

Characteristics of Neutral and Deleterious Protein-Coding Variation among Individuals and Populations

Wenqing Fu,^{1,*} Rachel M. Gittelman,¹ Michael J. Bamshad,^{1,2} and Joshua M. Akey^{1,*}

Whole-genome and exome data sets continue to be produced at a frenetic pace, resulting in massively large catalogs of human genomic variation. However, a clear picture of the characteristics and patterns of neutral and deleterious variation within and between populations has yet to emerge, given that recent large-scale sequencing studies have often emphasized different aspects of the data and sometimes appear to have conflicting conclusions. Here, we comprehensively studied characteristics of protein-coding variation in high-coverage exome sequence data from 6,515 European American (EA) and African American (AA) individuals. We developed an unbiased approach to identify putatively deleterious variants and investigated patterns of neutral and deleterious single-nucleotide variants and alleles between individuals and populations. We show that there are substantial differences in the composition of genotypes between EA and AA populations and that small but statistically significant differences exist in the average number of deleterious alleles carried by EA and AA individuals. Furthermore, we performed extensive simulations to delineate the temporal dynamics of deleterious alleles for a broad range of demographic models and use these data to inform the interpretation of empirical patterns of deleterious variation. Finally, we illustrate that the effects of demographic perturbations, such as bottlenecks and expansions, often manifest in opposing patterns of neutral and deleterious variation depending on whether the focus is on populations or individuals. Our results clarify seemingly disparate empirical characteristics of protein-coding variation and provide substantial insights into how natural selection and demographic history have patterned neutral and deleterious variation within and between populations.

Introduction

Mutations impose a substantial burden on fitness, disease, and longevity through the introduction of deleterious alleles into the population.^{1–5} A deeper understanding of deleterious variation in humans will have profound implications for disease-mapping studies, personal genomics, and predictive medicine. A considerable amount of theoretical work has been done to inform the dynamics of deleterious variation across a range of demographic models.^{6–8} Moreover, a large number of empirical studies in humans have been performed to survey patterns of deleterious variation within and between populations.^{9–13} For example, in a study of 15 African American (AA) and 20 European American (EA) individuals, Lohmueller et al.¹⁰ found that the European sample had an excess of putatively deleterious variants and through simulations demonstrated that this was most likely a consequence of the Out-of-Africa bottleneck. The proportional increase in deleterious variation in European versus African populations has also been observed in other studies.^{8,14–16} Furthermore, Casals et al.¹⁷ showed that recent founder effects in the French Canadian Quebec population have led to different patterns of deleterious variation between it and the French Canadian population. Moreover, in a clever study, Szpiech et al.¹⁸ found that deleterious alleles were enriched in runs of homozygosity and that variable levels of inbreeding can influence patterns of deleterious variation across populations.

However, not all studies have found a clear relationship between demographic history and empirical patterns of deleterious variation. For example, Tennesen et al.¹² noted that characteristics of deleterious variants in EA and AA individuals are sensitive to how deleterious sites are defined. In addition, through detailed simulations and analyses of derived allele frequency (DAF) in a large exome sequencing data set, Simons et al.¹¹ suggested that the deleterious-mutation load is insensitive to recent population history and that the average number of derived alleles per individual at putatively deleterious sites is not significantly different across populations. Similarly, Do et al.¹⁹ have recently argued that there are no differences in the per-genome accumulation of deleterious alleles across diverse human populations, which appears to contradict previous claims of differences in the proportion of deleterious variants across populations.^{10,14,15}

Thus, despite the substantial amount of work that has been devoted to documenting and interpreting patterns of neutral and deleterious protein-coding variation in humans, a number of outstanding questions remain. Here, we describe a comprehensive analysis of protein-coding variation in a previously described high-coverage exome sequence data set consisting of 6,515 individuals of European and African ancestry and generated as part of the NHLBI Exome Sequencing Project (ESP).¹⁴ Furthermore, we performed extensive simulations of neutral and deleterious variation to help interpret empirical patterns of protein-coding variation. We show that many seemingly

¹Department of Genome Sciences, University of Washington, Seattle, WA 98195, USA; ²Department of Pediatrics, University of Washington, Seattle, WA 98195, USA

*Correspondence: wqfu@uw.edu (W.F.), akeyj@uw.edu (J.M.A.)

<http://dx.doi.org/10.1016/j.ajhg.2014.09.006>. ©2014 by The American Society of Human Genetics. All rights reserved.

disparate observations of neutral and deleterious variation can be accounted for by opposing variation patterns that manifest depending on how variation is summarized and whether the focus is on individuals or populations. Our empirical and simulation results provide insight into how natural selection and demographic history have interacted to influence neutral and deleterious variation within and between populations.

Material and Methods

Analysis of Empirical Data

Analysis of Samples and Exome Sequencing Data

We analyzed the exomes of 6,515 individuals, including 4,298 EA individuals (1,879 males and 2,419 females) and 2,217 AA individuals (582 males and 1,635 females), from the NHLBI ESP.¹⁴ Exome data were subjected to standard quality-control filters as previously described.¹⁴ We further removed sites whose ancestral inference was inferred with low confidence in the six primate EPO (Enredo, Pecan, Ortheus) alignments.²⁰ The final data set consisted of 1,110,148 single-nucleotide variants (SNVs) in autosomes and X chromosomes.

To avoid biases caused by different sample sizes, for all population-level analyses, such as estimating the site-frequency spectrum (SFS), we randomly sampled 2,217 EA individuals to match the sample size in AA individuals. For all individual-level analyses, we defined SNVs in individuals as sites that are heterozygous or homozygous for the derived allele. We compared the per-individual number of SNVs, heterozygotes, derived homozygotes, and derived alleles between EA and AA individuals with Mann-Whitney tests. Furthermore, to account for heterogeneity in missing data among individuals, we normalized the per-individual number of derived alleles by the per-individual number of total alleles that passed filtering criteria (see Fu et al.¹⁴).

We also used an alternative method to account for misidentification of ancestral states. Specifically, we identified the putative ancestral state of each SNV by comparing it to the chimpanzee genome (panTro2), and we corrected ancestral misidentification by using a context-dependent mutation model.²¹ In brief, this method accounts for the probability of misidentifying the ancestral state of a SNV by modeling the observed number of derived alleles (or derived homozygotes) as a mixture of SNVs whose ancestral states were correctly identified and those that were misidentified under the context-dependent substitution process.²² In total, 1,148,406 SNVs were used in these analyses.

Moreover, we used Fisher's exact test to compare the average number of derived alleles per individual as a function of allele frequency between deleterious variants and neutral variants within populations, as well as between EA and AA populations for deleterious variants. For example, in the comparison of the enrichment of deleterious rare variants (DAF < 0.05%) in one population, the elements of the 2 × 2 table consisted of the average per-individual number of derived alleles of rare variants (DAF < 0.05%) and of variants with other frequency (DAF ≥ 0.05%) for both the deleterious and neutral sites.

Definition of Deleterious Variants

Quantifying evolutionary constraint through sequence conservation is widely used for identifying genomic regions that have been subject to purifying selection.^{23,24} We used PhyloP scores²⁵ to identify putatively deleterious variants. PhyloP scores were

calculated from 36 eutherian-mammal EPO alignments downloaded from the Ensembl Genome Browser (release 70) in enhanced metafile format (emf). These emf alignments were converted to multiple alignment format (maf) with the script "emf2maf.pl," also downloaded from Ensembl. Alignment blocks in maf were then sorted with the mafTools package. Finally, sorted maf alignments were converted to SS format with the msa_program in the PHAST package. To calculate scores, we ran PhyloP (PHAST package) with the following command line option: -msa-format SS -wig-scores -mode CONACC -method LRT.

The calculation of PhyloP scores also requires a neutral model of evolution. For this, we used the phylogenetic tree provided with the 36 mammalian alignments and the substitution-rate matrix and nucleotide frequencies from the placentalMammals.mod file downloaded from the UCSC Genome Browser. PhyloP scores in wiggle (wig) format were converted to bed files with the BEDOPS package.²⁶ PhyloP scores were calculated with and without the human reference sequence (denoted as PhyloP_H and PhyloP_{NH}, respectively). Conditional on the 36-way eutherian-mammal phylogeny, simulations were performed with the base_evolve program in the PHAST package.

Population-Genetics Simulations

Forward Population Simulation for Protein-Coding Sequences

We performed forward population simulations with the program SFS_CODE²⁷ under different demographic models and selective regimes. We considered three general demographic models, including a population bottleneck, recent accelerated growth, and a more complicated model, by using previously inferred parameters in the EA and AA samples.¹² For the bottleneck model, a bottlenecked population was simulated from a constant population with effective size $N_e = 10,000$. This population experienced a bottleneck 50 ka ago, where the population size was reduced to 10% (a close approximation of the Out-of-Africa bottleneck)²⁸ and 1% of the original size, and recovered from the bottleneck 25 ka ago. The Out-of-Africa bottleneck has also been modeled as a shorter and more severe bottleneck.²⁹ In this model, a constant population ($N_e = 10,000$) experienced a bottleneck 118 ka ago and a quick recovery 108 ka ago. During the bottleneck, the population size was reduced to 7.57% of the original. We also simulated data under this model to study how robust our results are to particular implementations of the Out-of-Africa bottleneck.

For the model of recent population growth, a population started expanding from a constant population ($N_e = 10,000$) 5 ka ago. We considered different growth rates, including 0%, 2.0% (a close approximation for the population with European ancestry),¹² and 3.0% per generation.

In the more realistic demographic model, European and African populations split 51 ka ago, and the European lineage incurred two bottleneck events (the Out-of-Africa bottleneck 51 ka ago and the split of non-African populations 23 ka ago) and an initial population expansion with a growth rate of 0.307% per generation, whereas the African population evolved as a constant population during this period.²⁸ Beginning 5.115 ka ago, accelerated population growth occurred for both European and African populations with growth rates of 1.95% and 1.66%, respectively.¹² The simulated AA population is a result of recent admixture from European (20%) and African (80%) populations.

A total of 2,500 individuals were simulated for each parameter combination. For each individual, we simulated 5,000 independent genes, each with four 500 bp exons that are equally spaced with 2,000 bp introns (sequences for the introns were not

simulated). The mutation rate was set to 1.5×10^{-8} per base per generation, and the recombination rate was set to 10^{-8} per base per generation. Additive purifying selection was assumed to act on each nonsynonymous mutation. The selection coefficient $|s|$ follows a gamma distribution $\Gamma(\alpha, \beta)$, where the mean selection coefficient for newly arisen deleterious mutations is 0.03 (i.e., $|s| = \alpha/2N_e\beta$).³⁰ Specifically, in bottleneck models and recent-growth models ($N_e = 10,000$), the shape parameter α and rate parameter β were set to 0.206 and 0.000343, respectively. In the complicated demographic models for EA and AA samples ($N_e = 7,310$), α and β were set to 0.206 and 0.00047, respectively. All of the simulations were based on the finite-site model. Mutations that were fixed during the burn-in period were not recorded.

An important measurement of population fitness is mutation load, defined as the proportion by which the population fitness is decreased by deleterious mutations.^{31–33} Under the assumption of no epistasis, linkage equilibrium, and additive fitness effects across sites, the total mutation load for multiple sites is the sum of mutation load across sites:

$$\begin{aligned} L &= \sum_l L_l = \sum_l (1 - w_l) \\ &= \sum_l [1 - (w_{AA}f_{AA} + 2w_{AD}f_{AD} + w_{DD}f_{DD})] \\ &= \sum_l \{1 - [p^2 + 2pq(1 + hs) + q^2(1 + s)]\} \\ &= \sum_l [- (2pqhs + q^2s)]. \end{aligned}$$

For each site l , w_{AA} , w_{AD} , and w_{DD} are the fitness of different genotypes determined by the selection coefficient (s) and the dominance coefficient (h) (i.e., $w_{AA} = 1$, $w_{AD} = 1 - hs$, and $w_{DD} = 1 - s$); the genotype frequency follows $f_{AA} = p^2$, $f_{AD} = 2pq$, and $f_{DD} = q^2$ under Hardy-Weinberg equilibrium, where p and q are the ancestral allele frequency and DAF, respectively. Thus, mutation load is determined by the number of deleterious variants carried by individuals, their effect size, and the model relating genotypes to fitness (i.e., the value of h). Unless otherwise noted, we considered an additive selection model ($h = 0.5$).

Results

Reconciling Disparate Empirical Patterns of SNVs and Alleles

We first investigated broad-scale patterns of protein-coding variation within and between populations in 6,515 individuals sequenced as part of the NHLBI ESP.¹⁴ Recently, Simons et al.¹¹ reported that the total number of derived alleles per individual was remarkably similar between individuals of European and African ancestry. This observation is striking because, superficially, it appears to contradict other well-known differences in levels of diversity between African and European populations. For example, in the ESP data, nucleotide diversity is significantly higher in AA than in EA individuals (0.030% and 0.023%, respectively; bootstrap, $p < 10^{-5}$). Similarly, it is well known that the average number of SNVs (defined here as sites that are heterozygous or homozygous for the derived allele) per individual is higher in individuals of African ancestry.^{9,34} Indeed, the average number of SNVs per individual is significantly greater in AA than in EA individuals (30,124.3 and 28,192.5, respectively; Mann-Whitney test,

$p < 10^{-15}$; Figure 1A) in the ESP data. However, the average number of SNVs per individual masks the underlying opposing patterns of heterozygous and homozygous genotypes between samples, as previously noted by Lohmueller et al.¹⁰ Specifically, the average number of heterozygous genotypes per individual is significantly higher in AA than in EA individuals (16,310.1 and 12,334.6, respectively; Mann-Whitney test, $p < 10^{-15}$; Figure 1A), whereas the average number of derived homozygous genotypes per individual is higher in EA than in AA individuals (16,049.4 and 13,930.6, respectively; Mann-Whitney test, $p < 10^{-15}$; Figure 1A).

The contrasting patterns of heterozygous and homozygous genotypes in individuals of European and African ancestry are primarily a consequence of the differences in demographic history, most notably the Out-of-Africa bottleneck, between populations. To illustrate this, we constructed a modified SFS of all variant sites, including sites where the derived allele is absent in one population and polymorphic in the other, as well as sites where the derived allele is fixed in one population and polymorphic in the other (Figure 1B). The modified SFS differs between EA and AA populations—the EA sample has more nearly fixed and fixed variants, and the AA sample contains more rare variants (Figure 1B). The modified SFS shows that the higher number of derived homozygous genotypes in EA individuals is largely due to the increased number of fixed or nearly fixed variants, as expected in bottlenecked populations because of stronger genetic drift. Similarly, the higher number of heterozygous genotypes in AA individuals is primarily attributable to the number of variants that had either been lost in EA individuals because of drift or arose after population splitting (Figure 1B). However, these two opposing patterns effectively cancel each other out such that the mean DAFs of the modified SFS in the EA and AA samples are nearly identical (0.02014 for EA and 0.02018 for AA populations). Thus, the total number of derived alleles per individual is approximately equal between EA and AA individuals (44,050.4 and 43,938.6, respectively; Mann-Whitney test, $p > 0.05$; Figure 1C).

In short, genetic diversity of protein-coding sequences is higher in AA than in EA individuals,³⁵ and AA individuals have on average more SNVs than EA individuals; however, the total number of derived alleles per individual is nearly identical. The differences in empirical patterns of SNVs and derived alleles between EA and AA individuals reflect the unique demographic history of each population, which has influenced the SFS and composition of genotypes among individuals.

Unbiased Identification of Putatively Deleterious SNVs

Many different approaches have been developed for identifying putatively deleterious variation, particularly in protein-coding sequences.^{25,36–41} The most common functional-prediction methods were recently found to have a strong bias in identifying deleterious variants depending

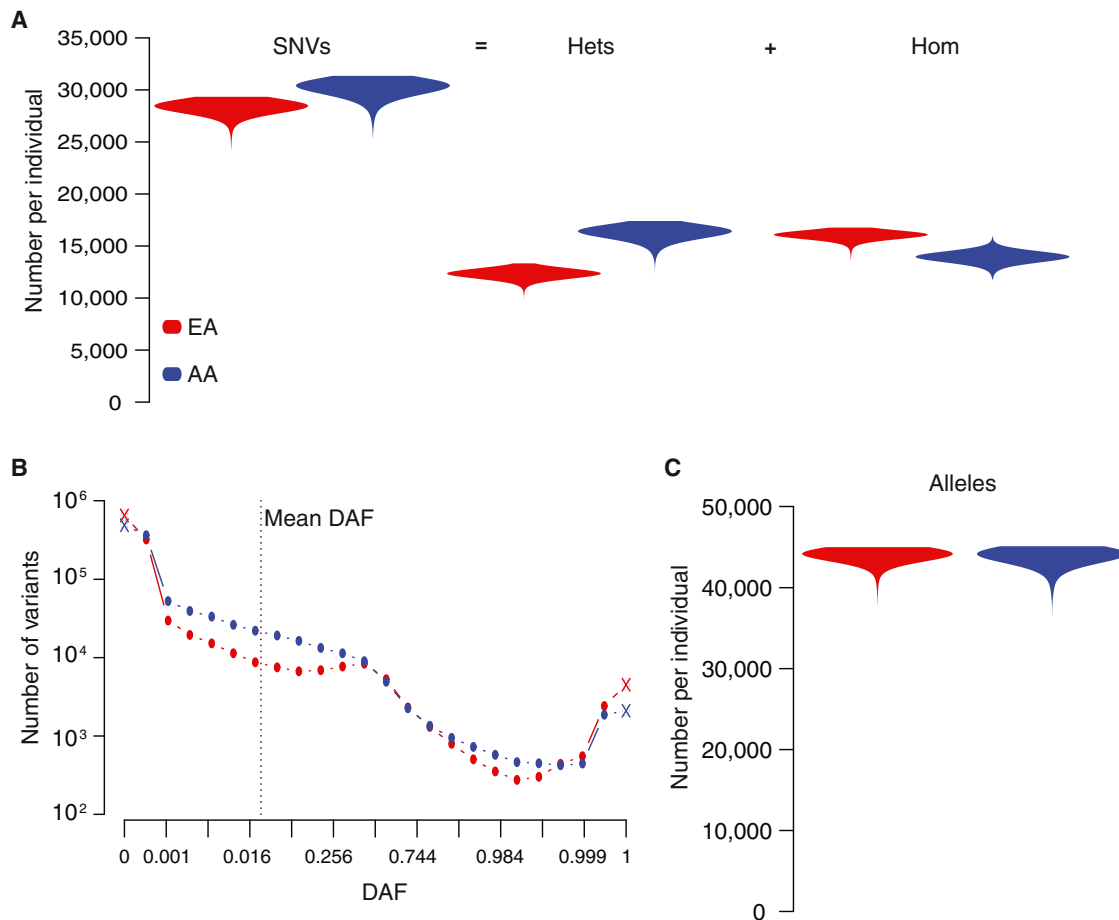


Figure 1. Patterns of Protein-Coding SNVs and Alleles among Populations and Individuals

(A) Violin plot of the number of SNVs per individual in EA and AA populations. Number of SNVs can be decomposed into whether individuals are heterozygous or homozygous for the derived allele. (B) Modified SFS (see text) in EA and AA populations. The mean DAF (dashed line) in EA and AA populations is nearly identical. (C) Violin plot of the number of derived alleles per individual in EA and AA populations. Note that the average number of derived alleles per individual is nearly identical in EA and AA populations.

on whether the reference human genome sequence contains the ancestral or derived allele.¹¹ In particular, the probability of calling a variant as deleterious is much lower for sites where the reference sequence carries the derived allele than for sites where the reference sequence carries the ancestral allele.

We hypothesized that because conservation is an integral component of nearly all functional-prediction methods, the strong reference bias is caused by assessing levels of conservation with alignments that contain the human reference sequence; thus, reference-derived sites appear to be less conserved. To test this hypothesis, we first simulated 1 Mb sequence alignments by conditioning on the 36-way eutherian-mammal phylogeny. We then calculated PhyloP with and without the simulated human sequence (denoted PhyloP_H and PhyloP_{NH} , respectively). The distribution of PhyloP_H scores was significantly different between human ancestral and derived sites (Mann-Whitney test, $p < 10^{-15}$; Figure 2A). However, this bias was completely absent in the PhyloP_{NH} scores (Mann-Whitney test, $p = 0.87$; Figure 2A), and therefore removing the human

sequence before calculating conservation is an effective strategy for mitigating reference-sequence bias.

Next, we recalculated PhyloP scores in the real data by using the 36-way eutherian-mammal alignments with and without the human reference sequence. We defined deleterious variants as those with PhyloP scores exceeding the 90% percentile of the empirical distribution of all PhyloP scores. As expected, we observed a strong reference bias for PhyloP_H , such that the proportion of variants called deleterious was markedly different for reference ancestral and derived sites as a function of DAF (Figure 2B). PhyloP_{NH} largely eliminated this bias (Figure 2B), as predicted by the simulation results. All of the analyses described below were based on the 107,736 sites identified as putatively deleterious with the use of PhyloP_{NH} , and the vast majority (85%) were nonsynonymous.

Patterns of Deleterious SNVs among Individuals and Populations

We next considered the average number of deleterious SNVs per individual in the EA and AA samples. Overall,

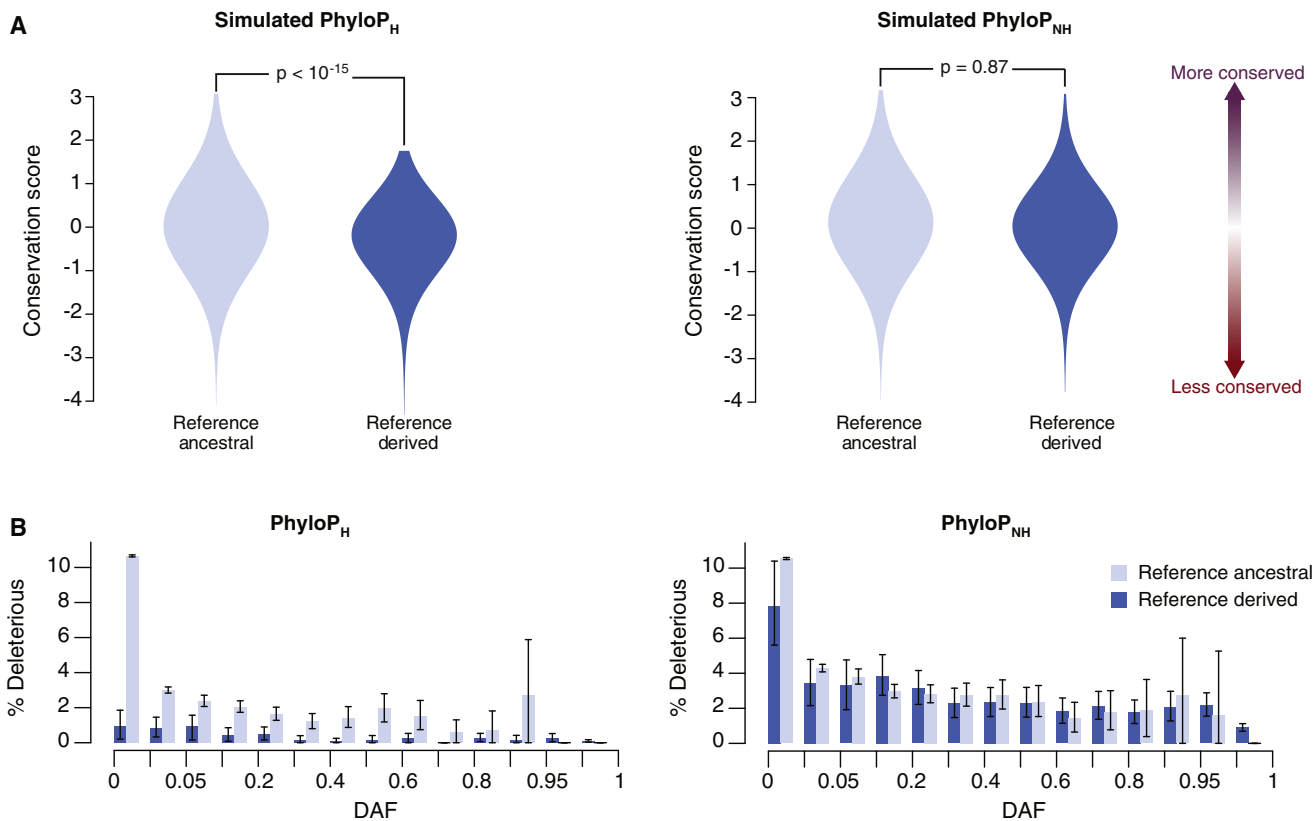


Figure 2. Identifying and Correcting Reference Bias of PhyloP

(A) Distribution of PhyloP scores calculated from sequence data that we simulated by conditioning on the branch lengths and topology of the 36-way eutherian-mammal phylogeny. PhyloP_H (left) and PhyloP_{NH} (right) were calculated on alignments including and excluding the human sequence, respectively. In each plot, the distribution of conservation scores is shown for ancestral and derived sites from the human reference sequence. Note that removing the simulated human sequence before calculating PhyloP mitigates the strong reference-bias effect.

(B) Proportion of deleterious variants in the observed data for reference ancestral and derived sites when conservation scores are calculated on alignments including (PhyloP_H, left) or excluding (PhyloP_{NH}, right) the human sequence as a function of DAF. Error bars indicate approximate 95% confidence intervals.

patterns of deleterious SNVs among individuals followed the trends observed for all SNVs (Figure 3A). Specifically, AA individuals had on average more deleterious SNVs per individual than did EA individuals (686.2 and 643.8, respectively; Mann-Whitney test, $p < 10^{-15}$; Figure 3A). Similarly, when deleterious SNVs were decomposed into genotypes, AA individuals had on average more heterozygous deleterious genotypes than did EA individuals (479.1 and 381.9, respectively; Mann-Whitney test, $p < 10^{-15}$; Figure 3A), whereas EA individuals had on average more deleterious derived homozygous genotypes than did AA individuals (262.2 and 207.3, respectively; Mann-Whitney test, $p < 10^{-15}$; Figure 3A).

Small but Significant Increase in the Number of Deleterious Alleles in EA Individuals

Recently, Simons et al.¹¹ studied patterns of deleterious variation in the ESP data and found that the mean frequency of derived deleterious alleles was not significantly different between EA and AA samples. We first replicated this observation and found no significant difference in

the mean frequency of derived deleterious alleles between EA and AA individuals (0.00423 and 0.00419, respectively; *t* test, $p = 0.82$).

To explore patterns of deleterious alleles more directly, we next compared the distributions of the number of derived deleterious alleles per individual for the EA and AA samples (Figure 3A). We found that the average number of derived deleterious alleles was slightly but significantly higher in EA individuals than in AA individuals (905.7 and 893.2, respectively; Mann-Whitney test, $p < 10^{-15}$). This result was robust to (1) different PhyloP_{NH} thresholds used to define deleterious SNVs, (2) correction for heterogeneity in levels of missing data among individuals, (3) removal of hypermutable CpG sites, (4) recalculation of PhyloP_{NH} by removal of SNVs within 10 bp of a gap in the eutherian-mammal alignment (to mitigate potential subtle biases caused by indels), and (5) sequence errors by removal of derived singletons. Furthermore, we also used a context-dependent mutation-rate model to correct for ancestral misidentification (see [Material and Methods](#)), which also resulted in small but statistically significant

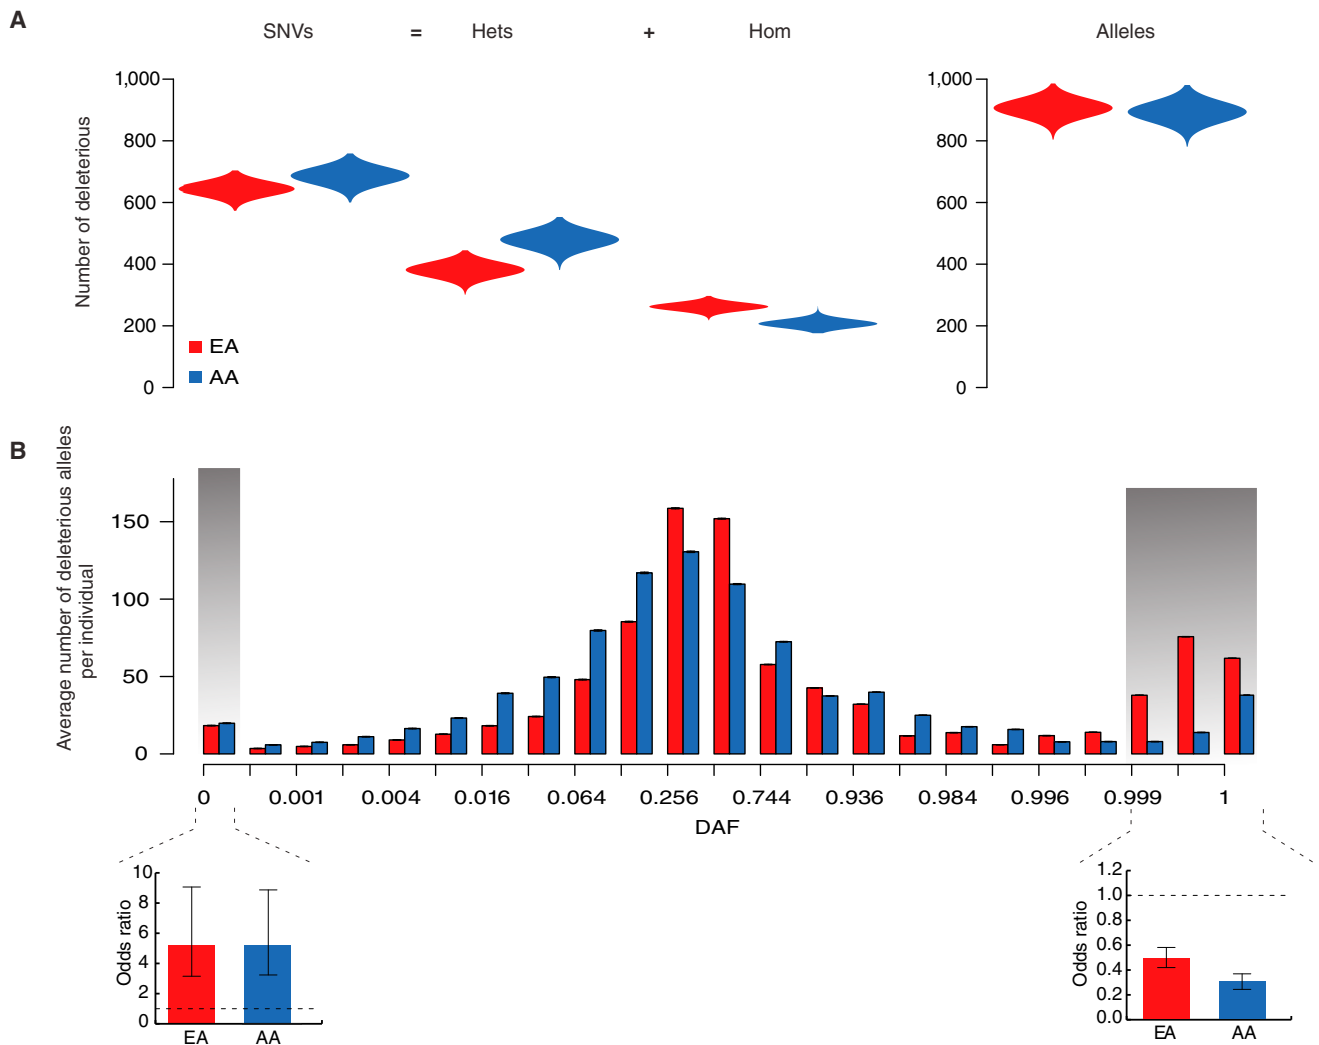


Figure 3. Empirical Patterns of Deleterious Protein-Coding Variants Carried by Individuals

(A) The average number of deleterious SNVs, heterozygous genotypes, homozygous genotypes, and derived alleles per individual. The average number of deleterious alleles per individual is small but significantly different between EA and AA individuals. (B) The average number of deleterious alleles per individual in EA and AA samples as a function of population DAF. The inset bar plots compare the odds that a derived allele is deleterious to the odds that a derived allele is neutral per individual for variants with a DAF $\leq 0.05\%$ (left) and $\geq 99.9\%$ (right) in EA and AA individuals, respectively. Error bars denote the 95% confidence interval of the odds ratio. Note that on average, EA individuals carry significantly more nearly fixed or fixed deleterious alleles than do AA individuals ($p = 8.63 \times 10^{-16}$).

differences in the average number of derived deleterious alleles carried by EA and AA individuals (994.8 and 978.0, respectively; Mann-Whitney test, $p < 10^{-15}$).

To help interpret this observation, we analyzed the average number of derived deleterious alleles per individual as a function of DAF (Figure 3B). Although most deleterious SNVs (and all SNVs) were rare in frequency at the population level, most deleterious SNVs (and all SNVs)¹² carried by individuals were common (Figure 3B). Both EA and AA individuals had on average significantly more rare (DAF $< 0.05\%$) deleterious alleles than putatively neutral alleles (Fisher's exact test, $p = 2.02 \times 10^{-8}$ and 3.61×10^{-9} for EA and AA individuals, respectively; Figure 3B). Similarly, both EA and AA individuals had on average significantly fewer nearly fixed or fixed deleterious

alleles (DAF $\geq 99.9\%$) than neutral alleles (Fisher's exact test, $p = 2.59 \times 10^{-21}$ and 2.29×10^{-28} for EA and AA individuals, respectively; Figure 3B). However, on average, EA individuals had significantly more nearly fixed or fixed deleterious alleles than did AA individuals (Fisher's exact test, $p = 8.63 \times 10^{-16}$; Figure 3B). Importantly, this small increase in the number of nearly fixed or fixed deleterious alleles was responsible for the higher average number of deleterious derived alleles in EA individuals.

Proportionally More Deleterious Variation in EA Individuals

Previously, it was observed that European populations possessed proportionally more deleterious SNVs than did African populations.^{10,34,42} However, this observation

might have been influenced by the previously unrecognized reference bias in methods of predicting deleterious variants. To this end, we used PhyloP_{NH} to calculate the proportion of deleterious SNVs (while correcting for sample size) in the EA and AA populations. Consistent with previous studies, the proportion of deleterious SNVs in EA individuals was significantly higher than that in AA individuals (9.36% and 8.28%, respectively, Fisher's exact test, $p = 1.06 \times 10^{-85}$), and therefore this finding was not an artifact of reference bias. We also found a small but statistically significant increase in the average proportion of deleterious alleles per individual in the EA and AA samples (2.06% and 2.03%, respectively; Mann-Whitney test, $p < 10^{-15}$).

Overview of Simulations

To better understand the empirical patterns of neutral and deleterious protein-coding variation described above, we performed extensive simulations. As above, all the variants that segregated in the combined simulated samples were included in the analysis. We first focused on demographic models with either population bottlenecks or recent accelerated growth to delineate how each of these processes influences patterns of deleterious variation. We then considered a more complicated and realistic model of human history that features both bottlenecks and recent growth. For each demographic model and parameter combination, we simulated 2,500 individuals, and for each individual we simulated 5,000 independent genes with an exonic length of 2 kb by using previously inferred parameters for the distribution of fitness effects of protein-coding variation³⁰ and assuming an additive model of selection (see [Material and Methods](#)). We used the simulated data to investigate characteristics of present-day deleterious SNVs in individuals and populations, when they arose, and the distribution of their fitness effects. Unless otherwise noted, we considered variants in the simulations to be deleterious if their effect size was $|s| \geq 10^{-4}$ (i.e., $2N_e|s| \gg 1$) because effects less than this are effectively neutral.

Temporal Dynamics of Deleterious Alleles in Bottleneck Models

We simulated protein-coding variation in bottleneck models where population size decreased to $0.1\times$ and $0.01\times$ 50 ka ago and then recovered to prebottleneck levels 25 ka ago ([Figure 4A](#)). Note that the decrease in size to $0.1\times$ closely approximates the estimated reduction in size during the Out-of-Africa bottleneck.²⁸ We summarized levels of neutral and deleterious variation as the mean number of derived alleles per individual per kilobase and compared the density of mutations found in contemporary individuals in bottlenecked and constant-sized populations as a function of DAF, when they arose, and their selection coefficient ([Figure 4](#)).

Several salient points emerge from [Figure 4A](#), in which the density of protein-coding variants in contemporary individuals is decomposed into the time in which the var-

iants arose. First, deleterious variants that occurred before or during the bottleneck constitute a small proportion of deleterious variants in the population (9.5% for constant-sized populations and 3.9% and 1.4% for $0.1\times$ and $0.01\times$ bottlenecked populations, respectively). However, almost all fixed and the majority of common deleterious variants in present-day individuals arose during this time period ([Figure 4A](#)), and as a result the number of derived deleterious alleles carried by these individuals was slightly higher than that carried by individuals from constant-sized populations ([Figure 4A](#)). For example, the mean density of derived deleterious alleles per individual in the bottleneck model with intensity $0.1\times$ was 0.0167/kb, whereas in individuals from constant-sized populations, it was 0.0151/kb (1.1-fold increase). Second, although individuals from a bottlenecked population carried more fixed deleterious mutations, they were nearly always of weak effect ($|s| \sim 10^{-4}$; [Figure 4A](#)). Third, strongly deleterious ($|s| > 0.01$) mutations carried by present-day individuals always arose recently and were rare in the population ([Figure 4A](#)).

The proportion of deleterious alleles in individuals sampled from the current generation was modestly affected by bottlenecks (3.05% for individuals from constant-sized populations and 3.34% and 4.67% for $0.1\times$ and $0.01\times$ bottlenecked populations, respectively; [Figure 4B](#)). However, the relative accumulation of deleterious and neutral variants exhibited different patterns depending on when mutations arose. Specifically, deleterious variants that arose before the bottleneck tended to be lost because of drift, but those that survived tended to increase in frequency; as a result, the proportion of derived deleterious alleles per individual slightly increased for variants that arose in this period ([Figure 4B](#)). Deleterious variants that occurred during the bottleneck increased both in number and in frequency because of less efficient purging, so that proportionally more derived deleterious alleles arising in this epoch accumulated in present-day individuals ([Figure 4B](#)). In contrast, and as expected, the accumulation of derived deleterious alleles that arose after the bottleneck did not change as a function of bottleneck intensity ([Figure 4B](#)).

We also considered a different parameterization of the Out-of-Africa bottleneck.²⁹ In this model, a constant population ($N_e = 10,000$) reduced its size to 7.57% 118 ka ago and recovered to the original population size after 10 ka. Patterns of neutral and deleterious variation in this bottleneck model with a quick recovery were similar to that described above. For instance, fewer than 3% of deleterious variants in the contemporary population arose before or during the bottleneck, but they comprised 47.8% of derived deleterious alleles carried by individuals and were of modest effect size ($10^{-4} \leq |s| < 0.01$). The mean density (proportion) of derived deleterious alleles per individual in this bottleneck model was 0.0176/kb (3.29%), compared to 0.0158/kb (3.00%) in constant-sized populations.

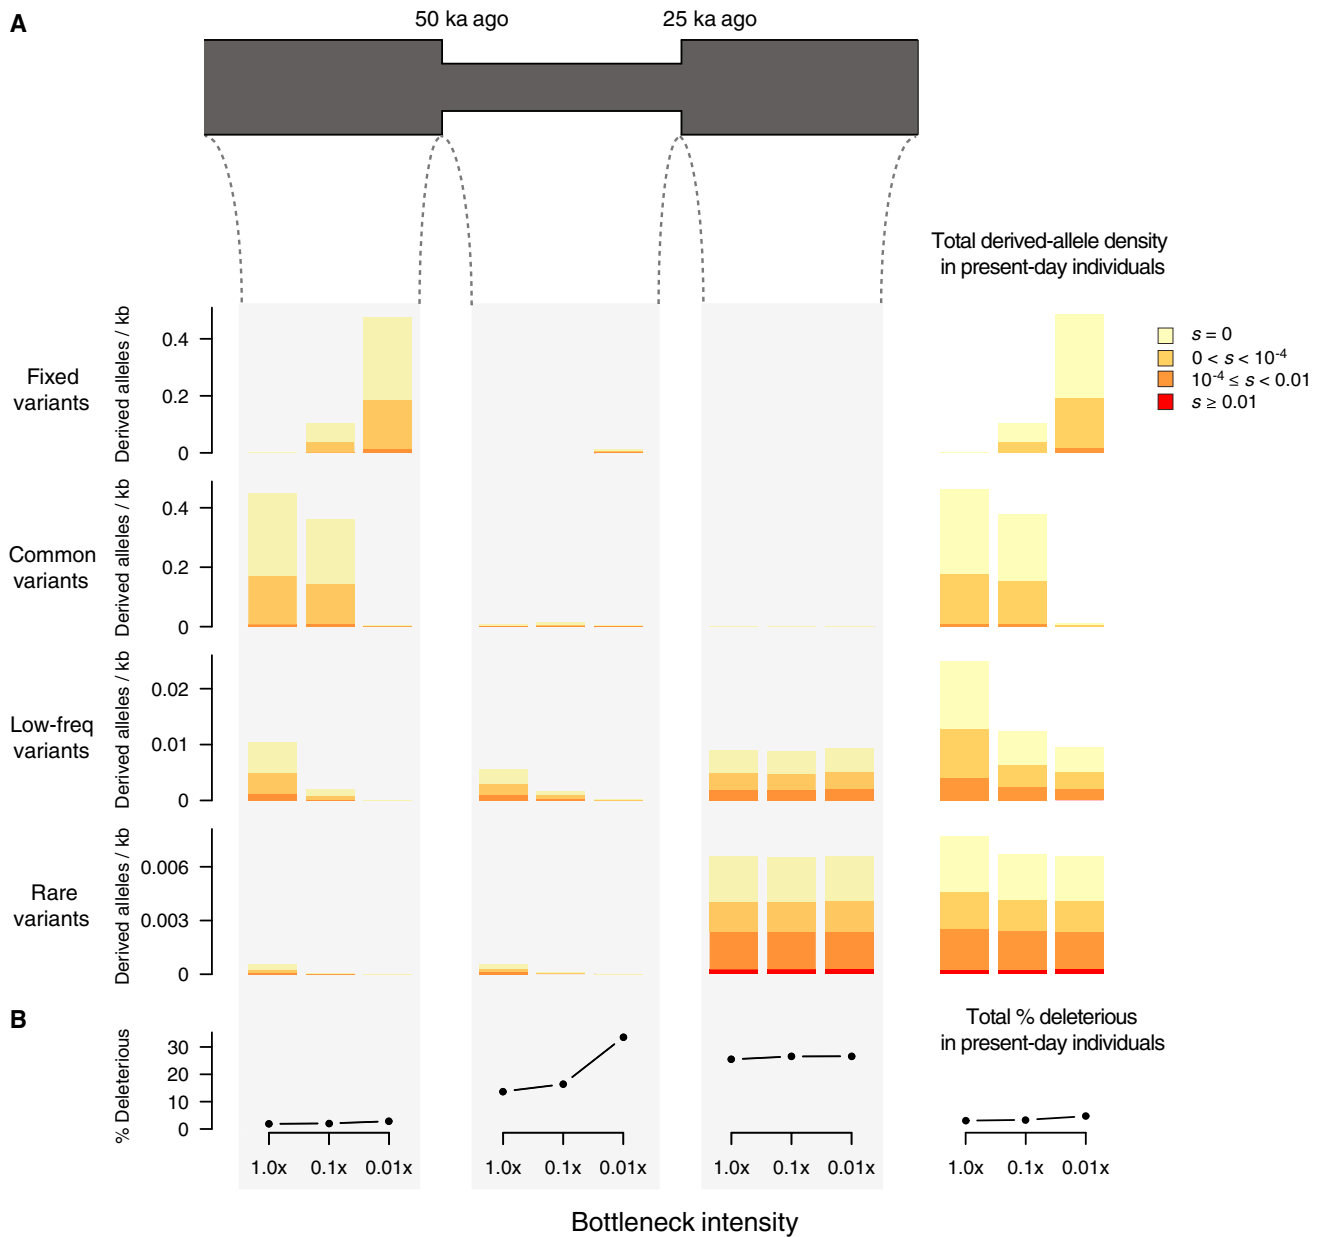


Figure 4. Temporal Decomposition of Neutral and Deleterious Variation among Present-Day Individuals in Bottleneck Models

(A) Diagram of the simulated bottleneck model, in which population size decreased to $0.1\times$ or $0.01\times$ 50 ka ago and then recovered to prebottleneck levels with a size of 10,000 25 ka ago. Bar plots on the far right show the average number of derived alleles per individual per kilobase in present-day individuals as a function of DAF (fixed DAF = 1, common DAF ≥ 0.05 , low-frequency $0.01 \leq$ DAF < 0.05 , and rare DAF < 0.01) and selection coefficients. We decomposed variant density in present-day individuals according to when the mutation arose (before, during, or after the bottleneck), as shown in the bar plots under the demographic model. Thus, we can obtain variant density of present-day individuals for each DAF category (rows) by simply adding variant density across the three time epochs considered.

(B) The proportion of derived deleterious alleles ($|s| \geq 10^{-4}$) per individual in present-day individuals and as a function of when the mutation arose.

Temporal Dynamics of Deleterious Alleles in Recent-Growth Models

To better understand how recent growth has influenced the burden of deleterious alleles among individuals, we simulated protein-coding variation in models where populations started expanding 5 ka ago (with growth rates of 2% or 3%) and compared patterns of neutral and deleterious derived alleles per individual to those of constant-size pop-

ulation models (Figure 5A). Note that a growth rate of 2% approximately corresponds to that estimated by Tennesen et al.¹² for European populations. We also considered a higher growth rate because it is likely that recent studies have underestimated this parameter.^{43–46}

As expected, population expansions led to a significant skew toward rare variation in the population—61.3%, 85.4%, and 87.7% of all SNVs were rare in constant-size,




Figure 5. Temporal Decomposition of Neutral and Deleterious Variation among Present-Day Individuals in Recent-Growth Models (A) Diagram of simulated models of recent accelerated growth. Here, populations started growing at a rate of 2% or 3% per generation from a constant population size of 10,000 individuals 5 ka ago. As in Figure 4, the far right column shows the mean number of derived alleles per individual per kilobase in present-day individuals as a function of DAF (fixed DAF = 1, common DAF ≥ 0.05 , low-frequency $0.01 \leq \text{DAF} < 0.05$, and rare DAF < 0.01) and selection coefficients. Similarly, we decomposed variant density in present-day individuals according to when the mutation arose (before or during growth).

(B) The proportion of deleterious alleles ($|s| \geq 10^{-4}$) per individual in present-day individuals and as a function of when the mutation arose.

2% growth, and 3% growth models, respectively. In population-expansion models, a significant amount of rare variation arose during the growth epoch. At the level of individuals, however, recent population expansions had much more subtle consequences. For example, the average density of rare derived alleles per individual per kilobase in constant-size, 2% growth, and 3% growth models was 0.0076, 0.0103, and 0.0106, respectively (Figure 5A). Furthermore, the average density of rare deleterious ($|s| \geq 10^{-4}$) alleles per individual in constant-size, 2% growth, and 3% growth models was 0.0024, 0.0032,

and 0.0034, respectively (Figure 5A). Note that in both constant-sized and growth models, nearly all strongly deleterious variants ($|s| > 0.01$) arose recently (Figure 5A).

The proportion of derived deleterious alleles per individual was slightly lower in expanding populations than in constant-sized populations (3.20%, 3.15%, and 3.18% for constant-sized, 2% growth, and 3% growth models, respectively; Figure 5B). Thus, as predicted by population-genetics theory,⁴⁷ population growth allows deleterious alleles to be purged more efficiently. Deleterious mutations have a lower probability of fixation in an exponentially

growing population than in a constant-sized population; also, most deleterious mutations will eventually be lost,⁴⁸ predicting a proportional reduction of deleterious mutations at the individual level. However, for mutations that reach fixation, the expected fixation time increases with population size. Weakly deleterious mutations can segregate longer at lower frequency in the population and result in a proportional increase in deleterious variants at the population level. For example, in the simulated data, the proportion of deleterious variants in the population increased from 30.2% in populations with no growth to 36.4% and 35.8% in the 2% and 3% growth models, respectively, consistent with previous observations.⁴⁹

We note that the patterns of deleterious variation among individuals and populations in our simulations did not change monotonously as a function of growth rate. This was most likely a consequence of sampling variation, given that the effects of the growth rates considered here (2% versus 3%) on patterns of variation were quantitatively similar in the short amount of time (~200 generations) that elapsed since the onset of population growth.

Temporal Dynamics of Deleterious Alleles in Complicated Demographic Models

The results described above allowed us to isolate the effects of bottlenecks and population expansions on patterns of deleterious variation. Here, we consider the combined effects of bottlenecks, recent accelerated growth, population structure, and admixture (Figure 6A) by using previously inferred parameters for the EA and AA samples by Tennessen et al.¹²

On average, EA individuals carried more fixed and common deleterious alleles than did AA individuals (0.0140/kb and 0.0076/kb in EA and AA individuals, respectively; Figure 6A); all of them were of modest effect size (i.e., $10^{-4} \leq |s| < 0.01$), and over 95% arose ≥ 23 ka ago. In contrast, AA individuals carried on average more rare and low-frequency derived deleterious alleles than did EA individuals (0.0043/kb and 0.0087/kb in EA and AA individuals, respectively; Figure 6A) because of their larger effective population size between 5.1 and 51 ka ago. In both EA and AA individuals, all of the deleterious mutations with strong effect ($|s| \geq 0.01$) were rare in the population, and the majority (96.1% in EA and 95.8% in AA) arose in the last 5.1 ka (Figure 6A), consistent with previous simulation results.⁵⁰ Overall, EA individuals carried on average more deleterious alleles per individual than did AA individuals (0.0182/kb and 0.0164/kb in EA and AA individuals, respectively; Figure 6A), in agreement with empirical observations in the empirical data (Figure 3).

Consistent with the expected effects of bottlenecks and recent accelerated growth as described above, the proportion of derived deleterious alleles per individual exhibited differences between simulated EA and AA individuals in different time epochs (Figure 6B). Specifically, for variants that arose before or during the European bottlenecks, the proportion of derived deleterious alleles per individual

was higher in EA than in AA individuals (Figure 6B). However, for variants that arose during recent accelerated population growth, the proportion of derived deleterious alleles per individual was slightly lower in EA than in AA individuals (Figure 6B). Overall, the average proportion of derived deleterious alleles in present-day individuals in the simulated data was 4.39% for EA individuals and 3.96% for AA individuals (Figure 6B).

Opposing Patterns of Variation in Populations and Individuals

The temporal decomposition of present-day variation into different time epochs (Figures 4, 5, and 6) intimates that patterns of deleterious variation manifest differently depending on whether the focus is on populations or individuals. To more clearly articulate this important point, we calculated the density and proportion of deleterious variants and derived alleles in the simulated data for populations and individuals.

At the population level, the density of deleterious variants in bottleneck models decreased from 1.01/kb in constant-size models to 0.97/kb and 0.94/kb in bottlenecks of intensity $0.1\times$ and $0.01\times$, respectively (Figure 7A). However, in individuals, the density of derived deleterious alleles increased from 0.0151/kb in constant-size models to 0.0167/kb and 0.0240/kb in bottlenecks of intensity $0.1\times$ and $0.01\times$, respectively (Figure 7B). Note that in the bottleneck model with a $0.1\times$ decrease in size, this corresponds to approximately 10.2% more derived deleterious alleles per individual than in constant-size models. This increase was largely due to the bottleneck-mediated shift in the SFS toward higher allele frequencies and, in some cases, fixation.

Similarly, recent population growth led to opposing density patterns of derived deleterious alleles in populations and individuals. Recent growth led to a substantial increase in the density of deleterious variants in populations, largely because of SNVs arising during the period of recent growth (Figure 7A). For example, the density of deleterious variants in populations increased from 1.01/kb in constant-size models to 3.18/kb and 3.75/kb in models with 2% and 3% growth, respectively (Figure 7A). Conversely, the average density of derived deleterious alleles in individuals decreased slightly from 0.0160/kb in constant-sized populations to 0.0158/kb and 0.0159/kb in models with 2% and 3% growth, respectively (Figure 7B). In the 2% growth model, this corresponds to approximately 1.5% fewer deleterious alleles per individual than in constant-size models. These results are consistent with Gazave et al.,⁵¹ who found the same trend of growth decreasing the density of deleterious alleles per individual for variants with a selection coefficient of $|s| \geq 10^{-4}$.

We further compared the load of deleterious mutations in constant, bottleneck, and recent-growth models (Figure 7C) by assuming additive and independent effects among deleterious alleles. In the bottleneck model, load was modestly higher than in constant-size populations.

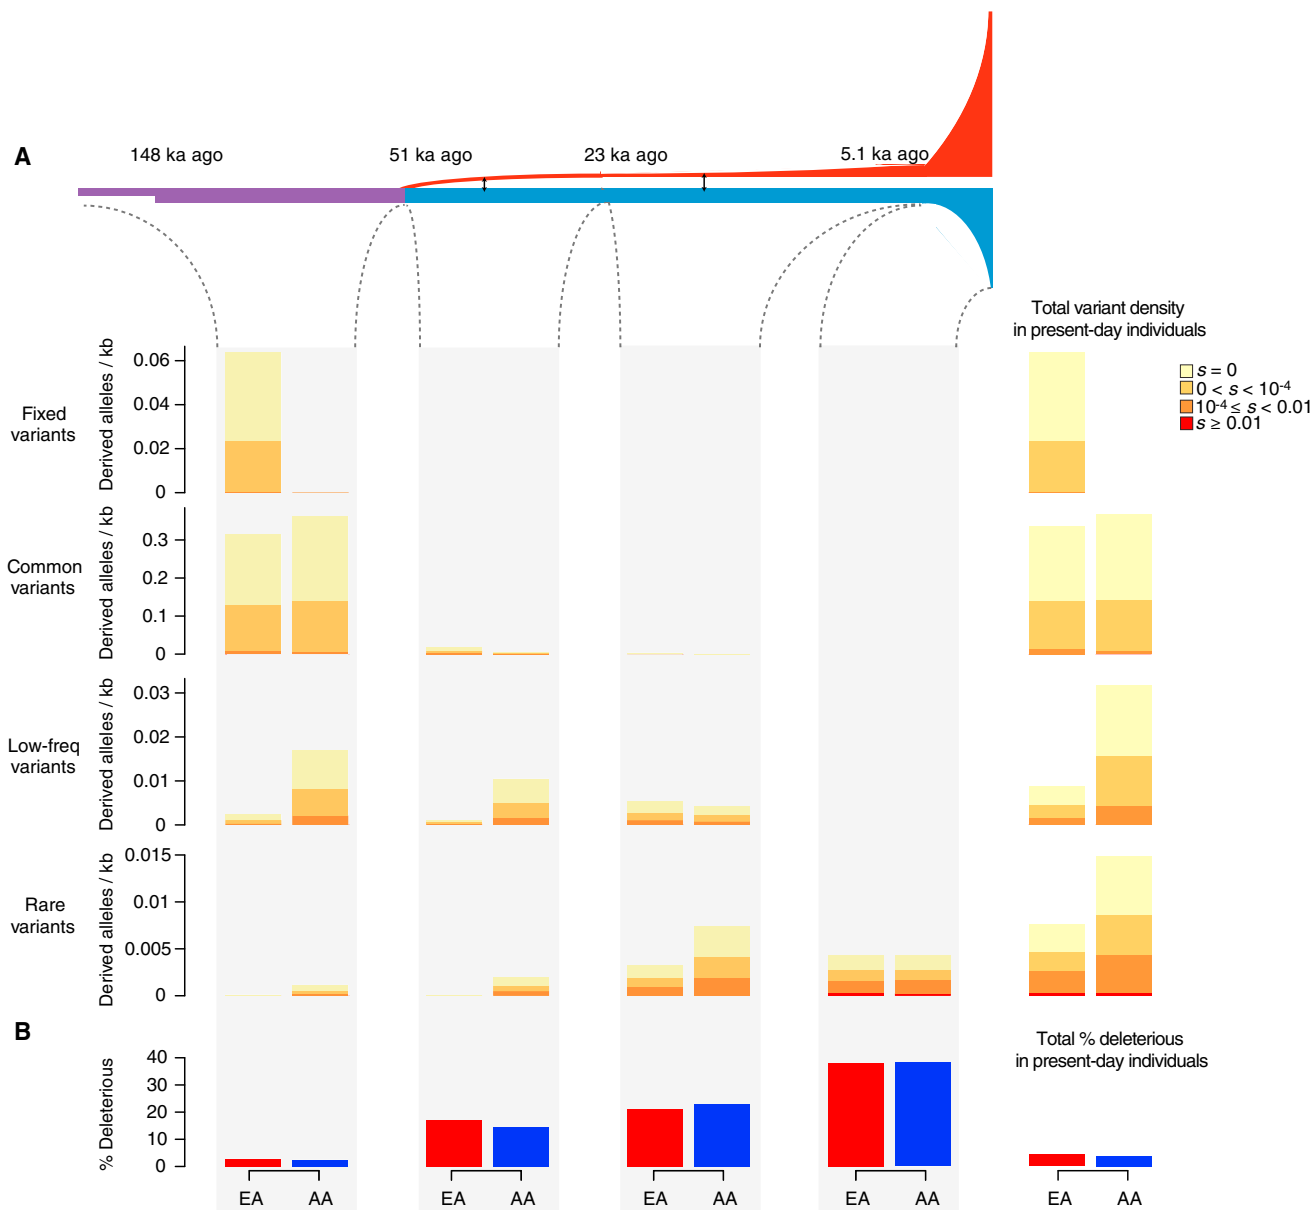


Figure 6. Temporal Decomposition of Neutral and Deleterious Variation among Present-Day Individuals in a More Realistic Demographic Model

(A) Diagram of a more realistic demographic model for EA and AA populations. This model involves multiple bottlenecks in the EA lineage, recent accelerated growth, and admixture as inferred by Tennesen et al.¹² As in Figure 4, the far right column shows the mean number of derived alleles per individual per kilobase in present-day individuals as a function of DAF (fixed DAF = 1, common DAF ≥ 0.05 , low-frequency $0.01 \leq \text{DAF} < 0.05$, and rare DAF < 0.01) and selection coefficients. Similarly, we decomposed variant density in present-day individuals according to when the mutation arose (before the Out-of-Africa bottleneck, during the bottleneck, during the initial period of growth in EA individuals, or during the recent accelerated growth in both EA and AA individuals).

(B) The proportion of deleterious alleles ($|s| \geq 10^{-4}$) per individual in present-day individuals and as a function of when the mutation arose.

For example, in bottlenecks with a $0.1\times$ intensity (corresponding to the Out-of-Africa event), load was on average 4.5% higher than in constant-size populations. It is important to note that the accumulation of fixed deleterious variants contributed to 2.2% of mutation load in this model but only 0.1% in the constant-sized model. In models with a bottleneck intensity of $0.01\times$, fixed deleterious variants accounted for 43.2% of the mutation load.

Thus, although the increase in load was modest, it had significant evolutionary implications, given that fixed variants cannot be purged from the population. In contrast, recent population expansions had negligible effects on load (Figure 7C), consistent with previous studies.^{11,51}

Finally, for comparison, we also summarized the average density of deleterious variants and alleles at the population and individual levels, as well as mutation load, for

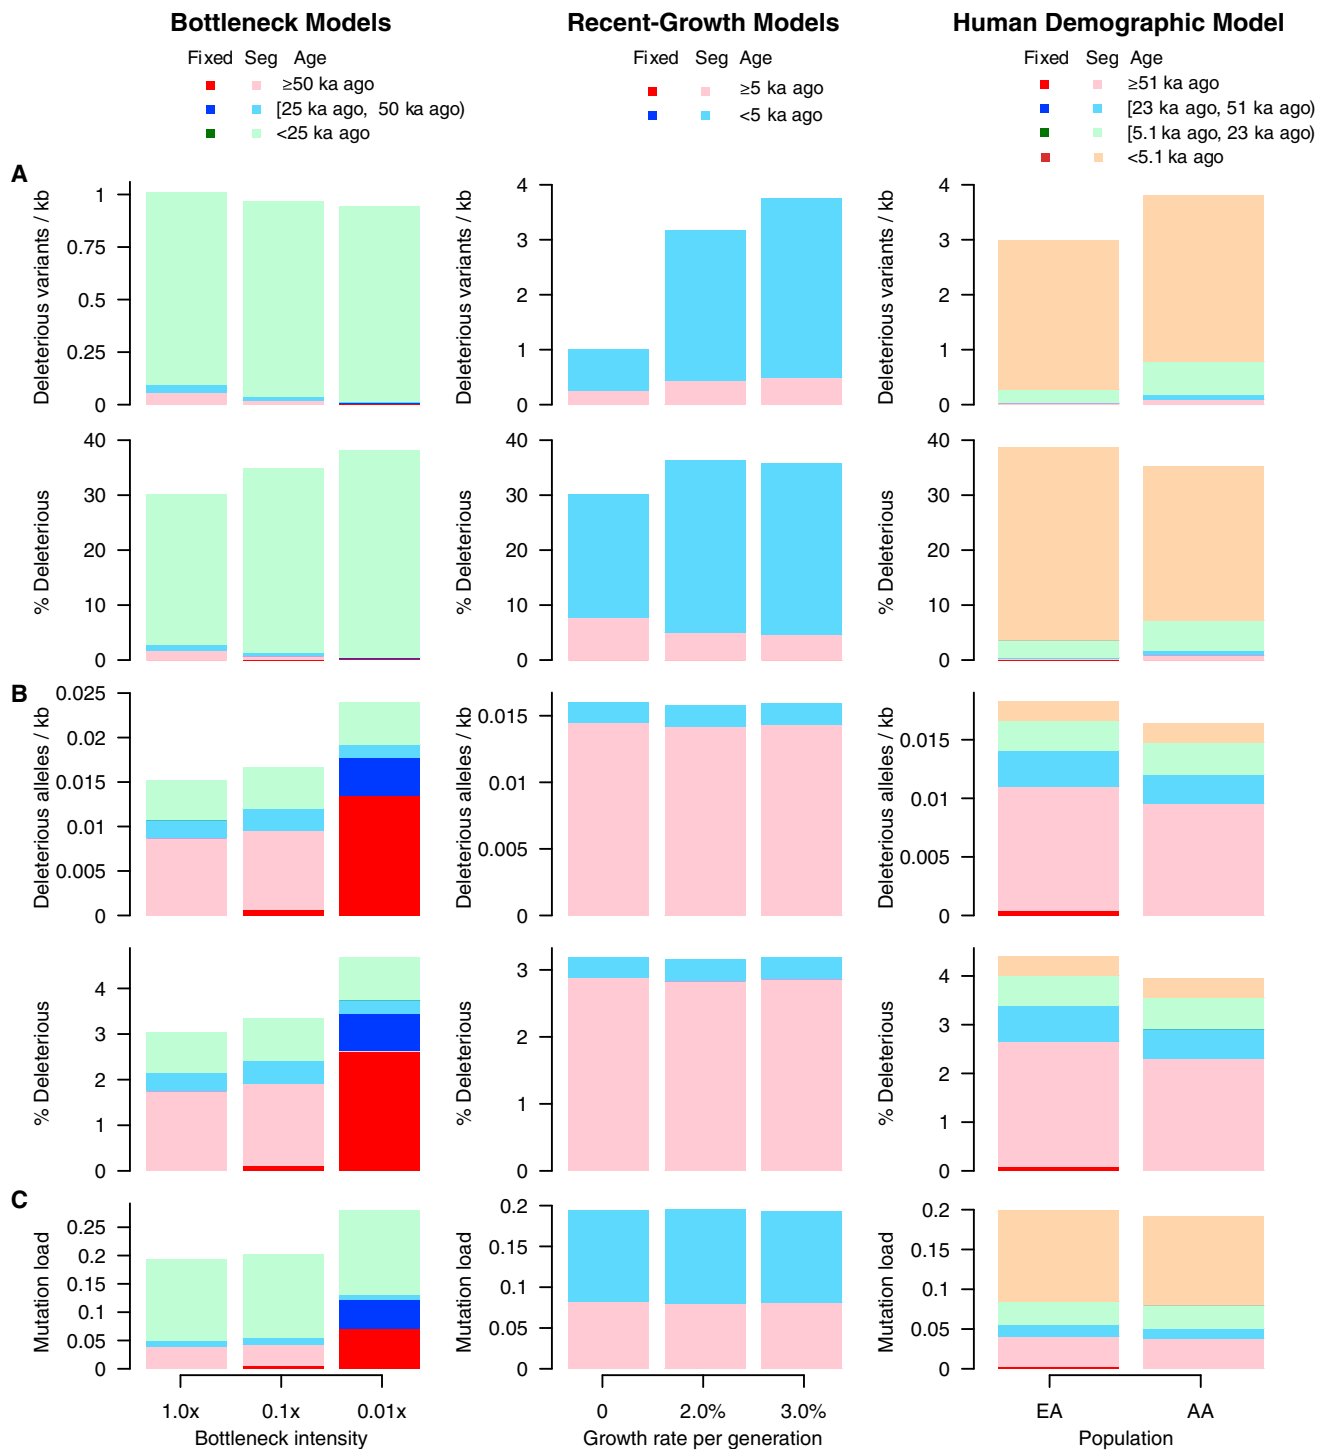


Figure 7. Oposing Patterns of Deleterious SNVs and Alleles in Individuals and Populations

(A) The number and proportion of deleterious variants per kilobase in populations as a function of mutation age and DAF. (B) The mean number and proportion of derived deleterious alleles per individual per kilobase as a function of mutation age and DAF. (C) The mutation load as a function of mutation age and DAF. Note that the Tennesen et al.¹² model refers to the demographic model shown in Figure 6. Variants are colored on the basis of whether they are fixed or segregating (“seg”) and when they arose (“age”).

the more realistic demographic model of EA and AA populations. As expected, for equal sample sizes, more variants were observed in the simulated AA population than in the EA population (Figure 7A). The average density of derived

deleterious alleles in EA individuals was 0.0182/kb, which is approximately 11.3% higher than that for AA individuals (0.0164/kb; Figure 7B). Notably, 21.4% of the difference in the density of derived deleterious alleles was

attributable to the increase in fixed deleterious mutations in EA individuals (Figure 7B). As a result, mutation load is predicted to be slightly inflated in EA individuals (20.0% in EA individuals versus 19.2% in AA individuals; Figure 7C).

Discussion

Recent large-scale sequencing studies have afforded unique insights into the spectrum of human genomic variation.^{9–12,19,45,51} However, these studies have emphasized different aspects of the data and, in some cases, appear to have come to conflicting conclusions, particularly for deleterious variants.^{10–12,19} To better understand empirical patterns of neutral and deleterious variation, we performed a comprehensive analysis of a large exome sequencing data set. Furthermore, we leveraged extensive simulations on the dynamics of neutral and deleterious alleles in nonequilibrium populations to facilitate interpretations of the exome data. Importantly, we developed an approach to mitigate the previously described strong reference-bias effects observed in nearly all methods of defining deleterious variation, which should be useful in a wide variety of applications.

Our empirical and simulation results show that characteristics of neutral and deleterious variation can result in opposing patterns depending on how variation is summarized. For example, the total number of SNVs in the AA population is significantly higher than that in the EA population, and AA individuals carry significantly more SNVs on average than do EA individuals. However, in EA individuals, the strongly skewed SFS due to population bottlenecks caused some SNVs to be lost and others to drift to higher frequency. Thus, the composition of genotypes in the EA and AA samples is considerably different (Figures 1 and 3)—EA individuals have on average more homozygous derived sites, and AA individuals have more heterozygous sites. Therefore, although there can be large differences in the average number of SNVs per individual between populations, the average number of derived alleles is approximately equal.

Furthermore, opposing patterns of variation can also manifest depending on whether the focus is on individuals or populations, which is the primary factor accounting for seemingly disparate empirical characteristics of protein-coding variation in recent studies. For instance, although recent accelerated growth has profoundly influenced protein-coding variation in populations, it has thus far had limited impact on individuals. For example, 80.8% and 73.4% of deleterious SNVs in the EA and AA populations, respectively, have a population DAF < 0.05%, but they comprise only 2.0% and 2.2% of derived deleterious alleles carried by EA and AA individuals, respectively. In contrast, although only 0.20% and 0.06% of deleterious variants are nearly fixed or fixed (DAF ≥ 99.9%) in the EA and AA populations, respec-

tively, they comprise 19.4% and 6.7% of derived deleterious alleles carried by EA individuals and AA individuals, respectively. Thus, although most SNVs are rare in the population (i.e., have a DAF < 0.05%), most of the variants carried by individuals are common.

Our simulation results also show that population bottlenecks and expansions have opposing effects on patterns of variation in populations and individuals. Compared to populations of constant size, bottlenecks reduce, whereas recent growth increases, the number of neutral and deleterious variants in the population. Similarly, bottlenecks and expansions have opposite effects on the SFS—the former leads to a skew toward common variation, and the latter results in a skew toward rare variation.^{47,48,52} As a result, compared with a constant-size population, individuals from populations that have experienced bottlenecks tend to carry more deleterious alleles, whereas individuals from expanding populations carry slightly fewer derived deleterious alleles.

It is of interest to note that in the bottleneck simulations, we also considered a severe reduction of population size to 0.01× during the bottleneck. In this case, 0.93% of deleterious variants were fixed in the population, which contributed to a substantial proportion (73.3%) of derived deleterious alleles carried by individuals. Although this is more extreme than the estimated intensity of the Out-of-Africa bottleneck, it might be relevant for some populations that experienced stronger founder effects.^{53–56}

Our simulation results recapitulate many of the qualitative patterns of neutral and deleterious variation observed in the empirical data. For example, the Tennesen et al.¹² demographic model (Figure 6) predicts that EA individuals will on average carry more derived deleterious alleles than will AA individuals. Furthermore, the simulations also show that the proportion of derived deleterious alleles per individual is predicted to be higher in EA than in AA individuals. Despite the qualitative agreement between the empirical and simulated data, we do observe quantitative differences. Specifically, simulations suggest that the average density of derived deleterious alleles in EA individuals is ~11.3% higher than in AA individuals, whereas in the empirical data, it is only ~1.4% higher in EA individuals. A number of factors most likely contribute to this quantitative difference, including differences in demography and distribution of mutational effect sizes in real and simulated data. Another important difference is that in simulations, deleterious variants can be identified precisely, whereas in empirical data they are inferred with considerable error. Indeed, a substantial amount of the higher density of deleterious alleles in EA individuals in the simulated data is attributable to weakly deleterious mutations ($|s| \approx 10^{-4}$), and it is probably more difficult to identify these in empirical data. Thus, developing methods for predicting the functional and evolutionary significance of human genetic variation remains an important endeavor.

Recently, Simons et al.¹¹ used allele-frequency data from the same data set and found that the average DAF of deleterious variants was not significantly different between EA and AA populations. Overall, our empirical results are qualitatively similar and in broad agreement with their findings. However, we observed that the average number of deleterious alleles was slightly but significantly higher in EA individuals than in AA individuals (905.7 and 893.2 for EA and AA individuals, respectively; Mann-Whitney test, $p < 10^{-15}$; Figure 3). These differences in the average number of deleterious alleles per individual and distribution of effect sizes observed in the simulated models suggest that there might be subtle differences in mutation load between EA and AA populations and that they might be governed by mutations with small fitness effects. Simons et al. also observed that mutation load was slightly higher in EA populations than in AA populations (see their Figure S10¹¹) as a result of mutations of weak effect. However, we note that the distinct genotypic composition between EA and AA individuals suggests that differences in mutation load might be larger if deleterious alleles are dominant or recessive.¹¹ Furthermore, variability in mutation load might be greater in populations that experienced more intense bottlenecks (Figures 4 and 7). More generally, as eloquently noted by Lohmueller,⁵⁷ homogeneity of load between populations does not imply homogeneity in patterns of deleterious variation, as illustrated in our empirical and simulation data (Figure 7). Finally, although our simulation and empirical results demonstrate that patterns of deleterious SNVs and alleles vary among populations, these differences are small in relation to the magnitude of variability among individuals. Furthermore, the majority of deleterious alleles carried by individuals are predicted to have selection coefficients on the order of $|s| \approx 10^{-4}$. Although such weakly deleterious alleles are important over evolutionary timescales, their contribution to the burden of human disease is largely unknown. Thus, it would be imprudent to directly relate differences in patterns of evolutionarily defined deleterious variation and mutation load to differences in disease burden. Nonetheless, recent demographic history does influence the characteristics and spectrum of rare and common variation within and between populations, which has important implications in the design and interpretation of disease-mapping studies.⁴⁹ More generally, the continued development of experimental and computational methods of identifying and predicting the functional consequences of genetic variation is ultimately the most direct and accurate means of assessing individual disease risk.

Acknowledgments

We thank members of the J.M.A. laboratory, Kirk E. Lohmueller, Jonathan K. Pritchard, Guy Sella, Carlos D. Bustamante, Shamil R. Sunyaev, and David Reich for helpful feedback and discussions related to this work.

Received: June 2, 2014

Accepted: September 11, 2014

Published: October 2, 2014

Web Resources

The URLs for data presented herein are as follows:

1000 Genomes, <http://www.1000genomes.org/>
 BEDOPS, <https://code.google.com/p/bedops/>
 Ensembl Genome Browser, <http://useast.ensembl.org/index.html>
 mafTools, <https://github.com/dentearl/mafTools>
 NHLBI Exome Sequencing Project (ESP) Exome Variant Server, <http://evs.gs.washington.edu/EVS/>
 PHAST, <http://compgen.bscb.cornell.edu/phast/>
 Phylo_{NH}, <http://akeylab.gs.washington.edu/downloads.shtml>
 SFS_CODE, http://sfscode.sourceforge.net/SFS_CODE/index/index.html
 UCSC Genome Browser, <http://genome.ucsc.edu>

References

- Haldane, J.B.S. (1937). The effect of variation on fitness. *Am. Nat.* 71, 337–349.
- Crow, J.F. (1970). Genetic Loads and the Cost of Natural Selection. In *Mathematical Topics in Population Genetics*, K.-i. Kojima, ed. (New York: Springer-Verlag), pp. 128–177.
- Crow, J.F. (2000). The origins, patterns and implications of human spontaneous mutation. *Nat. Rev. Genet.* 1, 40–47.
- Gavrilov, L.A., Gavrilova, N.S., Kroutko, V.N., Evdokushkina, G.N., Semyonova, V.G., Gavrilova, A.L., Lapshin, E.V., Evdokushkina, N.N., and Kushnareva, Y.E. (1997). Mutation load and human longevity. *Mutat. Res.* 377, 61–62.
- Lynch, M. (2010). Rate, molecular spectrum, and consequences of human mutation. *Proc. Natl. Acad. Sci. USA* 107, 961–968.
- Kimura, M., Maruyama, T., and Crow, J.F. (1963). The Mutation Load in Small Populations. *Genetics* 48, 1303–1312.
- Ohta, T. (1973). Slightly deleterious mutant substitutions in evolution. *Nature* 246, 96–98.
- Peischl, S., Dupanloup, I., Kirkpatrick, M., and Excoffier, L. (2013). On the accumulation of deleterious mutations during range expansions. *Mol. Ecol.* 22, 5972–5982.
- Abecasis, G.R., Auton, A., Brooks, L.D., DePristo, M.A., Durbin, R.M., Handsaker, R.E., Kang, H.M., Marth, G.T., and McVean, G.A.; 1000 Genomes Project Consortium (2012). An integrated map of genetic variation from 1,092 human genomes. *Nature* 491, 56–65.
- Lohmueller, K.E., Indap, A.R., Schmidt, S., Boyko, A.R., Hernandez, R.D., Hubisz, M.J., Sninsky, J.J., White, T.J., Sunyaev, S.R., Nielsen, R., et al. (2008). Proportionally more deleterious genetic variation in European than in African populations. *Nature* 451, 994–997.
- Simons, Y.B., Turchin, M.C., Pritchard, J.K., and Sella, G. (2014). The deleterious mutation load is insensitive to recent population history. *Nat. Genet.* 46, 220–224.
- Tennessen, J.A., Bigham, A.W., O'Connor, T.D., Fu, W., Kenny, E.E., Gravel, S., McGee, S., Do, R., Liu, X., Jun, G., et al.; Broad GO; Seattle GO; NHLBI Exome Sequencing Project (2012). Evolution and functional impact of rare coding variation from deep sequencing of human exomes. *Science* 337, 64–69.

13. Hodgkinson, A., Casals, F., Idaghdour, Y., Grenier, J.C., Hernandez, R.D., and Awadalla, P. (2013). Selective constraint, background selection, and mutation accumulation variability within and between human populations. *BMC Genomics* *14*, 495.
14. Fu, W., O'Connor, T.D., Jun, G., Kang, H.M., Abecasis, G., Leal, S.M., Gabriel, S., Rieder, M.J., Altshuler, D., Shendure, J., et al.; NHLBI Exome Sequencing Project (2013). Analysis of 6,515 exomes reveals the recent origin of most human protein-coding variants. *Nature* *493*, 216–220.
15. Kidd, J.M., Gravel, S., Byrnes, J., Moreno-Estrada, A., Musharoff, S., Bryc, K., Degenhardt, J.D., Brisbin, A., Sheth, V., Chen, R., et al. (2012). Population genetic inference from personal genome data: impact of ancestry and admixture on human genomic variation. *Am. J. Hum. Genet.* *91*, 660–671.
16. Subramanian, S. (2012). The abundance of deleterious polymorphisms in humans. *Genetics* *190*, 1579–1583.
17. Casals, F., Hodgkinson, A., Hussin, J., Idaghdour, Y., Bruat, V., de Maillard, T., Grenier, J.C., Gbeha, E., Hamdan, F.F., Girard, S., et al. (2013). Whole-exome sequencing reveals a rapid change in the frequency of rare functional variants in a founding population of humans. *PLoS Genet.* *9*, e1003815.
18. Szpiech, Z.A., Xu, J., Pemberton, T.J., Peng, W., Zöllner, S., Rosenberg, N.A., and Li, J.Z. (2013). Long runs of homozygosity are enriched for deleterious variation. *Am. J. Hum. Genet.* *93*, 90–102.
19. Do, R., Balick, D., Li, H., Adzhubei, I., Sunyaev, S., and Reich, D. (2014). No evidence that natural selection has been less effective at removing deleterious mutations in Europeans than in West Africans. *arXiv*, arXiv:1402.4896, <http://arxiv.org/abs/1402.4896>.
20. Abecasis, G.R., Altshuler, D., Auton, A., Brooks, L.D., Durbin, R.M., Gibbs, R.A., Hurles, M.E., and McVean, G.A.; 1000 Genomes Project Consortium (2010). A map of human genome variation from population-scale sequencing. *Nature* *467*, 1061–1073.
21. Hernandez, R.D., Williamson, S.H., and Bustamante, C.D. (2007). Context dependence, ancestral misidentification, and spurious signatures of natural selection. *Mol. Biol. Evol.* *24*, 1792–1800.
22. Hwang, D.G., and Green, P. (2004). Bayesian Markov chain Monte Carlo sequence analysis reveals varying neutral substitution patterns in mammalian evolution. *Proc. Natl. Acad. Sci. USA* *101*, 13994–14001.
23. Pennacchio, L.A., Olivier, M., Hubacek, J.A., Cohen, J.C., Cox, D.R., Fruchart, J.C., Krauss, R.M., and Rubin, E.M. (2001). An apolipoprotein influencing triglycerides in humans and mice revealed by comparative sequencing. *Science* *294*, 169–173.
24. Boffelli, D., McAuliffe, J., Ovcharenko, D., Lewis, K.D., Ovcharenko, I., Pachter, L., and Rubin, E.M. (2003). Phylogenetic shadowing of primate sequences to find functional regions of the human genome. *Science* *299*, 1391–1394.
25. Pollard, K.S., Hubisz, M.J., Rosenbloom, K.R., and Siepel, A. (2010). Detection of nonneutral substitution rates on mammalian phylogenies. *Genome Res.* *20*, 110–121.
26. Neph, S., Kuehn, M.S., Reynolds, A.P., Haugen, E., Thurman, R.E., Johnson, A.K., Rynes, E., Maurano, M.T., Vierstra, J., Thomas, S., et al. (2012). BEDOPS: high-performance genomic feature operations. *Bioinformatics* *28*, 1919–1920.
27. Hernandez, R.D. (2008). A flexible forward simulator for populations subject to selection and demography. *Bioinformatics* *24*, 2786–2787.
28. Gravel, S., Henn, B.M., Gutenkunst, R.N., Indap, A.R., Marth, G.T., Clark, A.G., Yu, F., Gibbs, R.A., and Bustamante, C.D.; 1000 Genomes Project (2011). Demographic history and rare allele sharing among human populations. *Proc. Natl. Acad. Sci. USA* *108*, 11983–11988.
29. Keinan, A., Mullikin, J.C., Patterson, N., and Reich, D. (2007). Measurement of the human allele frequency spectrum demonstrates greater genetic drift in East Asians than in Europeans. *Nat. Genet.* *39*, 1251–1255.
30. Boyko, A.R., Williamson, S.H., Indap, A.R., Degenhardt, J.D., Hernandez, R.D., Lohmueller, K.E., Adams, M.D., Schmidt, S., Sinsky, J.J., Sunyaev, S.R., et al. (2008). Assessing the evolutionary impact of amino acid mutations in the human genome. *PLoS Genet.* *4*, e1000083.
31. Morton, N.E., Crow, J.F., and Muller, H.J. (1956). An Estimate of the Mutational Damage in Man from Data on Consanguineous Marriages. *Proc. Natl. Acad. Sci. USA* *42*, 855–863.
32. Crow, J.F. (1958). Some possibilities for measuring selection intensities in man. *Hum. Biol.* *30*, 1–13.
33. Muller, H.J. (1950). Our load of mutations. *Am. J. Hum. Genet.* *2*, 111–176.
34. Nickerson, D.A., Taylor, S.L., Weiss, K.M., Clark, A.G., Hutchinson, R.G., Stengård, J., Salomaa, V., Vartiainen, E., Boerwinkle, E., and Sing, C.F. (1998). DNA sequence diversity in a 9.7-kb region of the human lipoprotein lipase gene. *Nat. Genet.* *19*, 233–240.
35. Campbell, M.C., and Tishkoff, S.A. (2008). African genetic diversity: implications for human demographic history, modern human origins, and complex disease mapping. *Annu. Rev. Genomics Hum. Genet.* *9*, 403–433.
36. Kircher, M., Witten, D.M., Jain, P., O'Roak, B.J., Cooper, G.M., and Shendure, J. (2014). A general framework for estimating the relative pathogenicity of human genetic variants. *Nat. Genet.* *46*, 310–315.
37. Chun, S., and Fay, J.C. (2009). Identification of deleterious mutations within three human genomes. *Genome Res.* *19*, 1553–1561.
38. Davydov, E.V., Goode, D.L., Sirota, M., Cooper, G.M., Sidow, A., and Batzoglou, S. (2010). Identifying a high fraction of the human genome to be under selective constraint using GERP++. *PLoS Comput. Biol.* *6*, e1001025.
39. Adzhubei, I.A., Schmidt, S., Peshkin, L., Ramensky, V.E., Gerasimova, A., Bork, P., Kondrashov, A.S., and Sunyaev, S.R. (2010). A method and server for predicting damaging missense mutations. *Nat. Methods* *7*, 248–249.
40. Schwarz, J.M., Rödelsperger, C., Schuelke, M., and Seelow, D. (2010). MutationTaster evaluates disease-causing potential of sequence alterations. *Nat. Methods* *7*, 575–576.
41. Kumar, P., Henikoff, S., and Ng, P.C. (2009). Predicting the effects of coding non-synonymous variants on protein function using the SIFT algorithm. *Nat. Protoc.* *4*, 1073–1081.
42. Lachance, J., Vernot, B., Elbers, C.C., Ferwerda, B., Froment, A., Bodo, J.M., Lema, G., Fu, W., Nyambo, T.B., Rebbeck, T.R., et al. (2012). Evolutionary history and adaptation from high-coverage whole-genome sequences of diverse African hunter-gatherers. *Cell* *150*, 457–469.
43. Coventry, A., Bull-Otterson, L.M., Liu, X., Clark, A.G., Maxwell, T.J., Crosby, J., Hixson, J.E., Rea, T.J., Muzny, D.M., Lewis, L.R., et al. (2010). Deep resequencing reveals excess rare recent variants consistent with explosive population growth. *Nat. Commun.* *1*, 131.

44. Keinan, A., and Clark, A.G. (2012). Recent explosive human population growth has resulted in an excess of rare genetic variants. *Science* 336, 740–743.
45. Nelson, M.R., Wegmann, D., Ehm, M.G., Kessner, D., St Jean, P., Verzilli, C., Shen, J., Tang, Z., Bacanu, S.A., Fraser, D., et al. (2012). An abundance of rare functional variants in 202 drug target genes sequenced in 14,002 people. *Science* 337, 100–104.
46. United Nations, Department of Economic and Social Affairs, Population Division (2011). World Population Prospects: The 2010 Revision, Volume I: Comprehensive Tables. ST/ESA/SERA/313. http://esa.un.org/wpp/documentation/pdf/WPP2010_Volume-I_Comprehensive-Tables.pdf.
47. Kimura, M. (1955). Stochastic processes and distribution of gene frequencies under natural selection. *Cold Spring Harb. Symp. Quant. Biol.* 20, 33–53.
48. Waxman, D. (2011). A unified treatment of the probability of fixation when population size and the strength of selection change over time. *Genetics* 188, 907–913.
49. Lohmueller, K.E. (2014). The impact of population demography and selection on the genetic architecture of complex traits. *PLoS Genet.* 10, e1004379.
50. Maher, M.C., Uricchio, L.H., Torgerson, D.G., and Hernandez, R.D. (2012). Population genetics of rare variants and complex diseases. *Hum. Hered.* 74, 118–128.
51. Gazave, E., Chang, D., Clark, A.G., and Keinan, A. (2013). Population growth inflates the per-individual number of deleterious mutations and reduces their mean effect. *Genetics* 195, 969–978.
52. Nei, M., Maruyama, T., and Chakraborty, R. (1975). The bottleneck effect and genetic variability in populations. *Evolution* 29, 1–10.
53. Atzmon, G., Hao, L., Pe'er, I., Velez, C., Pearlman, A., Palamara, P.F., Morrow, B., Friedman, E., Oddoux, C., Burns, E., and Ostrer, H. (2010). Abraham's children in the genome era: major Jewish diaspora populations comprise distinct genetic clusters with shared Middle Eastern Ancestry. *Am. J. Hum. Genet.* 86, 850–859.
54. Palamara, P.F., Lencz, T., Darvasi, A., and Pe'er, I. (2012). Length distributions of identity by descent reveal fine-scale demographic history. *Am. J. Hum. Genet.* 91, 809–822.
55. Sajantila, A., Salem, A.H., Savolainen, P., Bauer, K., Gierig, C., and Pääbo, S. (1996). Paternal and maternal DNA lineages reveal a bottleneck in the founding of the Finnish population. *Proc. Natl. Acad. Sci. USA* 93, 12035–12039.
56. Wang, S.R., Agarwala, V., Flannick, J., Chiang, C.W., Altshuler, D., and Hirschhorn, J.N.; GoT2D Consortium (2014). Simulation of Finnish population history, guided by empirical genetic data, to assess power of rare-variant tests in Finland. *Am. J. Hum. Genet.* 94, 710–720.
57. Lohmueller, K.E. (2014). The distribution of deleterious genetic variation in human populations. *bioRxiv*, <http://dx.doi.org/10.1101/005330>.