not rule out the possibility that at least some sites were generated by a mechanism of simultaneous mutation. Indeed, most of the information about a mechanism generating an excessive number of triallelic sites is contained in the absolute number of triallelic sites observed in the human genome. We did not use this quantity; instead we *conditioned* on the observed number M of triallelic sites and addressed a slightly different question: Given the excessive number of triallelic sites in the genome, is the frequency spectrum of some fraction of these sites consistent with the two derived alleles being generated simultaneously within a single individual? Even if this proposition is rejected, the question of why there is such an excess of triallelic sites remains. We are hopeful that the coming flood of (subpopulation-labeled) genomic data will enable us to address these questions with much improved power.

One lesson of our work is that rare alleles are as vital when looking at triallelic sites as when looking at diallelic ones; singleton alleles at *diallelic* sites are already the focus of other coalescent-based tests of neutrality (Fu and Li 1993; Achaz 2009). While earlier genotyping projects often chose to exclude sites below a given minor allele frequency, typically 5%, more recent trends—improved sequencing technologies, larger sample sizes, and an interest in rare variants in their own right (Cirulli and Goldstein 2010; Coventry *et al.* 2010; Keinan and Clark 2012; Nelson *et al.* 2012; Tennessen *et al.* 2012)—should serve to make the appropriate data more readily available. As this trend continues, we expect our results to find further use as recurrent mutations manifest themselves more and more commonly.

Acknowledgments

We thank Christoph Theunert for helpful discussions and Alan Hodgkinson for providing ancestral allele information to accompany the triallelic site data. We also thank John Wakeley for helpful suggestions for improving the exposition of this article. This research is supported in part by National Institutes of Health grant R01-GM094402, an Alfred P. Sloan Research Fellowship, and a Packard Fellowship for Science and Engineering.

Literature Cited

- Achaz, G., 2009 Frequency spectrum neutrality tests: one for all and all for one. *Genetics* 183: 249–258.
- Bhaskar, A., J. A. Kamm, and Y. S. Song, 2012 Approximate sampling formulae for general finite-alleles models of mutation. *Adv. Appl. Probab.* 44: 408–428.
- Boyko, A., S. H. Williamson, A. R. Indap, J. D. Degenhardt, R. D. Hernandez *et al.*, 2008 Assessing the evolutionary impact of amino acid mutations in the human genome. *PLoS Genet.* 4(5): e1000083.
- Burnham, K. P., and D. R. Anderson, 2002 *Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach*, Ed. 2. Springer-Verlag, Berlin/Heidelberg, Germany/New York.
- Bustamante, C. D., J. Wakeley, S. Sawyer, and D. L. Hartl, 2001 Directional selection and the site-frequency spectrum. *Genetics* 159: 1779–1788.
- Campbell, C. D., J. X. Chong, M. Malig, A. Ko, B. L. Dumont *et al.*, 2012 Estimating the human mutation rate using autozygosity in a founder population. *Nat. Genet.* 44(11): 1277–1281.
- Chan, A. H., P. A. Jenkins, and Y. S. Song, 2012 Genome-wide fine-scale recombination rate variation in *Drosophila melanogaster*. *PLoS Genet.* 8(12): e1003090.
- Chen, H., 2012 The joint allele frequency spectrum of multiple populations: a coalescent approach. *Theor. Popul. Biol.* 81: 179–195.
- Cirulli, E. T., and D. B. Goldstein, 2010 Uncovering the roles of rare variants in common disease through whole-genome sequencing. *Nat. Rev. Genet.* 11: 415–425.
- Coventry, A., L. M. Bull-Otterson, X. Liu, A. G. Clark, T. J. Maxwell *et al.*, 2010 Deep resequencing reveals excess rare recent variants consistent with explosive population growth. *Nat. Commun.* 1: 131.
- Crawford, D. C., C. L. Sanders, X. Qin, J. D. Smith, C. Shephard *et al.*, 2006 Genetic variation is associated with C-reactive protein levels in the Third National Health and Nutrition Examination Survey. *Circulation* 114: 2458–2465.
- Desai, M. M., and J. B. Plotkin, 2008 The polymorphism frequency spectrum of finitely many sites under selection. *Genetics* 180: 2175–2191.
- Evans, S. N., Y. Shvets, and M. Slatkin, 2007 Non-equilibrium theory of the allele frequency spectrum. *Theor. Popul. Biol.* 71: 109–119.
- Fu, Y.-X., 1995 Statistical properties of segregating sites. *Theor. Popul. Biol.* 48: 172–197.
- Fu, Y.-X., and W.-H. Li, 1993 Statistical tests of neutrality of mutations. *Genetics* 133: 693–709.
- Griffiths, R. C., 2003 The frequency spectrum of a mutation, and its age, in a general diffusion model. *Theor. Popul. Biol.* 64: 241–251.
- Griffiths, R. C., and S. Tavaré, 1994 Sampling theory for neutral alleles in a varying environment. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 344: 403–410.
- Griffiths, R. C., and S. Tavaré, 1998 The age of a mutation in a general coalescent tree. *Stoch. Models* 14: 273–295.
- Gutenkunst, R. N., R. D. Hernandez, S. H. Williamson, and C. D. Bustamante, 2009 Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genet.* 5(10): e1000695.
- GVS, 2011 Genome Variation Server 134, August 2011. <http://gvs.gs.washington.edu>.
- Hobolth, A., and C. Wiuf, 2009 The genealogy, site frequency spectrum and ages of two nested mutant alleles. *Theor. Popul. Biol.* 75: 260–265.
- Hodgkinson, A., and A. Eyre-Walker, 2010 Human triallelic sites: evidence for a new mutational mechanism? *Genetics* 184: 233–241.

- Hudson, R. R., 2002 Generating samples under a Wright-Fisher neutral model of genetic variation. *Bioinformatics* 18: 337–338.
- Hüebner, C., I. Petermann, B. L. Browning, A. N. Shelling, and L. R. Ferguson, 2007 Triallelic single nucleotide polymorphisms and genotyping error in genetic epidemiology studies: *MDR1 (ABCB1) G2677/T/A* as an example. *Cancer Epidemiol. Biomarkers Prev.* 16(6): 1185–1192.
- Jenkins, P. A., and Y. S. Song, 2011 The effect of recurrent mutation on the frequency spectrum of a segregating site and the age of an allele. *Theor. Popul. Biol.* 80(2): 158–173.
- Keinan, A., and A. G. Clark, 2012 Recent explosive human population growth has resulted in an excess of rare genetic variants. *Science* 336(6082): 740–743.
- Kent, W. J., C. W. Sugnet, T. S. Furey, K. M. Roskin, T. H. Pringle *et al.*, 2002 The human genome browser at UCSC. *Genome Res.* 12(6): 996–1006. Available at: <http://genome.ucsc.edu> (hg19, August 2011).
- Kong, A., M. L. Frigge, G. Masson, S. Besenbacher, P. Sulem *et al.*, 2012 Rate of *de novo* mutations and the importance of father's age to disease risk. *Nature* 488(7412): 471–475.
- Kullback, S., and R. A. Leibler, 1951 On information and sufficiency. *Ann. Math. Stat.* 22: 79–86.
- Lukić, S., and J. Hey, 2013 Demographic inference using spectral methods on SNP data, with an analysis of the human out-of-Africa expansion. *Genetics* 192: 619–639.
- Lukić, S., J. Hey, and K. Chen, 2011 Non-equilibrium allele frequency spectra via spectral methods. *Theor. Popul. Biol.* 79: 203–219.
- Nelson, M. R., D. Wegmann, M. G. Ehm, D. Kessner, P. St. Jean *et al.*, 2012 An abundance of rare functional variants in 202 drug target genes sequenced in 14,002 people. *Science* 337(6090): 100–104.
- NIEHS SNPs, 2011 NIEHS Environmental Genome Project. University of Washington, Seattle. Available at: <http://egp.gs.washington.edu>. Accessed: August 2011.
- Polanski, A., and M. Kimmel, 2003 New explicit expressions for relative frequencies of single-nucleotide polymorphisms with application to statistical inference on population growth. *Genetics* 165: 427–436.
- Polanski, A., A. Bobrowski, and M. Kimmel, 2003 A note on distributions of times to coalescence, under time-dependent population size. *Theor. Popul. Biol.* 63: 33–40.
- Sargsyan, O., 2006 Analytical and simulation results for the general coalescent. Ph.D. Thesis, University of Southern California, Los Angeles.
- Sawyer, S. A., and D. L. Hartl, 1992 Population genetics of polymorphism and divergence. *Genetics* 132: 1161–1176.
- SeattleSNPs, 2011 SeattleSNPs. NHLBI Program for Genomic Applications, SeattleSNPs, Seattle. Available at: <http://pga.gs.washington.edu>. Accessed: August 2011.
- Self, S. G., and K.-Y. Liang, 1987 Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under non-standard conditions. *J. Am. Stat. Assoc.* 82(398): 605–610.
- Slatkin, M., and R. R. Hudson, 1991 Pairwise comparisons of mitochondrial DNA sequences in stable and exponentially growing populations. *Genetics* 129: 555–562.
- Song, Y. S., and M. Steinrücken, 2012 A simple method for finding explicit analytic transition densities of diffusion processes with general diploid selection. *Genetics* 190: 1117–1129.
- Steinrücken, M., Y. Wang, and Y. S. Song, 2013 An explicit transition density expansion for a multi-allelic Wright-Fisher diffusion with general diploid selection. *Theor. Popul. Biol.* 83: 1–14.
- Tennessen, J. A., A. W. Bigham, T. D. O'Connor, W. Fu, E. E. Kenny *et al.* 2012 Evolution and functional impact of rare coding variation from deep sequencing of human exomes. *Science* 337(6090): 64–69.
- Watterson, G. A., 1975 On the number of segregating sites in genetical models without recombination. *Theor. Popul. Biol.* 7: 256–276.
- Williamson, S. H., R. Hernandez, A. Fledel-Alon, L. Zhu, and C. D. Bustamante, 2005 Simultaneous inference of selection and population growth from patterns of variation in the human genome. *Proc. Natl. Acad. Sci. USA* 102(22): 7882–7887.
- Wiuf, C., and P. Donnelly, 1999 Conditional genealogies and the age of a neutral mutant. *Theor. Popul. Biol.* 56: 183–201.
- Wooding, S., and A. Rogers, 2002 The matrix coalescent and an application to human single-nucleotide polymorphisms. *Genetics* 161: 1641–1650.
- Živković, D., and T. Wiehe, 2008 Second-order moments of segregating sites under variable population size. *Genetics* 180: 341–357.

Communicating editor: J. Wakeley

Appendix

Proof of (2). The arguments given here mimic in part those found in Jenkins and Song (2011), and we refer the reader to that work for further details. In that article we work with *unordered* configurations, $\mathbf{n}$; here, the argument is easier to illustrate using *ordered* configurations. We denote a random vector of n alleles consistent with the unordered configuration $\mathbf{n}$ by $\mathbf{v}_n$; by sampling exchangeability there are $n!/\prod_{i=1}^K n_i!$ equiprobable such vectors. First, denote the event that a single mutation occurred in the genealogical history relating the sample and that it gave rise to a derived allele b by $E_1^{(b)}$. Now, inspection of Figure 2 tells us that any particular coalescent history with leaf configuration $\mathbf{v}_n$ and consistent with $E_1^{(b)}$ must exhibit the following sequence of events going back in time, for some $l_a \in \{1, 2, \dots, n_a\}$: a collection of $n_a - l_a$ coalescence events of type a lineages and a collection of $n_b - 1$ coalescence events of type b lineages, with the two collections in some interspersed ordering; a mutation event taking the single remaining type b lineage to a type a lineage; and coalescence of the remaining $l_a + 1$ type a lineages to a most recent common ancestor of the sample.

The quantity l_a represents the number of extant lineages ancestral to samples of type a at the time of the single mutation event. Suppose the first coalescence event is between two lineages whose allele at the leaves of the tree is b ; such a coalescence occurs (among all possible coalescence events) with probability

$$\frac{n_b(n_b - 1)}{n(n - 1)}.$$

Continuing in this vein back to the most recent common ancestor, we find that the probability of a given sequence $\mathcal{C}_{\mathbf{v}_n, l_a}$ of $n - 1$ alleles corresponding to each coalescence event, which is consistent with $\mathbf{v}_n$ and with $E_1^{(b)}$, satisfies

$$\mathbb{P}(\mathcal{C}_{\mathbf{v}_n, l_a}) = \frac{n_a!(n_a - 1)!n_b!(n_b - 1)!}{n!(n - 1)!} l_a(l_a + 1).$$

For example, in Figure 2, $\mathcal{C}_{\mathbf{v}_n, l_a} = (b, a, a, b, b, a, a, a)$. In fact, for a fixed l_a this probability is invariant with respect to permutations of the alleles in this sequence [i.e., with respect to the way the coalescence events are interspersed (Jenkins and Song 2011)]. There are $\binom{n - l_a - 1}{n_b - 1}$ ways to intersperse the first $n - l_a - 1$ coalescence events (Jenkins and Song 2011, Lemma 3.1).

Now, given $\mathcal{C}_{\mathbf{v}_n, l_a}$, we require the probability that a single mutation event occurs on the correct branch of the tree, an event we denote by $M_{\mathbf{v}_n}$. Since mutation events occur along the branches as a Poisson process of rate $\theta/2$, we have that

$$\mathbb{P}(M_{\mathbf{v}_n} | \mathcal{C}_{\mathbf{v}_n, l_a}) = \mathbb{E} \left[\frac{\theta}{2} L_n e^{-(\theta/2)L_n} \cdot \frac{T_{l_a+1}}{L_n} \right] = \frac{\theta}{2} \mathbb{E}[T_{l_a+1}] + O(\theta^2),$$

as $\theta \rightarrow 0$, where $L_n = \sum_{j=2}^n jT_j$ is the total branch length.

Putting all this together,

$$\begin{aligned} \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b, E_1^{(b)}) &= \binom{n}{n_b} \mathbb{P}(\mathbf{v}_n, E_1^{(b)}), \\ &= P_{ab} \binom{n}{n_b} \sum_{l_a=1}^{n_a} \binom{n - l_a - 1}{n_b - 1} \mathbb{P}(\mathcal{C}_{\mathbf{v}_n, l_a}) \mathbb{P}(M_{\mathbf{v}_n} | \mathcal{C}_{\mathbf{v}_n, l_a}), \\ &= \frac{\theta P_{ab}}{2} \sum_{l_a=1}^{n_a} \binom{n_a - 1}{l_a - 1} \binom{n - 1}{l_a}^{-1} \mathbb{E}[T_{l_a+1}] + O(\theta^2). \end{aligned} \tag{A1}$$

Summing over n_a , we find

$$\mathbb{P}(E_1^{(b)}) = \sum_{n_a=1}^{n-1} \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b, E_1) = \frac{\theta P_{ab}}{2} \sum_{l_a=1}^{n-1} (l_a + 1) \mathbb{E}[T_{l_a+1}] + O(\theta^2),$$

and hence, noting that

$$\phi(i) := \lim_{\theta \rightarrow 0} \mathbb{P}[\mathbf{n} = (n-i)\mathbf{e}_a + i\mathbf{e}_b | E_1^{(b)}] = \lim_{\theta \rightarrow 0} \frac{\mathbb{P}[\mathbf{n} = (n-i)\mathbf{e}_a + i\mathbf{e}_b, E_1^{(b)}]}{\mathbb{P}(E_1^{(b)})},$$

substituting for the numerator and denominator and letting $\theta \rightarrow 0$ yields the given result. $\square$

Henceforward we denote the trinomial coefficient by $\binom{n}{i,j,k}$

Proof of Lemma 1. The argument here is very similar to that of the *Proof* of (2). This time a coalescent tree compatible with $\mathbf{v}_n$ and with $E_{2N}^{(b,c)}$ must exhibit the following sequence of events (Jenkins and Song 2011, Lemma 4.1): $n_a - l_y$ coalescence events of type a alleles, $n_b - m$ coalescence events of type b alleles, and $n_c - 1$ coalescence events of type c alleles, in some interspersed ordering; a mutation event reverting the sole remaining type c allele to a type b ; $l_y - l_o$ coalescence events of type a alleles and m coalescence events of type b alleles; a mutation event reverting the sole remaining type b allele to a type a ; and coalescence of the remaining $l_o + 1$ type a lineages to a most recent common ancestor of the sample.

Here, l_y is the number of extant type a lineages at the time of the younger mutation event, l_o is the number of type a lineages at the time of the older mutation event, and m is the number of type b lineages at the time of the younger mutation event (Figure 3). Arguing as above, any one of the $\binom{n-l_y-m-1}{n_a-l_y, n_b-m, n_c-1} \binom{m+l_y-l_o}{m}$ compatible sequences $\mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}$ satisfies

$$\mathbb{P}(\mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}) = \frac{n_a!(n_a-1)!n_b!(n_b-1)!n_c!(n_c-1)!}{n!(n-1)!} m(m+1)l_o(l_o+1),$$

and the two mutation events land on the correct pair of branches with probability

$$\mathbb{P}(M_{\mathbf{v}_n} | \mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}) = \mathbb{E} \left[\left(\frac{\theta}{2} L_n \right)^2 \frac{e^{-(\theta/2)L_n}}{2!} \cdot 2 \frac{T_{m+l_y+1}}{L_n} \frac{T_{l_o+1}}{L_n} \right] = \frac{\theta^2}{4} \mathbb{E} [T_{m+l_y+1} T_{l_o+1}] + O(\theta^3).$$

Hence, following the reasoning that led to (A1), we find

$$\begin{aligned} & \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c, E_{2N}^{(b,c)}) \\ &= P_{ab} P_{bc} \binom{n}{n_a, n_b, n_c} \\ & \times \sum_{l_y=1}^{n_a} \sum_{l_o=1}^{l_y} \sum_{m=1}^{n_a} \binom{n-l_y-m-1}{n_a-l_y, n_b-m, n_c-1} \binom{m+l_y-l_o}{m} \mathbb{P}(\mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}) \mathbb{P}(M_{\mathbf{v}_n} | \mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}). \end{aligned}$$

Substituting for each term in the summand and simplifying leads to (4), and summing over triallelic configurations $\mathbf{n}$ yields (6).

The nonnested case is similar. There are $\binom{n-m-l_y-1}{n_a-l_y, n_b-m, n_c-1} \binom{m+l_y-l_o}{m-1}$ possible $\mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}$ (Jenkins and Song 2011, Lemma 4.2), each with probability

$$\mathbb{P}(\mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}) = \frac{n_a!(n_a-1)!n_b!(n_b-1)!n_c!(n_c-1)!}{n!(n-1)!} l_y(l_y+1)l_o(l_o+1).$$

The two mutation events land on the correct pair of branches and are in the correct age order with probability

$$\begin{aligned} \mathbb{P}(M_{\mathbf{v}_n}^{(b,c)} | \mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}) &= \mathbb{E} \left[\left(\frac{\theta}{2} L_n \right)^2 \frac{e^{-(\theta/2)L_n}}{2!} \cdot 2 \frac{T_{m+l_y+1}}{L_n} \frac{T_{l_o+1}}{L_n} \right] \cdot \frac{1}{1 + \delta_{m+l_y, l_o}} \\ &= \frac{\theta^2}{4} \frac{\mathbb{E} [T_{m+l_y+1} T_{l_o+1}]}{1 + \delta_{m+l_y, l_o}} + O(\theta^3). \end{aligned} \tag{A2}$$

The additional factor on the right-hand side of (A2) accounts for the fact that, should the two mutations arise during the same epoch ($m + l_y + 1 = l_o + 1$), then only with probability 1/2 is the mutation giving rise to the b allele the elder one. Finally,

$$\begin{aligned}
& \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c, E_{2, \mathcal{N}' \mathcal{N}}^{(b,c)}) \\
&= P_{ab} P_{ac} \binom{n}{n_a, n_b, n_c} \\
&\times \sum_{l_y=1}^{n_a} \sum_{l_o=1}^{l_y+1} \sum_{m=1}^{n_b} \binom{n-m-l_y-1}{n_a-l_y, n_b-m, n_c-1} \binom{m+l_y-l_o}{m-1} \mathbb{P}(\mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}) \mathbb{P}(M_{\mathbf{v}_n}^{(b,c)} | \mathcal{C}_{\mathbf{v}_n, m, l_y, l_o}),
\end{aligned}$$

which yields (5) after substituting and simplifying. Summing over triallelic configurations $\mathbf{n}$ yields (7). $\square$

Proof of Theorem 2. Conditional on the event E'_1 , the single simultaneous mutation event occurs uniformly on the $n-1$ coalescence nodes of the coalescent tree. Hence, immediately prior to (i.e., more recently than) the coalescence event at which the mutation event occurred, there are L lineages ancestral to the present-day sample, with $L \sim \text{Uniform}\{2, \dots, n\}$. Given $L = l$, it is well known (e.g., Griffiths and Tavaré 1998) that the distribution of the number of leaves subtending each of these l lineages is uniform on the $\binom{n-1}{l-1}$ possible compositions of n . Of these, $\binom{n_a-1}{l-3}$ compositions have n_b leaves and n_c leaves, respectively, subtending the two lineages coalescing into the node that experiences the simultaneous mutation event. Hence,

$$\mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c | E'_1, \{L = l\}) = Q_{a,\{b,c\}} \frac{\binom{n_a-1}{l-3}}{\binom{n-1}{l-1}},$$

for $3 \leq l \leq n$ (and 0 otherwise). Also note that

$$\mathbb{P}(O_3 | E'_1) = \frac{n-2}{n-1}.$$

This is the probability that the simultaneous mutation event does not occur at the oldest of the coalescence nodes ($l = 2$), which would lead to a sample containing no copies of the ancestral allele.

Putting all this together we have

$$\begin{aligned}
\mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c | O_3, E'_1) &= \frac{\mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c, O_3 | E'_1)}{\mathbb{P}(O_3 | E'_1)}, \\
&= \frac{\sum_{l=2}^n \mathbb{P}(L = l | E'_1) \mathbb{P}(\mathbf{n} = n_a \mathbf{e}_a + n_b \mathbf{e}_b + n_c \mathbf{e}_c, O_3 | E'_1, \{L = l\})}{\mathbb{P}(O_3 | E'_1)}, \\
&= \frac{n-1}{n-2} \sum_{l=3}^n \frac{1}{n-1} \cdot Q_{a,\{b,c\}} \binom{n_a-1}{l-3} \binom{n-1}{l-1}^{-1},
\end{aligned}$$

which simplifies to Equation 9. $\square$

GENETICS

Supporting Information

<http://www.genetics.org/lookup/suppl/doi:10.1534/genetics.113.158584/-/DC1>

General Triallelic Frequency Spectrum Under Demographic Models with Variable Population Size

Paul A. Jenkins, Jonas W. Mueller, and Yun S. Song

File S1

Supporting Information

1 Tests for the existence of a mechanism of simultaneous mutation

In the main text we describe an analysis of the simultaneous mutation mechanism under a population model in which there is random mating and for which the population size varies in time. In this supplement we investigate the effect of varying our modelling assumptions on the conclusions of this test.

1.1 Test for a population of constant size.

Consider a simpler model in which $N_t \equiv N_0$. We applied our likelihood ratio test to the data described in the main text, and obtained $\ln \tilde{\Lambda} = 5.71$ and $\hat{\lambda} = 0.21$. Because we have obtained closed-form expressions for $\tilde{\phi}_0$, $\tilde{\phi}_S$ and therefore also $\tilde{\phi}_\lambda$, it is straightforward to compute the whole likelihood curve for λ (Figure S1). We estimated a 95% confidence interval of $[0.07, 0.35]$ for λ by parametric bootstrap. In other words, we simulated a complete dataset of $M = 96$ triallelic sites with the three observed alleles a_i , b_i , c_i , and the sample size $n^{(i)}$, fixed to match their observed values at the i th site, for each $i = 1, \dots, M$. Sample counts were then drawn from $\tilde{\phi}_\lambda$ for each site, and a bootstrapped MLE, $\hat{\lambda}^*$, was computed for this simulated sample. The entire procedure was repeated 1000 times to obtain an empirical distribution $\{\hat{\lambda} - \hat{\lambda}_k^*\}_{k=1, \dots, 1000}$, which serves as an approximation of the distribution of $\lambda - \hat{\lambda}^*$. The confidence interval is obtained from the central 95th percentile of this distribution (centered around $\hat{\lambda}$) (Figure S1). The upper limit of our confidence interval for $\hat{\lambda}$ is slightly lower than an estimate of $\lambda \approx 0.5$ by Hodgkinson and Eyre-Walker (2010), who used a different method but also assumed a population of constant size in their calculation.

We further estimated the null distribution of $\tilde{\Lambda}$ by bootstrap simulation; because the parameter of interest is not necessarily an interior point of the interval $[0, 1]$, the usual χ_1^2 approximation does not apply. We used the same procedure for the simulation of bootstrap samples as described above, except this time drawing sample counts from $\tilde{\phi}_0$ rather than $\tilde{\phi}_\lambda$. This method yielded an empirical distribution of test statistics, $\{\ln \tilde{\Lambda}_k\}_{k=1, \dots, 1000}$ (Figure S2), whose 95th% percentile tail provides an appropriate p -value. In fact, the largest empirical $\tilde{\Lambda}$ value was 4.03, much less than our observed value. From this we conclude that $p < 0.001$ and we reject the null hypothesis that $\lambda = 0$. This is consistent with $\lambda = 0$ lying outside our 95% confidence interval for $\hat{\lambda}$.

One important point to note is that nine of the 96 sites in the data overwhelmingly contribute to the significance of the test statistic. These nine sites are the SNPs for which $n_b = n_c = 1$. Using our formulas for the sampling frequency spectrum under the two possible mutation mechanisms, we have $\tilde{\phi}_S(n^{(i)} - 2, 1, 1) > 0.33$ and $\tilde{\phi}_0(n^{(i)} - 2, 1, 1) < 0.04$ across all sites i with this configuration. Thus the sample counts observed at these sites are roughly more than 10 times more likely under the simultaneous mutation mechanism. However, it could be argued that these highly influential doubly-singleton triallelic sites are also prime candidates for being the result of sequencing error. To examine the dependence of the conclusion of our test on these doubly-singleton sites, we did the following. We first ranked the 9 sites from most influential to least influential, based on the ratio

$$\frac{\tilde{\phi}_S(n^{(i)} - 2, 1, 1)}{\tilde{\phi}_0(n^{(i)} - 2, 1, 1)},$$

(where $n^{(i)}$ is the sample size of the i th site). We then removed the topmost influential sites in

succession and re-computed the test statistics via the previous methodology. Results are shown in Table S1.

From the table we can see that for our test statistic to lose its significance at the 5% level, we must remove six or more of the most influential sites. Is it reasonable that at least six of these nine sites are the result of sequencing errors? One of the attractive aspects of the SeattleSNPs project and Environmental Genome Project datasets is that they are resequenced to a high quality, with each SNP confirmed in multiple reactions. Assuming a quality score of $Q > 25$, the probability of a base-calling error in two independent reactions is 10^{-5} (Hodgkinson and Eyre-Walker, 2010), and it therefore seems reasonable to conclude that it is very unlikely that the presence of more than five of these nine sites are the result of sequencing error. Thus, we reject the null hypothesis that the mixture parameter $\lambda = 0$ even if we allow for at least some of these sites to have resulted from sequencing errors.

In order to test the robustness of our results to the choice of $\mathbf{P}$, we also redid our analysis with a “noninformative” choice of $\mathbf{P}$ given by

$$\mathbf{P} = \begin{matrix} & \begin{matrix} \text{A} & \text{T} & \text{C} & \text{G} \end{matrix} \\ \begin{matrix} \text{A} \\ \text{T} \\ \text{C} \\ \text{G} \end{matrix} & \begin{pmatrix} 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \end{pmatrix} \end{matrix}.$$

This time we find $\hat{\lambda} = 0.22$ and $\ln \tilde{\Lambda} = 7.02 > 6.15 = \max_k \ln \tilde{\Lambda}_k^*$, and thus $p < 0.001$ as before. We conclude that our test exhibits good robustness to the perturbations in the exact specification of $\mathbf{P}$, at least in the direction of the alternative choice of $\mathbf{P}$ investigated here.

1.2 Test for a subdivided population of nonconstant size with migration.

To extend our test for the mechanism that gives rise to triallelic sites to a subdivided population with migration, we proceed by simulation. There are two complications which make this approach more difficult. First, the sample space for the joint frequency spectrum is over collections of vectors, where each vector in the collection records the allele counts in one of the subpopulations. This vastly increases the size of the space over which simulation must be performed, dramatically increasing computation time. Second, to compute the frequency spectrum under this model requires all of our samples to be subpopulation-labelled. As we describe below, this is far from the case for data from SeattleSNPs and NIEHS. Discarding samples whose subpopulation membership is unknown reduces the sample size and hence the power of any test.

We now describe our Monte Carlo procedure for approximating the joint frequency spectrum of (subpopulation-labelled) samples from a subdivided population. It is based on an algorithm of Hudson (2001) who was interested in two linked diallelic SNPs rather than a single triallelic SNP; we give a brief description and refer the reader there for further details. First consider generalizing $\tilde{\phi}_0(n_a, n_b, n_c)$ to a subdivided population, which we write in the form $\tilde{\phi}_0(\mathbf{m})$ with $\mathbf{m} = (n_a^{[u]}, n_b^{[u]}, n_c^{[u]})_{u=1, \dots, U}$ denoting the collection of sample counts for each subpopulation, $u = 1, \dots, U$. (We focus on a single site, and the superscript refers to the subpopulation label rather than a site index. The site is assumed to be triallelic in the combined sample, but it need not be triallelic within each subpopulation.) In order to take a Monte Carlo approximation we first write

the relevant sampling probabilities in the following form:

$$\begin{aligned}\mathbb{P}(\mathbf{m}, O_3^{(a,b,c)}) &= \mathbb{E} \left[\sum_{j,k} \mathbb{I}\{\mathbf{m}, \tau, j, k\} (1 - e^{-\theta l_j/2})(1 - e^{-\theta l_k/2}) e^{-\theta(L_n - l_j - l_k)/2} P_{a(j)b} P_{a(k)c} \right], \\ &= \frac{\theta^2}{4} \mathbb{E} \left[\sum_{j,k} \mathbb{I}\{\mathbf{m}, \tau, j, k\} l_j l_k P_{a(j)b} P_{a(k)c} \right] + O(\theta^3).\end{aligned}$$

Here the expectation is taken with respect to random genealogies drawn according to the given complex demographic history. The function $\mathbb{I}\{\mathbf{m}, \tau, j, k\}$ is an indicator that ensures we count only those genealogies for which a mutation on branch j from allele $a(j) \mapsto b$ and a mutation on branch k from allele $a(k) \mapsto c$ gives rise to the configuration $\mathbf{m}$. The notation $a(\cdot)$ is taken to mean “the allele at the ancestral node of this branch”. The summation is a double-summation over all branches of τ , which considers each possible placement of the mutations giving rise to allele b and allele c . Branch j has length l_j , branch k has length l_k , and the total tree length is L_n . Ignoring terms of $O(\theta^3)$ discards the possibility of more than two mutations, as usual.

Similarly, we can write

$$\mathbb{P}(O_3^{(a,b,c)}) = \frac{\theta^2}{4} \mathbb{E} \left[\sum_{j,k} \mathbb{I}\{O_3^{(a,b,c)}, \tau, j, k\} l_j l_k P_{a(j)b} P_{a(k)c} \right] + O(\theta^3),$$

where $\mathbb{I}\{O_3^{(a,b,c)}, \epsilon, j, k\}$ is an indicator for the event that a mutation on branch j from $a(j) \mapsto b$ and a mutation on branch k from $a(k) \mapsto c$ gives rise to a triallelic sample comprising the three alleles a , b , and c . Finally, we have that

$$\begin{aligned}\tilde{\phi}_0(\mathbf{m}) &= \lim_{\theta \rightarrow 0} \mathbb{P}(\mathbf{m} | O_3^{(a,b,c)}) = \lim_{\theta \rightarrow 0} \frac{\mathbb{P}(\mathbf{m}, O_3^{(a,b,c)})}{\mathbb{P}(O_3^{(a,b,c)})}, \\ &= \frac{\mathbb{E} \left[\sum_{j,k} \mathbb{I}\{\mathbf{m}, \tau, j, k\} l_j l_k P_{a(j)b} P_{a(k)c} \right]}{\mathbb{E} \left[\sum_{j,k} \mathbb{I}\{O_3^{(a,b,c)}, \tau, j, k\} l_j l_k P_{a(j)b} P_{a(k)c} \right]}, \\ &\approx \frac{\frac{1}{N} \sum_{i=1}^N \left[\sum_{j,k} \mathbb{I}\{\mathbf{m}, \tau^{(i)}, j, k\} l_j^{(i)} l_k^{(i)} P_{a(j)b} P_{a(k)c} \right]}{\frac{1}{N} \sum_{i=1}^N \left[\sum_{j,k} \mathbb{I}\{\mathbf{m}, \tau^{(i)}, j, k\} l_j^{(i)} l_k^{(i)} P_{a(j)b} P_{a(k)c} \right]}.\end{aligned}$$

The final step replaces each of the numerator and denominator with a Monte Carlo sample of N simulated genealogies, $\tau^{(1)}, \dots, \tau^{(N)}$ (which can be reused in both numerator in denominator, though this may introduce some bias).

The frequency spectrum under the simultaneous mutation mechanism can be obtained in a similar manner. We write

$$\mathbb{P}(\mathbf{m}, O_3^{(a,b,c)} | E'_1) = \mathbb{E} \left[\sum_{v \in V_\tau} \frac{1}{2} [\mathbb{I}\{\mathbf{m}, \tau, v, b, c\} + \mathbb{I}\{\mathbf{m}, \tau, v, c, b\}] \frac{1}{n-1} Q_{a, \{b,c\}} \right],$$

summing over all coalescence vertices V_τ of τ . Then vertex v is chosen with probability $1/(n-1)$, and the simultaneous mutation at v produces the correct derived alleles with probability $Q_{a, \{b,c\}}$. The factor of $1/2$ accounts for the two possible assignments of the alleles b and c to the branches

descending from the simultaneous mutation event, and there is an indicator corresponding to each of these assignments, which is 1 if the configuration at the leaves of the genealogical tree is $\mathbf{m}$ and zero otherwise. Finally, we obtain a Monte Carlo approximation of $\tilde{\phi}_S$ by noting that placing the simultaneous mutation event at the root vertex does not give rise to a triallelic site, so

$$\mathbb{P}(O_3^{(a,b,c)}|E'_1) = \frac{n-2}{n-1}Q_{a,\{b,c\}},$$

(see the proof of Theorem 4), and hence

$$\begin{aligned}\tilde{\phi}_S(\mathbf{m}) &= \frac{\mathbb{P}(\mathbf{m}, O_3^{(a,b,c)}|E'_1)}{\mathbb{P}(O_3^{(a,b,c)}|E'_1)} = \frac{1}{n-2}\mathbb{E}\left[\sum_{v \in V_\tau} \frac{1}{2}[\mathbb{I}\{\mathbf{m}, \tau, v, b, c\} + \mathbb{I}\{\mathbf{m}, \tau, v, c, b\}]\right], \\ &\approx \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \frac{1}{n-2} \sum_{v \in V_{\tau^{(i)}}} \frac{1}{2}[\mathbb{I}\{\mathbf{m}, \tau^{(i)}, v, b, c\} + \mathbb{I}\{\mathbf{m}, \tau^{(i)}, v, c, b\}].\end{aligned}$$

The above method is extremely general, in the sense that it requires only that one is able to simulate genealogical trees from the assumed model. Indeed, Hudson (2001) did not consider population subdivision, yet, other than the simulation of genealogical histories from $\mathbf{ms}$, the details of the algorithm are unchanged when subdivision is also considered.

To account for population subdivision, we assumed as the true demographic model that inferred by Gutenkunst et al. (2009). Their method extended that of Williamson et al. (2005) to allow for an ancestral population to split into modern subpopulations with rare but continuous migration between subpopulations and each subpopulation allowed to change in size with time. Using data from three HapMap populations: 12 Yoruba individuals from Ibadan, Nigeria (YRI); 22 CEPH Utah residents with ancestry from northern and western Europe (CEU); and 12 Han Chinese individuals from Beijing, China (CHB); Gutenkunst et al. (2009) fit a model governing the historical relationship of the three populations during the human expansion out of Africa. The data from SeattleSNPs and the Environmental Genome Project comprise samples from panels made up of individuals in HapMap populations and samples from other panels. Individuals in the latter are classified as of European, African-American, Asian, or Hispanic descent. For simplicity and following Hodgkinson and Eyre-Walker (2010), we focus on a model only for an African and a non-African (European) population, excluding Asian and Hispanic samples from our analysis. To obtain subpopulation labels (not provided in Hodgkinson and Eyre-Walker, 2010) we reanalyzed the data available from SeattleSNPs and the Environmental Genome Project, excluding sites in coding regions and so on as described above. This time we excluded from each site any individuals for whom subpopulation labels were not given, leaving $M = 34$ triallelic sites in total.

We applied the Monte Carlo method described above for approximating the frequency spectrum under the null and alternative models of mutation, given the complex demographic history of Gutenkunst et al. (2009). The Monte Carlo error at sites whose sample configurations were of very low probability could unduly affect the overall likelihood ratio statistic, so we varied the Monte Carlo sample size $\mathcal{N}$, from 10^6 to 10^{10} for different sites, according to a preliminary estimate of the magnitude of the probability of their sample configuration. (As an extreme example, sites with configurations of very low probability could potentially have this probability estimated as 0, which would also cause the overall likelihood ratio statistic to be 0.) Using this approach we obtained $\ln \tilde{\Lambda} = 2.01$ ($p \approx 0.42$) and $\hat{\lambda} = 0.16$. Due to the computationally-intensive nature of the simulation, the p -value is based on only a very coarse approximation of the null distribution of $\ln \hat{\Lambda}$ using 50 bootstrapped test statistics (Figure S3); nonetheless, it is clearly non-significant at the 5% level.

References

- Gutenkunst, R. N., Hernandez, R. D., Williamson, S. H., and Bustamante, C. D. (2009). Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genetics*, **5**, e1000695.
- Hodgkinson, A. and Eyre-Walker, A. (2010). Human triallelic sites: evidence for a new mutational mechanism? *Genetics*, **184**, 233–241.
- Hudson, R. R. (2001). Two-locus sampling distributions and their application. *Genetics*, **159**, 1805–1817.
- Williamson, S. H., Hernandez, R., Fledel-Alon, A., Zhu, L., and Bustamante, C. D. (2005). Simultaneous inference of selection and population growth from patterns of variation in the human genome. *Proceedings of the National Academy of Sciences of the United States of America*, **102**, 7882–7887.

2 Supporting Figures and Table

TABLE S1: The effect on the test statistic of triallelic sites for which both derived alleles are singletons. As more and more of these sites are removed from the data in succession, the mixture parameter MLE, $\hat{\lambda}$, and the significance of the test statistic, gradually decrease.

Number of sites		$\hat{\lambda}$	p -value
removed			
0		0.204	≤ 0.001
1		0.189	≤ 0.001
2		0.171	≤ 0.003
3		0.153	≤ 0.006
4		0.133	≤ 0.013
5		0.113	≤ 0.036
6		0.090	≤ 0.090
7		0.064	≤ 0.161
8		0.036	≤ 0.283
9		0.002	≤ 0.461

FIGURE S1: Likelihood function for λ , relative to its value at $\lambda = 0$, under the assumption of a panmictic population of constant size. The vertical dashed line shows the location of the maximum likelihood estimate $\hat{\lambda}$, and the histogram beneath the curve illustrates the empirical distribution of our simulated maximum likelihood estimates $\{\hat{\lambda}^*\}_{k=1, \dots, 1000}$. One can see that the distribution of $\hat{\lambda}^*$ strongly reflects the shape of the likelihood function.

FIGURE S2: Empirical null distribution of $\ln \tilde{\Lambda}$. A Gaussian kernel density estimate is superimposed over the histogram of $\ln \tilde{\Lambda}$, and the dashed lines indicate the locations of the largest of these values as well as the $\ln \tilde{\Lambda}$ value from our data.

FIGURE S3: Empirical distribution of $\ln \tilde{\Lambda}$ under a null model incorporating complex demography. The dashed line indicated the value of the test statistic computed from the actual data, $\ln \tilde{\Lambda} = 2.01$.