

UC Davis

UC Davis Previously Published Works

Title

The timing of human adaptation from Neanderthal introgression

Permalink

<https://escholarship.org/uc/item/0x8384th>

Journal

Genetics, 218(1)

ISSN

0016-6731

Authors

Yair, Sivan

Lee, Kristin M

Coop, Graham

Publication Date

2021-05-17

DOI

10.1093/genetics/iyab052

Peer reviewed

The timing of human adaptation from Neanderthal introgression

Sivan Yair,^{1,2,*} Kristin M. Lee,³ and Graham Coop^{1,2}

¹Center for Population Biology, University of California, Davis, Davis, CA 95616, USA,

²Department of Evolution and Ecology, University of California, Davis, Davis, CA 95616, USA, and

³Department of Biological Sciences, Columbia University, New York, NY 10027, USA

*Corresponding author: Department of Evolution and Ecology, University of California, Davis, One Shields Avenue, Davis, CA 95616, USA. ssyair@ucdavis.edu

Abstract

Admixture has the potential to facilitate adaptation by providing alleles that are immediately adaptive in a new environment or by simply increasing the long-term reservoir of genetic diversity for future adaptation. A growing number of cases of adaptive introgression are being identified in species across the tree of life, however the timing of selection, and therefore the importance of the different evolutionary roles of admixture, is typically unknown. Here, we investigate the spatio-temporal history of selection favoring Neanderthal-introgressed alleles in modern human populations. Using both ancient and present-day samples of modern humans, we integrate the known demographic history of populations, namely population divergence and migration, with tests for selection. We model how a sweep placed along different branches of an admixture graph acts to modify the variance and covariance in neutral allele frequencies among populations at linked loci. Using a method based on this model of allele frequencies, we study previously identified cases of adaptive Neanderthal introgression. From these, we identify cases in which Neanderthal-introgressed alleles were quickly beneficial and other cases in which they persisted at low frequency for some time. For some of the alleles that persisted at low frequency, we show that selection likely independently favored them later on in geographically separated populations. Our work highlights how admixture with ancient hominins has contributed to modern human adaptation and contextualizes observed levels of Neanderthal ancestry in present-day and ancient samples.

Keywords: adaptive introgression; ancient DNA; genetic hitchhiking

Introduction

Within the last decade, population genomic studies have revealed many cases of hybridization that have led to the introgression of genetic material between diverged populations. While these genetic introductions are often ecologically or developmentally maladaptive, a growing number of studies are providing evidence of adaptive introgression, whereby natural selection favored the spread of introgressed alleles (e.g., Whitney *et al.* 2006; The Heliconius Genome Consortium 2012; Jones *et al.* 2018; Oziolor *et al.* 2019). These introgressed alleles likely facilitated adaptation to new or changing environments, highlighting admixture as a potentially important source of genetic variation for fitness.

Patterns of archaic introgression in modern humans offer some of the most compelling examples of adaptive introgression (reviewed in Racimo *et al.* 2015). When a subset of modern humans spread out of Africa, likely in the past hundred thousand years, they encountered a broad range of novel environments, including reduced UV exposure, new pathogen pressures, and colder climates. At roughly the same time that modern humans outside of Africa experienced new conditions, they were also interbreeding with Neanderthals, who had been living in and adapting to these Eurasian environments for hundreds of

thousands of years (Hublin 2009). The early-generation hybrids of Neanderthals and modern humans may have had low fitness due to the accumulation of weakly deleterious alleles in Neanderthals, who had a small effective population size. The low fitness of these early hybrids, combined with the greater efficacy of purifying selection in modern humans, led to widespread selection against deleterious Neanderthal alleles and linked Neanderthal variation (Harris and Nielsen 2016; Juric *et al.* 2016). As a result, Neanderthal alleles segregate at around 1–3% frequency genome-wide in present-day modern humans, with a depletion in gene rich, regulatory, and low-recombination regions (Sankararaman *et al.* 2014; Schumer *et al.* 2018; Steinrücken *et al.* 2018; Petr *et al.* 2019; Telis *et al.* 2020).

While, on average, Neanderthal-derived alleles have been selected against, a few alleles persist at high frequency in present-day non-African populations and reflect putative cases of adaptive introgression (Abi-Rached *et al.* 2011; Khrameeva *et al.* 2014; Sankararaman *et al.* 2014; Vernot and Akey 2014; Dannemann *et al.* 2016; Deschamps *et al.* 2016; Gittelman *et al.* 2016; Quach *et al.* 2016; Sams *et al.* 2016; Racimo *et al.* 2017; Jagoda *et al.* 2018; Setter *et al.* 2020). These alleles have been identified based on the characteristic patterns left behind by adaptive introgression, namely high haplotype similarity between Eurasian modern

humans and Neanderthals, and the deep divergence of haplotypes within modern human populations. It appears that most sweeps on Neanderthal-introgressed alleles were partial, with selected alleles only reaching frequencies around 30–60%. Thus, dips in genetic diversity due to adaptive Neanderthal introgression are dampened relative to hard, full sweeps on *de novo* mutations.

Some of the selected Neanderthal alleles contribute to traits that may have been under selection during modern human expansion out of Africa, such as immunity, skin pigmentation, and metabolism (Abi-Rached *et al.* 2011; Khrameeva *et al.* 2014; Sankararaman *et al.* 2014; Vernot and Akey 2014; Dannemann *et al.* 2016; Gittelman *et al.* 2016; Quach *et al.* 2016; Racimo *et al.* 2017; Jagoda *et al.* 2018). Because modern humans mated with Neanderthals around the same time that they were exposed to new environments, a natural assumption is that Neanderthal variation immediately facilitated modern human adaptation. Alternatively, some of the selected Neanderthal-introgressed alleles may have contributed to the reservoir of standing genetic variation that became adaptive much later (Jagoda *et al.* 2018) as human populations were exposed to and created further novel environments. A recent study found evidence of this phenomenon for a Denisovan haplotype contributing to high-altitude adaptation on the Tibetan Plateau (Zhang *et al.* 2020), and so nonimmediate selection on introgressed archaic variation may be more common than previously thought.

Population genetics offers a number of approaches to date the timing of selection on alleles. The first broad category of approaches relies on the hitchhiking signal created in surrounding linked variants as the selected alleles sweep up a chunk of the haplotype on which they arose (or introgressed) (Maynard Smith and Haigh 1974). Second, ancient DNA now offers an opportunity to assess when these selected haplotypes rose in frequency, and the ancient populations in which they first achieved high frequency. Currently, methods that use ancient DNA to investigate the temporal history of selection focus on identifying significant increases in selected allele frequencies over time (*e.g.*, Mathieson *et al.* 2015; Schraiber *et al.* 2016; Mathieson and Mathieson 2018). These time series approaches are powerful because they can provide more direct evidence of selection driving allele frequency change, however they are limited to characterizing cases of selection that began after the oldest sampling time and among sampled populations. For studies of modern human evolution, we are currently restricted to learning about adaptation in mostly European populations and within the last 10,000 years, long after hybridization with Neanderthals.

Here, we leverage the advantages of both temporal sampling and patterns of neutral diversity at linked loci (the hitchhiking effect) to infer the time that Neanderthal-introgressed alleles became adaptive. Our investigation of the hitchhiking effect allows us to date selection older than the ancient samples, while ancient populations provide useful reference points closer to the sweep time. We use coalescent theory to describe patterns of sequence similarity around a selected site when selection favors introgressed variation at different times since admixture. Our predictions relate sequences among ancient and present-day populations by incorporating information about their history of divergence and migration. In addition, we utilize information from partial sweeps by modeling this hitchhiking effect among both selected and nonselected haplotypes. Our approach builds on the model-based inference framework introduced in Lee and Coop (2017), which connects predicted coalescent histories to their corresponding probability distribution of population allele

frequencies. Thus, we can distinguish among possible selection times by analyzing allele frequency data from ancient and present-day populations.

By providing the age of selection favoring Neanderthal alleles in specific regions, we can determine the context within which Neanderthal alleles facilitated modern human adaptation and in turn narrow their potential phenotypic contributions. If selection was immediate, then the Neanderthal variation was useful in the early Eurasian environments that modern humans experienced, possibly because Neanderthals had been adapting to those same conditions for a long period of time. Otherwise, Neanderthal haplotypes were selected during different time periods and in different populations. These latter haplotypes are particularly interesting candidates, as they shed light on the similarities and differences in selection pressures affecting human populations over space and time. Furthermore, they show that admixture between closely related species can provide an important source of standing genetic variation for future adaptation. In this study, we provide evidence for each of these scenarios in different genomic regions of adaptive introgression.

We begin by introducing the data that we analyzed, followed by an intuitive description of the hitchhiking patterns that we modeled before providing the mathematical details of our model and inference framework. We then show how our method performs on simulations and its results when applied to candidates of adaptive Neanderthal introgression.

Distribution of adaptive Neanderthal-introgressed haplotypes across populations

Choice of population samples

We applied our method to a set of eight population samples: one archaic (Neanderthals), three present-day modern human, and four ancient modern human. For the Neanderthal sample we used the Vindija 33.19 Neanderthal ($n=2$, for two chromosomes sampled from one individual) because it is high coverage and most closely related to the introgressing population (Prüfer *et al.* 2017; Mafessoni *et al.* 2020). Our results only change slightly when the high coverage Altai Neanderthal sample is included in the Neanderthal population, which we discuss later. The present-day populations are from Phase 3 of the 1000 Genomes Project (The 1000 Genomes Project Consortium 2015): the Yoruba in Ibadan, Nigeria (YRI; $n=216$), Han Chinese from Beijing (CHB; $n=206$), and Utah Residents (CEPH) with Northern and Western European Ancestry (CEU; $n=198$). We chose ancient populations based on previous work about their relationships with present-day populations. We included the West Eurasian Upper Paleolithic (EurUP; $n=12$), who are the oldest samples and basally related to all West Eurasians (mean sampling time $t \approx 34$ kya), and three populations known to be ancestral to present-day Europeans: the Mesolithic hunter gatherers from western Europe (WHG; $n=66$; $t \approx 9$ kya), the Neolithic Anatolian Farmers (EF; $n=134$; $t \approx 7$ kya), and the Bronze Age Steppe individuals (Steppe; $n=14$; $t \approx 5$ kya). Each population ancestral to present-day Europeans was composed of samples with $\geq 90\%$ inferred ancestry corresponding to that population. Inferred ancestry proportions are from Mathieson *et al.* (2018)'s supervised ADMIXTURE analysis using four clusters with fixed membership corresponding to each of these three ancient populations and Eastern Hunter Gatherers. See Supplementary File S2 for detailed information on the ancient samples we analyze. One reason for taking this set of ancient populations is that our method requires a known population history, namely the approximate timing and

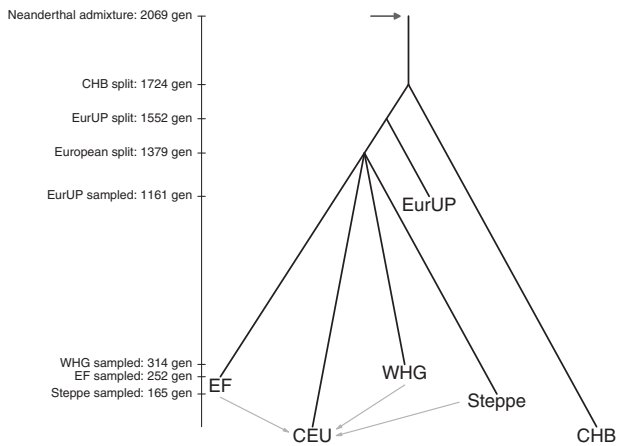


Figure 1 The primary admixture graph that we use in our method. Our method requires the divergence times among admixed modern human populations, the timing and proportion of admixture with Neanderthals, and the sampling time of each population. It also requires the timing and proportions of admixture among modern human populations, which we describe in detail in [Supplementary S1 in Supplementary File S1](#). We converted many of these estimates from years to generations assuming a generation time of 29 years ([Fenner 2005](#)). Our method is robust to modifications to this admixture graph, which we discuss later.

admixture proportions among populations, as well as the timing of divergence among Neanderthal-admixed populations. The primary admixture graph that we used is shown in [Figure 1](#) ([Skoglund et al. 2012](#); [Lazaridis et al. 2014](#); [Allentoft et al. 2015](#); [Haak et al. 2015](#); [Mathieson et al. 2015](#); [Fu et al. 2016](#); [Lipson et al. 2017](#); [Sikora et al. 2017](#); [Mathieson et al. 2018](#)). We acknowledge that we took a coarse approach by simplifying human demographic history to a set of ancestral populations, who themselves are products of mixtures of the past.

All ancient modern human samples were typed on a 1240k capture array. Because we needed dense sampling of the neutral loci surrounding the selected site, we imputed genotypes in the ancient modern human samples using Beagle 4.1 ([Browning and Browning 2007, 2016](#)). To ensure imputation accuracy, we selected ancient samples for analysis if they had $\geq 1\times$ coverage. This cutoff was chosen because at this point an imputation procedure very similar to ours can correctly recover around 80–95% of an ancient sample’s heterozygous sites and preserves properties of the genotypic data such as samples’ PCA locations relative to no imputation ([Mathieson 2016](#)). To perform the imputation we followed a procedure similar to that of [Mathieson et al. \(2015\)](#). At the 1240k sites, we computed genotype likelihoods from read counts according to binomial likelihoods of these counts and a small amount of sequencing error. We then ran Beagle twice, first to impute genotypes at just the 1240k sites using the `gl` and `impute = false` arguments, and then to impute the remaining sites using the `gt` and `impute = true` arguments. After each round of imputation, we removed polymorphic sites with allelic $r^2 < 0.8$ (a measure of imputation accuracy). These intermediate filtering steps can improve the imputation accuracy of even ultra-low coverage ancient samples ([Hui et al. 2020](#)). To account for some of the uncertainty in the imputation procedure, we calculated population allele frequencies by weighting sample genotypes by their posterior probability, rather than using the maximum likelihood genotypes. We imputed each population separately, used the HapMap Project’s genetic map ([The International HapMap Consortium 2007](#)), and used a reference panel consisting of all East Asian and European 1000 Genomes populations

(as multi-population panels can improve imputation accuracy in untyped populations; [Huang et al. 2009](#)). Finally, to assess the impact of imputation uncertainty on our results, we also ran our method on bootstrapped data sets, which we describe after introducing the method.

Choice of Neanderthal-introgressed regions

We analyzed genomic regions that were previously identified as putative candidates of adaptive introgression by [Racimo et al. \(2017\)](#). We chose regions with signals of adaptive introgression in European populations, as we have more information on their ancestral populations to infer the age of selection. We removed any regions whose introgressed sequences could not be distinguished from Denisovan ancestry (according to [Racimo et al. 2017](#)) and any regions on the X chromosome. To ensure that we were capturing the full window of potentially selected sites, we extended the 45 kb windows identified by [Racimo et al. \(2017\)](#) by 2×10^{-2} cM (20 kb if the window has a constant per base pair recombination rate of 10^{-8}) and collapsed those that were overlapping or directly adjacent. This resulted in 36 distinct regions, listed in [Supplementary Table S2 of Supplementary File S1](#), out of 50 45 kb windows with signals of adaptive Neanderthal introgression in Europeans.

Neanderthal ancestry over time

As a first step toward understanding the temporal history of selection, we investigated levels of Neanderthal ancestry in ancient and present-day admixed modern human populations within each of the previously specified regions. These levels were determined from ancestry informative sites, which we identified as biallelic sites in modern humans that have one allele fixed in Neanderthals (combined Vindija and Altai, two high coverage Neanderthals) and at less than 5% frequency in the Yoruba. We included the Altai Neanderthal sample here to increase our confidence that we identified true fixed differences between Neanderthal and modern human lineages. In [Figure 2](#), we show average Neanderthal allele frequencies across all ancestry informative sites with a Neanderthal allele frequency of at least 20% in at least one population. We use this subset of ancestry informative sites because we are interested in frequencies along the selected haplotype(s), which do not necessarily span the entire length of the previously described windows. If selection quickly favored introgressed alleles, we would at least expect all admixed modern human populations, whether sampled in the past or present, to show high levels of Neanderthal ancestry. However, in a number of these genomic regions, Neanderthal ancestry is almost or completely absent from East Asian and/or some ancient populations ([Figure 2](#)).

Method to estimate the timing of selection from introgressed haplotypes

Verbal description

We first describe the intuition behind our models, before laying out the mathematical framework. We focus on how patterns of haplotype similarity among modern human populations and Neanderthals change with the start time of selection. Because we describe selection favoring a Neanderthal-introgressed allele, the earliest selection can begin is the time of admixture between Neanderthals and modern humans. To accompany our description, in [Figure 3](#), we show a cartoon of haplotype diversity and introgressed ancestry under two scenarios: immediate selection and recent selection.




Figure 2 Average Neanderthal allele frequency at ancestry informative sites in ancient and present-day populations in genomic regions with signatures of adaptive introgression. Darker shades of purple correspond to higher average Neanderthal allele frequencies. The size of points correspond to the number of sites over which the average Neanderthal allele frequency was calculated. Moving left to right, regions have decreasing average Neanderthal allele frequencies among ancient populations.

Neanderthal alleles that introgress into modern human populations start on long Neanderthal ancestry tracts (entire chromosomes at the time of admixture) that shorten over time with recombination. Because the selected Neanderthal allele takes its linked variation to high frequency with it, the timing of selection determines the genetic distance over which neutral Neanderthal alleles sweep to high frequency as well. The earlier selection is, the larger the Neanderthal haplotypes that sweep to high frequency with the selected Neanderthal allele. Conversely, the later selection is, the shorter the Neanderthal haplotypes that sweep to high frequency, such that some haplotypes in modern humans that did not introgress from Neanderthals (which we call “modern human haplotypes”) also rise in frequency.

Because Neanderthal genetic diversity is very low, haplotypes that introgressed from Neanderthals into modern human populations should look almost identical to each other and haplotypes sampled from Neanderthals. Conversely, Neanderthal haplotypes are relatively distinct from modern human haplotypes compared to differences typically observed among modern human haplotypes. This is because the earliest time that modern human and Neanderthal haplotypes can coalesce is when Neanderthals and modern humans shared a common ancestor, about 16,000 generations ago. Therefore, when selection brings Neanderthal haplotypes to high frequency in those modern human populations where the Neanderthal allele experienced positive selection (hereafter “selected populations”), selected populations gain unusually high sequence similarity with Neanderthals and unusually low sequence similarity with other modern human populations in which Neanderthal haplotypes are rare in the genomic region affected by the sweep. These patterns of unusually high and low sequence similarity persist over greater genetic distances as selection begins closer to the time of introgression.

Within selected populations, selection increases sequence similarity because the sampled haplotypes descended from one or a few ancestral haplotypes that hitchhiked to high frequency during the sweep. When selection begins earlier, these ancestral haplotypes are mainly one of the few Neanderthal haplotypes introduced by

admixture. The later selection begins, the more likely these ancestral haplotypes are modern human haplotypes that became linked to the selected allele prior to the selection onset.

In the cases of adaptive introgression that we investigated, the putative selected Neanderthal allele did not reach fixation, either because the sweep is ongoing and selection is weak, selection pressures changed before the Neanderthal allele reached fixation, or the phenotypic response to selection was achieved by allele frequency changes at multiple loci. Our models allow for all of these possibilities, where if selection stopped favoring the Neanderthal allele we consider a neutral phase after the sweep. From this phase, we can further distinguish among selection times based on how recombination distributes Neanderthal ancestry among haplotypes within the same population that do and do not carry the selected Neanderthal allele. As the partial sweep finishes, the Neanderthal alleles that hitchhike to higher frequency with the selected allele have a higher chance of recombining onto the background of the nonselected allele. The earlier the onset of selection, the more time post-sweep for these recombination events to occur, and therefore the higher the probability that alleles on the nonselected background descended from Neanderthals.

Model background

We aim to distinguish among the possible scenarios of selection on introgressed variation. Each scenario is defined by the combination of two parameters we aim to infer in our model: the amount of time between Neanderthal admixture and the onset of selection, which we refer to as the waiting time until selection (t_b), and the additive strength of selection favoring the beneficial Neanderthal allele (s). We build on the model-based, statistical approach introduced in Lee and Coop (2017), which uses coalescent theory to describe how different selection scenarios modify the neutral variance and covariance of population allele frequencies surrounding the selected site. This allows us to describe a multivariate normal model of population allele frequencies for each scenario, which serves as a simple approximation for their probability distribution (Nicholson et al. 2002; Weir et al. 2002).




Figure 3 Top: Cartoon of haplotype diversity and ancestry at three time points when selection favors Neanderthal-introgressed alleles immediately ($t_b = 0$) and when selection favors Neanderthal-introgressed alleles more recently ($t_b > 0$). Our models use a specific value of t_b , but here our purpose is to visualize qualitative differences between early and recent selection. Purple tracts represent Neanderthal-introgressed haplotypes, whereas green tracts represent haplotypes that did not descend from Neanderthals. We display Neanderthal ancestry tracts with fewer shades to represent their lower genetic diversity. Bottom: Example frequency trajectory of a selected Neanderthal-introgressed allele when there is a waiting time until selection. The light gray lines represent 50 stochastic trajectories in which the selected allele reached a frequency greater than 20%. The dark gray line represents the trajectory we assume in our models using the same waiting time until selection (t_b), selection coefficient (s), and duration of the sweep phase (t_s) used to generate the stochastic trajectories.

Here, we review the framework laid out by Lee and Coop (2017). We model the change in allele frequency $\Delta x_i = x_i - x_a$ in our sampled populations (i) from their common ancestral population (a) at the root of the tree relating all populations we consider. For neutral alleles, the population allele frequency will on average be the same as the ancestral allele frequency ($E[\Delta x_i] = 0$) because drift and hitchhiking are direction-less on average. Drift and hitchhiking do however cause an increase in the variance in the change in allele frequencies within a population, and pairs of populations can also covary in their change in allele frequency from the ancestor if they have some shared population history or gene flow, i.e., their changes in allele frequency since the ancestral population are not independent of one another. These effects are captured by the population covariance between populations, i and j , given by

$$\text{Cov}[\Delta x_i, \Delta x_j] = x_a(1 - x_a)f_{ij}, \quad (1)$$

where f_{ij} is the probability that the ancestral lineages of an allele sampled in population i and an allele sampled in population j

coalesce before reaching the ancestral population (see Lee and Coop 2017, for details).

We can estimate the neutral probabilities of coalescing among populations from neutral allele frequency data genome-wide. If we are considering k populations in our analysis, we thus have a $k \times k$ matrix, \mathbf{F} , that describes probabilities of coalescing within and between these populations. See Appendix A2 for details on how we estimate this matrix. These estimated neutral probabilities of coalescing allow us to describe our neutral expectations of populations: they implicitly account for population size, divergence, and migration. We note, however, that while this flexible approach allows for arbitrary relationships among populations, our tree requires a root that our coalescent probabilities and frequency deviations will be relative to. We define this root by placing the two most genetically distant populations on opposite sides of it. This implies that in our models, they have no shared history since the ancestral population, and therefore no chance for coalescence between the ancestral lineages of alleles sampled

in each. In our analysis, the Neanderthals and Yoruba are the most distantly related, and therefore are placed on opposite sides of the root, without any subsequent gene flow. This does not negate the possibility of indirect gene flow between them; we simply define all other relationships relative to theirs.

When we incorporate the effects of selection, we model how probabilities of coalescing change with increasing genetic distance from the selected site. Importantly, all of our predictions converge to our neutral estimates because increasing recombination rates between selected and neutral sites allow for increasing independence of their coalescent history. At a far enough genetic distance, dynamics at the selected site become quickly disassociated with those at the neutral site, such that the probability of coalescing at the neutral site is that of any neutral allele unaffected by linked selection.

Haplotype partitioning and the null model

No identified case of adaptive Neanderthal introgression has resulted in the population fixation of a Neanderthal-derived allele. These partial sweeps weaken the effect of linked selection on surrounding neutral diversity because much of the original genetic diversity prior to selection persists in these populations. In order to increase the power of our approach, we divide admixed populations according to ancestry assignments at the putative selected site. For the 1000 Genomes samples, from which we have phased haplotypes, we create two partitions: one consisting of haplotypes that carry the beneficial Neanderthal allele at the selected site (B), and the other consisting of haplotypes that carry the non-Neanderthal allele at the selected site (b). For the ancient samples, we only have unphased genotype data, and so we divide the population samples into three partitions according to individuals' ancestry genotypes at the selected site: BB, Bb, and bb. In our models, we are concerned with a neutral allele's ancestry background at the selected site when it is sampled. Therefore, haplotype partition B is equivalent to genotype partition BB in that any neutral allele sampled in these partitions is linked to the Neanderthal allele at the selected site. Similarly, haplotype partition b is equivalent to genotype partition bb in that any neutral allele sampled in these partitions is linked to the non-Neanderthal allele at the selected site. From here on, we refer to them in our models as partition B or partition b. Predictions for genotype partition Bb are simply a linear combination of our predictions for the other partitions, which we describe in [Supplement S3.5 in Supplementary File S1](#).

Our haplotype partitioning of the data requires us to modify our null model from that of [Lee and Coop \(2017\)](#), who focused on full sweeps and simply used the genome-wide neutral \mathbf{F} matrix to parameterize their null model. The neutral probabilities of coalescing that we estimate between any pair of populations (i, j) are essentially averaged over all possible migration histories of our alleles. However, our population partitioning scheme alters the probability of each migration history because we split populations based on ancestry at the partition site. When we create a population that only carries Neanderthal alleles at the partition site (partition B), nearby sites will have a much higher frequency of Neanderthal alleles relative to the full population. Therefore, even under neutrality, the probability that a randomly sampled allele in partition B will coalesce with an allele sampled in Neanderthals is much higher than in the full population case.

Under population partitioning, we can describe our null coalescent probabilities as a function of the genome-wide neutral probabilities of coalescing and the recombination rate (r) between the partition site and neutral site of interest. We first consider

the relationship between Neanderthal population n and partition B of population i in which all haplotypes carry a Neanderthal allele at the partition site. We define t_i as the number of generations ago that Neanderthal alleles introgressed into modern human populations. We are interested in the probability that the neutral lineage remains linked to the same Neanderthal allele at the partition site that it was linked to at sampling by the time of admixture. We can approximate this probability as e^{-rt_i} . By remaining linked to the Neanderthal allele at the partition site, we know that it, too, introgressed from Neanderthals and thus will coalesce with the Neanderthal allele with approximately the same probability as any two alleles sampled from Neanderthals (f_{nn}). If it does recombine out at some point before admixture, it has the population's neutral probability of coalescing with Neanderthals. This neutral probability of coalescing still accounts for the possibility that the neutral allele descended from Neanderthals, when the average Neanderthal-introgressed allele frequency is the initial admixture fraction (g). We define probabilities of coalescing under the null model as $f_{in}^{(N)}$, where superscript (N) refers to predictions under our null model. In all we derive the probability of coalescing between a pair of lineages sampled from partition B of population i and Neanderthal population n to be

$$f_{in}^{(N)} = e^{-rt_i}f_{nn} + (1 - e^{-rt_i})f_{in}, \text{ if sampled from partition B.} \quad (2)$$

In partition b of the same population (i), all alleles are linked to the non-Neanderthal allele at the partition site when sampled. If an allele sampled in this partition never recombines out, we know it cannot have introgressed. Therefore, in order to have the chance at coalescing with the ancestral lineage of the sampled Neanderthal allele, it must recombine out before admixture. Forward in time, that means that a Neanderthal allele recombined onto the sequence carrying the non-Neanderthal allele at the partition site. Therefore, the ancestral lineage of a neutral allele sampled from partition b of population i has the following probability of coalescing with the neutral lineage sampled from Neanderthals:

$$f_{in}^{(N)} = (1 - e^{-rt_i})f_{in}, \text{ if sampled from partition b.} \quad (3)$$

Selection model

Selection favoring Neanderthal alleles in modern human populations modifies probabilities of coalescing from neutral expectations because it increases levels of Neanderthal ancestry around the selected site in selected populations. Thus, in the whole (nonpartitioned) population, the probability that an allele descended from Neanderthals is much higher than the original admixture proportion, as we described informally above in our discussion of [Figure 3](#).

In selected populations, we consider three phases in the selected allele frequency trajectory: the neutral phase following admixture in which the Neanderthal-derived variant is at frequency g for t_b generations (neutral phase I), the sweep phase in which the variant rises from frequency g to frequency x_s in time t_s , and the neutral phase between the sweep finish and the present in which the variant remains at frequency x_s (neutral phase II). When the relative fitness advantage of the selected allele is additive, such that heterozygotes have an advantage of s and homozygotes $2s$, the sweep duration

$$t_s = \frac{1}{s} \log \left(\frac{x_s(1-g)}{g(1-x_s)} \right), \text{ if } x_s < 1. \quad (4)$$

Since the sum of all three phase durations equals the time between the present day and admixture (t_i), the duration of neutral phase II equals $t_i - t_b - t_s$.

As the waiting time until selection t_b increases, and assuming the same t_b for all selected populations, we transition from describing a case in which the selected allele became beneficial in the common ancestor of all Eurasian populations, to the case in which the selected allele became beneficial independently in each selected population. If t_b is short enough such that selection began in the common ancestor of a group of selected populations, we must account for the possibility that this common ancestor also has descendent populations that do not carry the selected allele at high frequency, possibly due to subsequent drift or negative selection. According to t_b , we assign each population with very low frequency of the selected allele into one of two categories: (i) ancestors never selected or (ii) ancestors selected with subsequent loss of the selected allele in this population. We assign the first category if the low-frequency population does not share a common ancestor with any selected populations at the time selection starts, i.e., its divergence from all selected populations predates selection. Otherwise, we assign the second category.

Similar to the null model, we condition on ancestry at the putative selected site and consider how recombination modulates ancestry with increasing genetic distance. The most important difference between the null and selection model is neutral phase II: in the selection model, if a neutral lineage recombines, it has a high probability of recombining onto the selected allele's background. Therefore, there is a higher chance that the neutral lineage itself descended from Neanderthals.

Between Neanderthals and selected populations:

In this section, we derive the probabilities of coalescing between the haplotypes in each partition of selected human populations and Neanderthals. When we sample an allele, we know whether or not it is linked to the selected Neanderthal allele based on the partition it belongs to. Conditioning on this ancestry background at sampling, we focus on whether the ancestral lineage is linked to the selected Neanderthal allele when we transition between each phase looking backwards from the present to admixture.

First, we determine the probability that a neutral lineage is linked to the selected Neanderthal allele at the transition point between neutral phase II and the sweep completion (which occurs at time $t_i - t_b - t_s$). If there is never a recombination event between the selected and neutral site during neutral phase II, then the neutral lineage remains associated with its initial ancestry background at the time of the sweep completion. Alternatively, if at least one recombination event occurs, then the final recombination event determines the ancestry background that the neutral lineage is associated with at the time of the sweep completion. The frequency of the selected allele determines the probability that a recombining neutral lineage becomes associated with the selected background. Thus, with probability x_s , the neutral lineage becomes associated with the selected Neanderthal allele, and with probability $1 - x_s$ it becomes associated with the non-Neanderthal allele at the selected site. Therefore, the probability that a neutral lineage sampled from a selected population is linked to the selected Neanderthal allele at the transition point between neutral phase II and the sweep completion is

$$\begin{aligned} & \text{Pr}(\text{linked to selected allele at sweep finish}) \\ &= \begin{cases} \exp(-r(t_i - t_b - t_s)) + (1 - \exp(-r(t_i - t_b - t_s)))x_s, & \text{if sampled from partition B.} \\ (1 - \exp(-r(t_i - t_b - t_s)))x_s, & \text{if sampled from partition b.} \end{cases} \end{aligned} \quad (5)$$

in which an allele sampled from partition b can only be linked if it recombines off of its background.

Second, during the sweep phase, an allele that begins linked to the selected allele will always be associated with that background with probability

$$\begin{aligned} & \text{Pr}(\text{always linked to selected allele during sweep}) \\ & \approx \exp\left(-r \int_0^{t_i} 1 - r(1 - X(t))dt\right) \\ & = \left(\frac{x(1-g)}{g}\right)^{-\frac{t_i}{s}}, \end{aligned} \quad (6)$$

where $X(t)$ is the frequency of the selected allele in generation t of the sweep, following a deterministic, logistic trajectory. In words, this probability is approximately the product of the probabilities of not recombining onto the nonselected background each generation of the sweep. Similarly, an allele that begins the sweep phase linked to the nonselected allele will always be associated with that background during the sweep with probability

$$\begin{aligned} & \text{Pr}(\text{always linked to non - selected allele during sweep}) \\ & \approx \left(\frac{1}{x_s}\right)^{-\frac{t_i}{s}}, \text{ if } x_s < 1. \end{aligned} \quad (7)$$

The above associations persist if the lineage never recombines out of its background during neutral phase I, with approximate probability e^{-rt_b} . So, if a neutral lineage remains linked to the selected allele during the sweep phase and fails to recombine during neutral phase I, it must have descended from Neanderthals and thus coalesces with the lineage sampled from Neanderthals with approximately the same probability as any two lineages sampled from Neanderthals, f_{nn} . If the neutral lineage remains linked to the nonselected allele throughout the sweep phase and fails to recombine during neutral phase I, it definitely did not descend from Neanderthals and thus cannot coalesce with an allele sampled from Neanderthals. If a lineage becomes disassociated with its background at some point during the sweep phase, and/or recombines out at least once during neutral phase I, we assume it has its population's neutral probability of coalescing with Neanderthals. Our approximation ignores the possibility that a lineage linked to the nonselected allele can recombine onto the selected allele's background during the sweep. In total, the ancestral lineage of an allele sampled from an admixed, selected population coalesces with the ancestral lineage of an allele sampled from Neanderthals with probability

$$\begin{aligned} & f_{pn}^{(S)} \\ &= [\text{Pr}(\text{linked to selected allele at sweep finish}) \\ & \times \text{Pr}(\text{always linked to selected allele during sweep})e^{-rt_b}f_{nn}] \\ & + [\text{Pr}(\text{linked to selected allele at sweep finish}) \\ & \times (1 - \text{Pr}(\text{always linked to selected allele during sweep})e^{-rt_b}) \\ & + (1 - \text{Pr}(\text{linked to selected allele at sweep finish})) \\ & \times (1 - \text{Pr}(\text{always linked to non-selected allele during sweep}) \\ & \times e^{-rt_b})]f_{pn}, \end{aligned} \quad (8)$$

where the superscript S in $f_{pn}^{(S)}$ refers to our predictions under the selection model, and subscript p denotes any partition of any

selected population. In Supplement S3.1 and S3.2 (Supplementary File S1) we illustrate predictions for other population relationships under the selection model. We follow by describing modifications under both the null and selection models to incorporate ancient samples (Supplement S3.3), migration among admixed modern human populations (Supplement S3.4), and genotype partition Bb (Supplement S3.5). Our approximations provide a good fit to those obtained by simulations, see Figure 4.

Inference

For each model (null or selection) and combination of free parameters, we define a variance-covariance matrix of population allele frequencies for a given distance away from the selected site (Lee and Coop 2017). We approximate the joint probability distribution of population allele frequencies at a neutral locus (l) as being multivariate normal around the ancestral allele frequency (x_{al}) with covariance equal to $x_{al}(1-x_{al})$ times the modified matrix of coalescent probabilities ($\mathbf{F}^{(N)}$ or $\mathbf{F}^{(S)}$) (Nicholson et al. 2002; Weir et al. 2002; Samanta et al. 2009; Coop et al. 2010). In both the null and selection models, the modified matrices of coalescent probabilities depend on the genetic distance to the selected site (r_l), the genome-wide neutral matrix of coalescent probabilities (\mathbf{F}), and all of the admixture graph parameters of interest (A_G), which describe admixture proportions and timing among Neanderthal and admixed modern human populations as well as divergence and sampling times of admixed modern human populations (see Figure 1 and Supplementary Table S1 in Supplementary File S1). The selection model's $\mathbf{F}^{(S)}$ depends on three additional parameters: the strength of selection (s), the waiting time until selection (t_b), and the final frequency of the selected allele (x_s). Therefore, at a neutral polymorphic site with allele frequency data in our populations, we can estimate the probability of the observed population allele frequencies (\bar{x}_l) under the selection model as

$$P(\bar{x}_l | r_l, \mathbf{F}, A_G, s, t_b, x_s) \approx \mathcal{N}(\bar{x}_l | x_{al}, x_{al}(1-x_{al})\mathbf{F}^{(S)}(r_l, \mathbf{F}, A_G, s, t_b, x_s)), \quad (9)$$

where the set of parentheses following $\mathbf{F}^{(S)}$ denote the parameters of the function $\mathbf{F}^{(S)}$. When we predict $\mathbf{F}^{(S)}$, we categorize Neanderthal-admixed populations as “selected” if their frequency of the Neanderthal allele at the selected site is greater than or equal to 0.05. We set x_s to be the average selected Neanderthal allele frequency among all of these putative selected populations. Our method could be extended to allow x_s to vary among populations and to do model choice of the selected populations, however for simplicity we do not pursue those applications here. We estimate the composite likelihood of free parameters s and t_b given the allele frequency data across populations (D) in a window around a candidate selected site by taking the product of all likelihoods calculated at each locus to the left and right of the candidate selected site as follows,

$$\mathcal{L}_C(s, t_b; D) = \prod_{i=1}^{L_{\text{left}}} P(\bar{x}_i | r_i, \mathbf{F}, A_G, s, t_b, x_s) \times \prod_{j=1}^{L_{\text{right}}} P(\bar{x}_j | r_j, \mathbf{F}, A_G, s, t_b, x_s), \quad (10)$$

where L_{left} and L_{right} represent the number of polymorphic loci to the left and right of the candidate selected site. We calculate




Figure 4 Probabilities of coalescing between pairs of alleles at increasing genetic distances from the selected site. Analytical predictions (lines) closely match estimates averaged across 250 simulations (points). Here, we show results for a case in which the Neanderthal-derived allele became adaptive in the ancestors of Early Neolithic Farmers immediately upon introduction ($t_b = 0$, red) and after 600 generations (blue). The favorable allele experienced a selective advantage (s) of 0.01 and on average reached a frequency (x_s) of 0.7 before becoming neutral again. The genetic window corresponds to 1 Mb with a constant recombination rate of 10^{-8} . Partition B refers to alleles linked to the beneficial Neanderthal allele, while partition b refers to alleles linked to the non-Neanderthal counterpart at the selected site. We present categories of population relationships in which probabilities of coalescing vary the most with different waiting times until selection. See Supplementary Figures S12 and S13 in Supplementary File S1 for more population relationships, waiting times until selection, and final frequencies of the selected allele.

composite likelihoods among the proposed parameter combinations of s and t_b shown in Supplementary Table S3 in Supplementary File S1.

The composite likelihood under the null model takes the same form as selection; we simply replace $\mathbf{F}^{(S)}$ with $\mathbf{F}^{(N)}$ and remove

dependence on s , t_b , and x_s . Each partition site represents a potential selected site and thus differs in its genetic distance to each neutral locus. Therefore, each partition site has its own set of composite likelihoods for the null and selection models. We identify the “best” partition site for each region by selecting the site whose ratio of its maximum composite likelihood under the selection model to the null model is greatest, *i.e.*, by maximizing the level of support for selection at this site. From this partition site, we identify the maximum composite likelihood estimates \hat{f}_b and \hat{s} for the region. We note that in our application, these estimates tended to be very similar among partition sites.

Our method focuses on inferring the history of selection from haplotypes alone, and does not incorporate the distribution of the selected allele frequency across populations (as in [Racimo et al. 2018](#); [Refoyo-Martínez et al. 2019](#)). We could incorporate into our composite likelihoods the conditional probability of x_s across populations under a given parameterization of the selection model, but here our focus is on the information about the timing of selection contained in haplotype patterns.

We analyze bi-allelic sites that are polymorphic among our set of samples. Because our multivariate normal approximation works best when alleles segregate ancestrally at some intermediate frequency, we remove sites polymorphic in only one population with a minor allele frequency less than 0.01. In practice, we mean centered our observed allele frequencies, which removes dependence on the ancestral allele frequency x_{al} . We provide more detail on mean centering, sample size correction, and implementation in Appendix A3.

Distinguishing between immediate and nonimmediate selection:

To assess the timing of selection relative to introgression we first distinguish between immediate ($t_b = 0$) and nonimmediate ($t_b > 0$) selection. We do so by using the ratio of the maximum composite likelihood under all parameter combinations to the maximum composite likelihood when selection is immediate,

$$\text{CLR}_{t_b > 0} = \frac{\sup_{s, t_b} \mathcal{L}_C(s, t_b; D)}{\sup_s \mathcal{L}_C(s, t_b = 0; D)}. \quad (11)$$

The higher this ratio, the more support we have in favor of an initial neutral period relative to immediate selection. Because composite likelihoods ignore the correlation in allele frequencies across loci due to linkage disequilibrium, we cannot use traditional statistical methods for model selection. Therefore, we rely on simulations and reject the immediate selection case if the upper 97.5th percentile of $\text{CLR}_{t_b > 0}$ from immediate selection simulations does not exceed the observed value in a region. In other words, we reject the null hypothesis of immediate selection if we observe a value of $\text{CLR}_{t_b > 0}$ high enough to be unlikely if selection were truly immediate. The following sections contain more detail on the simulation procedure used to reject immediate selection.

Validation:

We ran our method on simulated data to evaluate its performance, using SLiM 3.0 for forward-in-time simulations ([Haller and Messer 2019](#)). We simulated 2 cM (2 Mb) loci with the selected mutation at the center of the locus. We simulated under the demographic history shown in [Figure 1](#), along with a divergence time between the Yoruba and Eurasian populations of 2500 generations and a divergence time between modern humans and Neanderthals of 16,000 generations. The Neanderthal population

was simulated with a population size of 3000, whereas all other populations were simulated with a size of 10,000.

Our SLiM simulations generated tree sequences, onto which we added neutral mutations and calculated population allele frequencies using the same sample sizes as in our real data in Python version 3.7.4 with `msprime` ([Kelleher et al. 2016](#)), `tskit` ([Kelleher et al. 2018](#)), and `pyslim` ([Haller et al. 2019](#)). We ran our method on these data in R version 3.4.4. The method assumed the same demographic history that we simulated and used a neutral \mathbf{F} estimated from neutral allele frequency data, also produced by SLiM simulations under the same demographic history as the selection simulations. See Appendix A1 for more simulation details.

Our first goal was to determine our power to detect selection that did not begin immediately, and so we established the significance cutoff using simulations under the null model of immediate selection. Running the method on simulations of immediate selection, we identified a positive correlation between x_s , the average selected allele frequency among all putative selected populations, and $\text{CLR}_{t_b > 0}$, the composite likelihood ratio quantifying the method’s support for nonimmediate selection relative to immediate selection (Supplementary Figure S1 in Supplementary File S1). This relationship between x_s and $\text{CLR}_{t_b > 0}$ among our immediate selection simulations does not differ among combinations of the selection coefficient (s) and the duration of the sweep (t_s) (Supplementary Figure S1 in Supplementary File S1). Thus, for each simulation of nonimmediate selection, we rejected immediate selection if its $\text{CLR}_{t_b > 0}$ was greater than the upper 97.5th percentile of $\text{CLR}_{t_b > 0}$ from the 500 immediate selection simulations with the closest x_s . We used a similar procedure to determine the significance cutoffs for the real data set.

Varying the onset time of selection in our simulations, we found that the method’s power to reject immediate selection increases with the true waiting time until selection (t_b) and x_s ([Figure 5](#)). Our method has reasonable power to reject immediate selection when selection begins >1000 generations after admixture, suggesting that we should correctly identify many of the regions that only recently contributed to adaptation. Overall, s is not well inferred, likely because we deal with old selection events. Therefore, while s appears in our selection model, we concentrate on inferring t_b in our application. The method yields somewhat biased estimates of the waiting time until selection (Supplementary Figures S2–S5 in Supplementary File S1). However, among the simulations in which the method rejected immediate selection, the method does a good job tracking the true waiting time until selection ([Figure 6](#)). This performance decreases when fewer populations are considered to be selected, but still, the method can identify that selection started much later in time. For more details on method performance, see Supplementary S2.1 in Supplementary File S1.

Data application to estimate the timing of selection on Neanderthal alleles

Procedure for identifying selected site and parameter estimates

For each region, we ran our method on the previously defined windows in addition to 1 cM flanking them on each side, because if selection is immediate with $s = 0.01$ it takes about 1 cM for the signals of selection to almost completely decay. We used the HapMap Project’s genetic map to identify the endpoints of genomic analysis windows and recombination rates between each neutral and selected site for the method (note that this means

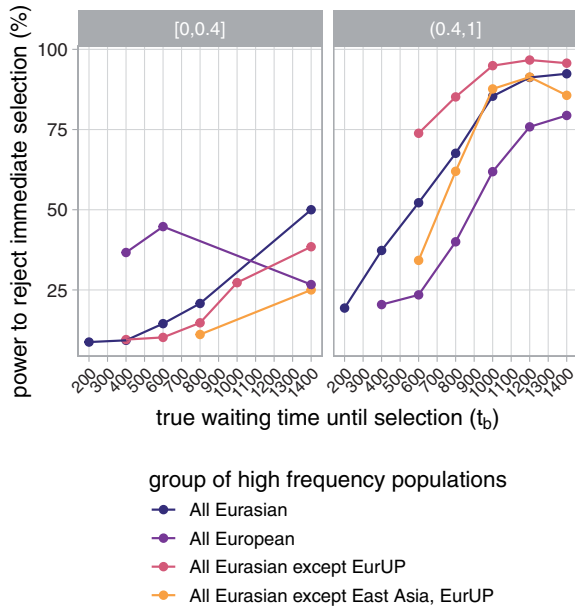


Figure 5 Power to reject immediate selection as the true waiting time until selection (t_b) increases. Lines correspond to simulations with different groups of high-frequency populations (each with the selected allele frequency greater than 0.05), i.e., groups of populations considered selected. Note that this does not perfectly correspond to groups of populations truly selected in our simulations, and so we have different numbers of observations contributing to each point. Points with fewer than 10 observations were removed. Panels correspond to bins of x_s , the selected allele’s average frequency among populations considered selected. Results for $x_s > 0.4$ are qualitatively similar, though we show finer bins in Supplementary Figure S7 in Supplementary File S1.

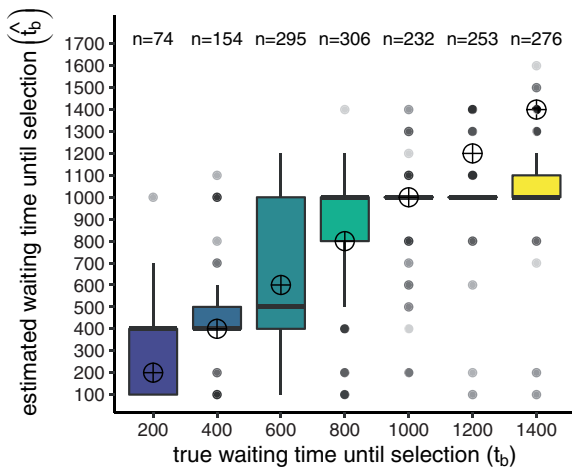


Figure 6 Estimated waiting time until selection (t_b) among simulations with all populations considered selected (each $x_s > 0.05$) and in which the method rejected immediate selection. Target points mark the true waiting time until selection that we aim to estimate. Earlier true waiting times have fewer observations because the method has less power to reject immediate selection in this parameter space. Results for more groups of putative selected populations are shown in Supplementary Figure S6 in Supplementary File S1. Estimated waiting times until selection under the immediate selection case can be found in Supplementary Figure S1 in Supplementary File S1.

that we need to assume that recombination rates have been relatively constant over this time period). We excluded sites without Neanderthal allele frequency data from our analysis, as they are

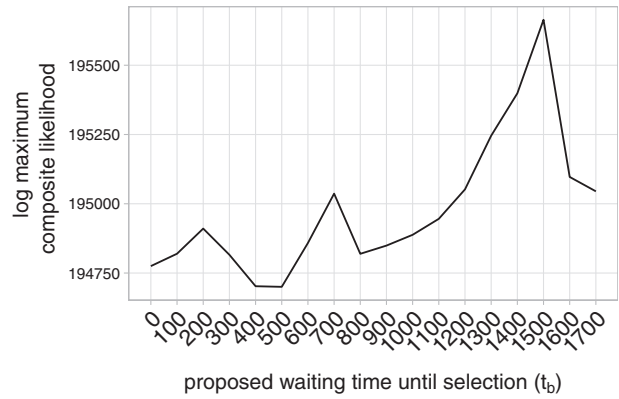


Figure 7 Profile composite likelihood surface of the waiting time until selection (t_b) in the region OAS1, OAS3.

less informative of introgression patterns and the method might instead pick up on unrelated signals of selection at these sites.

We estimated the neutral F from putative neutral allele frequency data genomewide. We chose regions at least 300 base pairs long, at least 0.4 cM away from a gene, and with a minimum per base pair recombination rate of 0.9×10^{-8} using the “Neutral Regions Explorer” (Arbiza et al. 2012). This led to 105,786 sites after filtering. While these sites may not be entirely unaffected by linked selection, we can still use them to represent background patterns of genetic diversity from which we distinguish strongly altered patterns caused by adaptive introgression.

Among the ancestry informative sites in each region, we chose partition sites (which represent potential selected sites) to be all sites with at least two ancient populations with data and a Neanderthal allele frequency $\geq 20\%$ in at least one European, East Asian, or South Asian 1000 Genomes population. For a region, we looped over possible partition sites to run the method on their corresponding data set, and then selected among these partition sites. Then from the “best” partition site we selected among parameter estimates of t_b and s . For example, in Figure 7, we show the profile composite likelihood surface of t_b at the maximum composite likelihood partition site (relative to the null model at this site) in the region OAS1, OAS3. The peak in this surface corresponds to our estimate of t_b , which in this region we found to be very high. As the relationships among our populations are only partially understood, we ran the method with the Altai Neanderthal sample included in the Neanderthal population (discussed in Supplement S2.2.1 and Supplementary Figure S10 in Supplementary File S1) and with multiple plausible modifications to the admixture graph in Figure 1 (discussed in Supplement S2.2.2 and Supplementary Figure S11 in Supplementary File S1). To account for the uncertainty in our imputation procedure, in each region we reran the method 40 times on bootstrapped data sets: rather than using the maximum likelihood genotype at the partition site to assign an ancient sample to a genotype partition, in each bootstrap run we randomly assigned the sample to one of the partitions according to its posterior genotype probabilities provided by the imputation algorithm (results shown in Supplementary Figure S9 in Supplementary File S1).

Procedure for distinguishing between immediate and nonimmediate selection

Among the regions in which the method estimated a nonimmediate waiting time until selection ($t_b > 0$), we rejected immediate




Figure 8 Top: Distribution of estimated waiting times until selection (\hat{t}_b) in the application of the method to 36 regions of adaptive introgression. Blue bars represent cases in which the method did not reject immediate selection, and orange bars represent cases in which the method did reject immediate selection. Blue bars are stacked on top of orange bars (they do not overlap). Bottom: Results among all regions with $\hat{t}_b > 0$. These represent 26 out of 36 cases. Regions in which the method did not reject immediate selection are colored blue, and regions in which the method rejected immediate selection are colored orange (17 cases).

selection if their $\text{CLR}_{\hat{t}_b > 0}$ exceeded the upper 97.5th percentile of those from the 500 immediate selection simulations with the closest x_s (the average selected allele frequency among all putative selected populations). Recall that we can only identify these significance thresholds with simulations. Because composite likelihoods may be sensitive to the spatial distribution of analyzed sites and their corresponding populations with known allele frequencies, we sampled sites to analyze in the simulated data using the same spatial patterning as the region of interest. Specifically, we binned the genetic positions of analyzed sites into 10^{-3} cM intervals to the left and right of the selected site. For each of these sites, we sampled a site in the simulated data from the same genetic distance bin and masked its allele frequencies from populations with no data. We used the same immediate selection simulations as in the Validation section, except this time evaluating them over a larger 3Mb (3 cM) locus.

Of the 36 regions we ran our method on, our method estimated $\hat{t}_b > 0$ in 26, and rejected immediate selection in 17 (Figure 8). Four of the regions in which we rejected immediate selection had a standard deviation greater than 100 generations of \hat{t}_b in their bootstraps over the imputation uncertainty: CHRM2, chr1:193885932, chr2:154493544, and chr12:84903554 (Supplementary Figure S9 in Supplementary File S1). For the remaining 13 regions in which we have greater confidence in our estimated waiting times, nine of these regions have $\hat{t}_b \leq 700$, whereas the remaining four have extremely high estimates: SEMA7A, UBL7 has $\hat{t}_b = 1300$ and CACNA1S, ASCL5; OAS1, OAS3; and chr4:189187062 have $\hat{t}_b = 1500$.

As with many selection scans, the function of alleles underlying adaptation is obviously unknown. Among the set of regions with a shorter initial period at low frequency, BNC2 is a better-studied candidate of adaptive Neanderthal introgression, where the likely selected Neanderthal variant is associated with lighter skin pigmentation (Dannemann et al. 2017). As for the other regions with early selection, mutations in EYS affect the development and maintenance of photoreceptors in the retina (Abd El-Aziz et al. 2008; Collin et al. 2008; Alfano et al. 2016), proteins encoded by EVC are involved in the regulation of Hedgehog signaling and mutations in this gene affect patterning during development (Ruiz-Perez et al. 2000; Blair et al. 2011; Caparrós-Martín et al. 2013; Pusapati et al. 2014), proteins encoded by CRMP1 may be involved in the development of the nervous system and epithelial sheets (Nakamura et al. 2014; Yu-Kemp et al. 2017), and SLC7A10 contributes to synaptic regulation in the central nervous system (Ehmsen et al. 2016; Palaćin et al. 2016; Mesuret et al. 2018). HELZ2 is involved in regulating the differentiation of adipocytes and could thus play a role in metabolism (Katano-Toki et al. 2013). PLA2R1, ITGB6, and SLIT3 putatively play a role in tumor suppression, however mutations at all of these loci are also known to influence variation in other phenotypes (Marlow et al. 2008; Hezel et al. 2012; Guo et al. 2013; Bernard and Vindrieux 2014; Yu et al. 2014; Ng et al. 2018; Yoshikawa and Asaba 2020). Most research on RHPN2 and PTK6 has focused on the role of their expression in cancer progression (e.g., Zheng et al. 2010; He et al. 2015). The remainder of regions in this category is poorly studied or intergenic and far from genes. As for the recently selected regions, SEMA7A plays a role in both the immune and nervous systems, regulating active T cells in the former and axon growth and formation during development in the latter (Pasterkamp et al. 2003; Suzuki et al. 2007; Gras et al. 2013). Polymorphisms at this locus are also associated with variation in bone mineral density (Koh et al. 2006). UBL7 may play a role in ubiquitin signaling, however little is known about the effects of variation at this locus (Zhang et al. 2017). CACNA1S encodes a protein known to contribute to the structure of calcium channels involved in skeletal muscle contraction (Tanabe et al. 1990; Ptáček et al. 1994; Ertel et al. 2000) and in some cases can influence body heat regulation (Beam et al. 2017). ASCL5 is less well studied, though mutations in this gene have been implicated to affect tooth development (He et al. 2019). OAS1 and OAS3 haplotypes can influence innate immunity, which we discuss in more detail later. Based on what we currently understand about mutations in these genomic regions, the nonimmediate benefits provided by Neanderthal alleles may have varied widely and were perhaps associated with different environmental factors.

Discussion

Here, we use patterns of linked selection to develop one of the first model-based methods that infers the time at which Neanderthal-introgressed alleles became adaptive. Our approach uses ancient DNA as well as the hitchhiking effect to investigate the temporal history of selection. We directly address partial sweeps, which likely reflect most cases of human adaptation and perhaps adaptive introgression more generally, and use them as a tool to time the onset of selection. While we require a known demographic history among Neanderthal-admixed populations, our method is robust to modifications in these assumptions (see Supplement S2.2 in Supplementary File S1 for discussion).

We found that in most regions that we analyzed, we could not reject a model of selection immediately favoring Neanderthal-introgressed alleles upon admixture. In some of these regions, East Asians and/or ancient populations do not contain signatures of adaptive introgression. Obviously some of these cases may represent our lack of power to rule out short onset times from haplotypes alone and so our results should not be taken as indicating that selection began immediately for many Neanderthal haplotypes. In addition, the regions that we ran our method on were identified from signals of unusual sharing with Neanderthals (Racimo et al. 2017), and so we may be missing some cases of recent selection on Neanderthal introgression that dragged up only very short blocks of Neanderthal ancestry. We also limited our analysis to regions with signatures of adaptive introgression in European populations, so there may be cases of recent selection in East Asian or other populations that we did not identify. Among the regions in which we did reject immediate selection, we estimated a wide range of waiting times until selection. From simulations, we found that we can generally classify these cases into earlier ($0 < \hat{t}_b \leq 800$) and more recent ($\hat{t}_b > 800$) selection. These time frames may be coarse, but with our estimates of recent selection being far from the boundary between early and recent selection, the transition between them possibly marks the end of the last glacial period (~11.5 kya). Thus, we are potentially distinguishing between adaptation to conditions during the Upper Paleolithic versus later periods defined by warming, technological innovation, the emergence of farming, and higher population densities.

We modeled that the beneficial allele was at a low frequency for some time followed by the onset of selection due to some environmental change. This initial period could correspond to the Neanderthal allele being neutral or balanced at low frequency. However, there are a few possibilities other than a truly delayed onset of selection that could explain our rejection of immediate selection. First, selection may have immediately favored an allele, but the allele was recessive and thus took a long time to rise in frequency. We chose to model additive selection because we do not have information on dominance, but for recessive alleles we may be timing its rise to intermediate frequency rather than the selection onset (see Jones et al. 2020, for a possible dominance related lag in an introgressed pigmentation allele in snowshoe hares). Given that the method does a poor job estimating the selection coefficient for these old sweeps from haplotype data alone, we likely cannot distinguish among different models of dominance. Second, selection may have immediately favored a beneficial Neanderthal allele, but it was only able to rise in frequency once it recombined off other tightly linked deleterious alleles on its Neanderthal haplotype background. There are two ways that this could lead us to incorrectly infer more recent selection: (i) the beneficial allele could have recombined earlier than expected, causing it to rise in frequency close to its introgression time but on a shorter Neanderthal haplotype than we predict under immediate selection, or (ii) a time lag between the selection onset and the recombination event that allowed the allele to rise in frequency could cause our timing estimates to mark the recombination event rather than the selection onset. Some studies have described the conditions in which an adaptive allele could introgress and eventually fix given its linked deleterious background upon introduction (e.g., Uecker et al. 2015), with some investigation into the timescale in which recombination would generate a haplotype with a net selective advantage (Sachdeva and Barton 2018). Finally, the sweep of Neanderthal alleles up in frequency may not be due to a beneficial Neanderthal allele, but

instead hitchhiking with a new beneficial allele that arose via mutation more recently than the Neanderthal introgression, by chance on the background of a Neanderthal haplotype. If Neanderthal-introgressed alleles are at ~2% frequency genome-wide, then we should expect about 2% of recent sweeps from new mutation to arise on a Neanderthal-introgressed background and sweep them to high frequency.

We identify two regions as having clear statistical support for very recent selection (OAS1, OAS3 and CACNA1S, ASCL5). These regions have very high \hat{t}_b , which our simulations show are very rarely estimated under cases of early selection. We are less convinced of alternative explanations for their late rise in frequency because they would need to be replicated in multiple populations. Specifically, given that the Neanderthal alleles rose in frequency so late, and that in both of these regions Neanderthal ancestry is at high frequency in present-day Europeans and East Asians, this rise must have begun after Europeans and East Asians diverged from each other, i.e., independently. The chance that the Neanderthal alleles both recombined off of their deleterious background or both hitchhiked on a sweep in the region is extremely small. Thus, for OAS1, OAS3 and CACNA1S, ASCL5, we have evidence that Neanderthal alleles were independently beneficial in Europeans and East Asians, well after admixture with Neanderthals.

In the CACNA1S, ASCL5 region, the estimated waiting time until selection ($\hat{t}_b = 1500$) implies selection began after the sampling time of the West Eurasian Upper Paleolithic, however this population does carry Neanderthal haplotypes at high frequency in this region. The Neanderthal allele may have drifted away from low frequency in this population after it diverged from the European ancestry populations that we study, which our simulations show happens reasonably often and does not confound our method. Alternatively, selection may have started earlier in the common ancestor of European ancestry populations and more recently in East Asian ancestry populations. We reran our method after removing East Asian samples but our estimates did not change, which could either reflect a selection onset more recent than when the West Eurasian Upper Paleolithic were sampled or poor performance of the method in the absence of East Asian samples. The latter possibility was beyond the scope of this paper to investigate, however future applications could investigate the effect of sampling and extend the method to allow for different onset times when selection independently favors alleles. In total, Neanderthal alleles in this region appear to have been independently favored in the ancestors of East Asians and the ancestors of Europeans, indicating that geographically separated populations experienced similar changes in environmental conditions, possibly at different times.

The selection pressure favoring the OAS1, OAS3 haplotype is thought to be related to innate immunity. Indeed, a recent study found that the Neanderthal haplotype at OAS1 is protective against COVID-19 severity, hospitalization, and susceptibility (Zhou et al. 2020; Zeberg and Pääbo 2021). It has previously been suggested that this gene contributes to traits under balancing selection because haplotypes at this locus have shown variable expression responses to different Flaviviruses (Sams et al. 2016) and a separate Denisovan haplotype segregates at high frequency at this locus in Melanesians (Mendez et al. 2012). Our recent date of selection, and the fact that genomic samples dated closer to the admixture time do not carry Neanderthal haplotypes in this region, are consistent with a recent onset of positive selection in a long-term cycle of balancing selection. A number of other genes known to influence immunity such as the TLR6-TLR1-TLR10

cluster and HLA class I genes harbor candidates of adaptive introgression, however, we did not analyze them as they contain both Neanderthal and Denisovan haplotypes in Eurasian populations (Abi-Rached *et al.* 2011; Dannemann *et al.* 2016). The timing of adaptation for the OAS1, OAS3 region is at odds with the hypothesis that the transfer of pathogens from Neanderthals to modern humans necessitated rapid human adaptation using Neanderthal-derived immunity alleles (Enard and Petrov 2018; Greenbaum *et al.* 2019). However, it would be interesting to test this idea more generally by attempting to date the start of adaptation for more of these introgressed, putative immunity-related haplotypes.

Our investigation contributes to a growing model-based understanding of the genomic patterns of adaptive introgression. Setter *et al.* (2020) introduced VolcanoFinder, a method that identifies cases of adaptive introgression by probing the recipient population's patterns of heterozygosity around the selected site. Using expected coalescence times and assuming fixation of the selected allele, they predict very low diversity in a small window around the selected site and very high levels of diversity at intermediate distances. Our predicted within-population probabilities of coalescing characterize similar patterns when combining results from both partitions of a selected population (weighted by their frequency in the population, *i.e.*, the frequency of the selected allele). However, the characteristic volcano pattern would disappear as the frequency of the selected allele decreases or waiting time until selection increases. Our work demonstrates that both of these situations are frequent, and thus the continued difficulty of developing relaxed tests for adaptive introgression while avoiding false positives. Shchur *et al.* (2020) also use a coalescent approach, in their case to model how adaptive introgression affects the distribution of introgression tract lengths. Fixing the time until introgression and admixture proportion, they show how stronger selection increases the introgression tract lengths. While we do not explicitly model tract lengths, these patterns are apparent in our coalescent predictions.

We implemented a flexible inference framework that could be further modified to investigate the spatio-temporal history of adaptation in modern humans and other systems, whether via introgression or derived mutations. Future applications could distinguish among groups of truly selected populations, allow the frequency of the selected allele to vary among populations, or to distinguish between drift and selection. Our models and estimation method for the timing of adaptation via introgression can be generalized to other organismal systems. In our current application, we assume that the introgressed lineage coalesces with the Neanderthal lineage due to the Neanderthal population's low effective population size, as this allows us to side step modeling the sweep in Neanderthals. However, in other systems, the donor population cannot be assumed to have a very low effective population size, but it should be simple to include the extra terms modeling the sweep-induced increase in coalescence between the recipient and donor populations close to the selected site. The application of methods like this to other species would allow a more general understanding of how often introgression is rapidly favored versus supplying genetic variation for future adaptation.

Using a combination of ancient DNA and haplotype-based timing, we documented spatially and temporally varying selection on Neanderthal alleles. Our results provide evidence that admixture between diverged populations can be a source of genetic variation for adaptation in the long term. They also allow us to better understand the historical and geographic contexts within

which selection favored Neanderthal introgression. As experimental work begins to identify the specific alleles and phenotypes potentially selected on, we can more fully flesh out how interbreeding with archaic populations tens of thousands of years ago has shaped the evolution of modern human populations.

Data availability

Scripts for imputation, simulations, and the method are provided at www.github.com/SivanYair/selTime_neanderthal_AI. The script that computes genotype likelihoods can be found at www.github.com/mathii/gdc3/blob/master/apulldown.py. Supplementary materials (Supplementary Files S1 and S2) are available at figshare: <https://doi.org/10.25386/genetics.14192909>.

Acknowledgments

The authors thank Iain Mathieson for helpful feedback and for sharing genotype likelihoods from the ancient samples that we included in our analysis. We also thank members of the Coop lab for helpful discussions and Erin Calfee, Chuck Langley, Pavitra Muralidhar, Matthew Osmond, Anita To, Michael Turelli, and Carl Veller for comments on earlier drafts of our manuscript. They thank Rasmus Nielsen and three anonymous reviewers for valuable feedback on earlier drafts as well. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) computing resource PSC Bridges through allocation TG-MCB200074 to S.Y.

Funding

This work was supported by the National Science Foundation (IOS-1353380 and PGR-1546719) and the National Institutes of Health (GM108779 and GM136290 to G.C. and GM121372 to Molly Przeworski).

Conflicts of interest

None declared.

Literature cited

- Abd El-Aziz MM, Barragan I, O'Driscoll CA, Goodstadt L, Prigmore E, *et al.* 2008. EYS, encoding an ortholog of *Drosophila* spacemaker, is mutated in autosomal recessive retinitis pigmentosa. *Nat Genet.* 40:1285–1287.
- Abi-Rached L, Jobin M, Kulkarni S, McWhinnie A, Dalva K, *et al.* 2011. The shaping of modern human immune systems by multiregional admixture with archaic humans. *Science.* 334:89–95.
- Alfano G, Kruczek PM, Shah AZ, Kramarz B, Jeffery G, *et al.* 2016. EYS is a protein associated with the ciliary axoneme in rods and cones. *PLoS One.* 11:e0166397.
- Allentoft ME, Sikora M, Sjögren K-G, Rasmussen S, Rasmussen M, *et al.* 2015. Population genomics of Bronze Age Eurasia. *Nature.* 522:167–172.
- Arbiza L, Zhong E, Keinan A. 2012. NRE: a tool for exploring neutral loci in the human genome. *BMC Bioinformatics.* 13:301.
- Beam TA, Loudermilk EF, Kisor DF. 2017. Pharmacogenetics and pathophysiology of CACNA1S mutations in malignant hyperthermia. *Physiol Genomics.* 49:81–87.
- Bernard D, Vindrieux D. 2014. PLA2R1: expression and function in cancer. *Biochim Biophys Acta.* 1846:40–44.

- Blair HJ, Tompson S, Liu YN, Campbell J, MacArthur K, et al. 2011. *Evc2* is a positive modulator of Hedgehog signalling that interacts with *Evc* at the cilia membrane and is also found in the nucleus. *BMC Biol.* 9:1–13.
- Browning BL, Browning SR. 2016. Genotype imputation with millions of reference samples. *Am J Hum Genet.* 98:116–126.
- Browning SR, Browning BL. 2007. Rapid and accurate haplotype phasing and missing-data inference for whole-genome association studies by use of localized haplotype clustering. *Am J Hum Genet.* 81:1084–1097.
- Caparrós-Martín JA, Valencia M, Reytor E, Pacheco M, Fernandez M, et al. 2013. The ciliary *Evc/Evc2* complex interacts with *Smo* and controls Hedgehog pathway activity in chondrocytes by regulating *Sufu/Gli3* dissociation and *Gli3* trafficking in primary cilia. *Hum Mol Genet.* 22:124–139.
- Collin RW, Littink KW, Klevering BJ, van den Born LI, Koenekoop RK, et al. 2008. Identification of a 2 Mb human ortholog of *Drosophila* eyes shut/spacemaker that is mutated in patients with Retinitis Pigmentosa. *Am J Hum Genet.* 83:594–603.
- Coop G, Witonsky D, Rienzo AD, Pritchard JK. 2010. Using environmental correlations to identify loci underlying local adaptation. *Genetics.* 185:1411–1423.
- Dannemann M, Andrés AM, Kelso J. 2016. Introgression of Neandertal- and Denisovan-like haplotypes contributes to adaptive variation in human toll-like receptors. *Am J Hum Genet.* 98:22–33.
- Dannemann M, Prüfer K, Kelso J. 2017. Functional implications of Neandertal introgression in modern humans. *Genome Biol.* 18:1–11.
- Deschamps M, Laval G, Fagny M, Itan Y, Abel L, et al. 2016. Genomic signatures of selective pressures and introgression from archaic hominins at human innate immunity genes. *Am J Hum Genet.* 98:5–21.
- Ehmsen JT, Liu Y, Wang Y, Paladugu N, Johnson AE, et al. 2016. The astrocytic transporter *SLC7A10* (*Asc-1*) mediates glycinergic inhibition of spinal cord motor neurons. *Sci Rep.* 6:35592.
- Enard D, Petrov DA. 2018. Evidence that RNA viruses drove adaptive introgression between Neanderthals and modern humans. *Cell.* 175:360–371.
- Ertel EA, Campbell KP, Harpold MM, Hofmann F, Mori Y, et al. 2000. Nomenclature of voltage-gated calcium channels voltage-gated. *Neuron.* 25:533–535.
- Fenner JN. 2005. Cross-cultural estimation of the human generation interval for use in genetics-based population divergence studies. *Am J Phys Anthropol.* 128:415–423.
- Fu Q, Posth C, Hajdinjak M, Petr M, Mallick S, et al. 2016. The genetic history of Ice Age Europe. *Nature.* 534:200–205.
- Gittelman RM, Schraiber JG, Vernot B, Mikacenic C, Wurfel MM, et al. 2016. Archaic hominin admixture facilitated adaptation to out-of-Africa environments. *Curr Biol.* 26:3375–3382.
- Gras C, Eiz-Vesper B, Seltsam A, Immenschuh S, Blasczyk R, et al. 2013. Semaphorin 7A protein variants differentially regulate T-cell activity. *Transfusion.* 53:270–283.
- Greenbaum G, Getz WM, Rosenberg NA, Feldman MW, Hovers E, et al. 2019. Disease transmission and introgression can explain the long-lasting contact zone of modern humans and Neanderthals. *Nat Commun.* 10:1–12.
- Guo H, Du G, Wang L, Wang D, Hu L, et al. 2013. Integrin alpha v beta 6 contributes to maintaining corneal epithelial barrier function. *Cell Biol Int.* 37:593–599.
- Haak W, Lazaridis I, Patterson N, Rohland N, Mallick S, et al. 2015. Massive migration from the steppe was a source for Indo-European languages in Europe. *Nature.* 522:207–211.
- Haller BC, Galloway J, Kelleher J, Messer PW, Ralph PL. 2019. Tree-sequence recording in SLiM opens new horizons for forward-time simulation of whole genomes. *Mol Ecol Resour.* 19:552–566.
- Haller BC, Messer PW. 2019. SLiM 3: forward genetic simulations beyond the wright-fisher model. *Mol Biol Evol.* 36:632–637.
- Harris K, Nielsen R. 2016. The genetic cost of Neanderthal introgression. *Genetics.* 203:881–891.
- He B, Chiba Y, Li H, de Vega S, Tanaka K, et al. 2019. Identification of the novel tooth-specific transcription factor *AmeloD*. *J Dent Res.* 98:234–241.
- He D, Ma L, Feng R, Zhang L, Jiang Y, et al. 2015. Analyzing large-scale samples highlights significant association between rs10411210 polymorphism and colorectal cancer. *Biomed Pharmacother.* 74:164–168.
- Hezel AF, Deshpande V, Zimmerman SM, Contino G, Alagesan B, et al. 2012. TGF- β and $\alpha v\beta 6$ integrin act in a common pathway to suppress pancreatic cancer progression. *Cancer Res.* 72:4840–4845.
- Huang L, Li Y, Singleton AB, Hardy JA, Abecasis G, et al. 2009. Genotype-imputation accuracy across worldwide human populations. *Am J Hum Genet.* 84:235–250.
- Hublin JJ. 2009. The origin of Neandertals. *Proc Natl Acad Sci USA.* 106:16022–16027.
- Hui R, D’Atanasio E, Cassidy LM, Scheib CL, Kivisild T. 2020. Evaluating genotype imputation pipeline for ultra-low coverage ancient genomes. *Sci Rep.* 10:8.
- Jagoda E, Lawson DJ, Wall JD, Lambert D, Muller C, et al. 2018. Disentangling immediate adaptive introgression from selection on standing introgressed variation in humans. *Mol Biol Evol.* 35:623–630.
- Jones MR, Mills LS, Alves PC, M CC, Alves JM, et al. 2018. Adaptive introgression underlies polymorphic seasonal camouflage in snowshoe hares. *Science.* 360:1355–1358.
- Jones MR, Mills LS, Jensen JD, Good JM. 2020. The origin and spread of locally adaptive seasonal camouflage in snowshoe hares. *Am Natural.* 196:316–332.
- Juric I, Aeschbacher S, Coop G. 2016. The strength of selection against Neanderthal introgression. *PLoS Genet.* 12:e1006340.
- Katano-Toki A, Satoh T, Tomaru T, Yoshino S, Ishizuka T, et al. 2013. *THRAP3* interacts with *HELZ2* and plays a novel role in adipocyte differentiation. *Mol Endocrinol.* 27:769–780.
- Kelleher J, Etheridge AM, McVean G. 2016. Efficient coalescent simulation and genealogical analysis for large sample sizes. *PLoS Comput Biol.* 12:e1004842.
- Kelleher J, Thornton KR, Ashander J, Ralph PL. 2018. Efficient pedigree recording for fast population genetics simulation. *PLoS Comput Biol.* 14:e1006581.
- Khrameeva EE, Bozek K, He L, Yan Z, Jiang X, et al. 2014. Neanderthal ancestry drives evolution of lipid catabolism in contemporary Europeans. *Nat Commun.* 5:1–8.
- Koh JM, Oh B, Lee JY, Lee JK, Kimm K, et al. 2006. Association study of semaphorin 7a (*sema7a*) polymorphisms with bone mineral density and fracture risk in postmenopausal Korean women. *J Hum Genet.* 51:112–117.
- Lazaridis I, Patterson N, Mittnik A, Renaud G, Mallick S, et al. 2014. Ancient human genomes suggest three ancestral populations for present-day Europeans. *Nature.* 513:409–413.
- Lee KM, Coop G. 2017. Distinguishing among modes of convergent adaptation using population genomic data. *Genetics.* 207:1591–1619.
- Lipson M, Szécsényi-Nagy A, Mallick S, Pósa A, Stégnár B, et al. 2017. Parallel palaeogenomic transects reveal complex genetic history of early European farmers. *Nature.* 551:368–372.

- Mafessoni F, Grote S, de Filippo C, Slon V, Kolobova KA, et al. 2020. A high-coverage Neandertal genome from Chagyrskaya Cave. *Proc Natl Acad Sci USA*. 117:15132–15136.
- Marlow R, Strickland P, Lee JS, Wu X, Pebenito M, et al. 2008. SLITs suppress tumor growth in vivo by silencing Sdf1/Cxcr4 within breast epithelium. *Cancer Res*. 68:7819–7827.
- Mathieson I. 2016. Imputation for Ancient Data. Last accessed: April 13, 2021. Retrieved from <https://mathii.github.io/2016/11/01/imputation-for-ancient-data>.
- Mathieson I, Alpaslan-Roodenberg S, Posth C, Szécsényi-Nagy A, Rohland N, et al. 2018. The genomic history of southeastern Europe. *Nature*. 555:197–203.
- Mathieson I, Lazaridis I, Rohland N, Mallick S, Patterson N, et al. 2015. Genome-wide patterns of selection in 230 ancient Eurasians. *Nature*. 528:499–503.
- Mathieson S, Mathieson I. 2018. FADS1 and the timing of human adaptation to agriculture. *Mol Biol Evol*. 35:2957–2970.
- Maynard Smith J, Haigh J. 1974. The hitch-hiking effect of a favourable gene. *Genet Res*. 23:23–35.
- Mendez FL, Watkins JC, Hammer MF. 2012. Global genetic variation at OAS1 provides evidence of archaic admixture in Melanesian populations. *Mol Biol Evol*. 29:1513–1520.
- Mesuret G, Khabbazzadeh S, Bischoff AM, Safory H, Wolosker H, et al. 2018. A neuronal role of the Alanine-Serine-Cysteine-1 transporter (SLC7A10, Asc-1) for glycine inhibitory transmission and respiratory pattern. *Sci Rep*. 8:8.
- Nakamura F, Kumeta K, Hida T, Isono T, Nakayama Y, et al. 2014. Amino- and carboxyl-terminal domains of Filamin-A interact with CRMP1 to mediate Sema3A signalling. *Nat Commun*. 5:5.
- Ng L, Chow AK, Man JH, Yau TC, Wan TM, et al. 2018. Suppression of Slit3 induces tumor proliferation and chemoresistance in hepatocellular carcinoma through activation of GSK3 β / β -catenin pathway. *BMC Cancer*. 18:13.
- Nicholson G, Smith AV, Jónsson F, Gústafsson Ó, Stefánsson K, et al. 2002. Assessing population differentiation and isolation from single-nucleotide polymorphism data. *J R Stat Soc B*. 64:695–715.
- Ozioloz EM, Reid NM, Yair S, Lee KM, VerPloeg SG, et al. 2019. Adaptive introgression enables evolutionary rescue from extreme environmental pollution. *Science*. 364:455–457.
- Palaćin M, Errasti-Murugarren E, Rosell A. 2016. Heteromeric amino acid transporters. In search of the molecular bases of transport cycle mechanisms. *Biochem Soc Trans*. 44:745–752.
- Pasterkamp RJ, Peschon JJ, Spriggs MK, Kolodkin AL. 2003. *Semaphorin 7A* promotes axon outgrowth through integrins and MAPKs. *Nature*. 424:398–405.
- Petr M, Pääbo S, Kelso J, Vernot B. 2019. Limits of long-term selection against Neandertal introgression. *Proc Natl Acad Sci USA*. 116:1639–1644.
- Prüfer K, Filippo CD, Grote S, Mafessoni F, Korlevi P, et al. 2017. A high-coverage Neandertal genome from Vindija Cave in Croatia. *Science*. 358:655–658.
- Ptáček LJ, Tawil R, Griggs RC, Engel AG, Layzer RB, et al. 1994. Dihydropyridine receptor mutations cause Hypokalemic Periodic Paralysis. *Cell*. 77:863–868.
- Pusapati GV, Hughes CE, Dorn KV, Zhang D, Sugianto P, et al. 2014. EFCAB7 and IQCE regulate Hedgehog signaling by tethering the EVC-EVC2 complex to the base of primary cilia. *Dev Cell*. 28:483–496.
- Quach H, Rotival M, Pothlichet J, Kelso J, Albert ML, et al. 2016. Genetic adaptation and neandertal admixture shaped the immune system of human populations. *Cell*. 167:643–656.
- Racimo F, Berg JJ, Pickrell JK. 2018. Detecting polygenic adaptation in admixture graphs. *Genetics*. 208:1565–1584.
- Racimo F, Marnetto D, Huerta-Sánchez E. 2017. Signatures of archaic adaptive introgression in present-day human populations. *Mol Biol Evol*. 34:296–317.
- Racimo F, Sankararaman S, Nielsen R, Huerta-Sánchez E. 2015. Evidence for archaic adaptive introgression in humans. *Nat Rev Genet*. 16:359–371.
- Refoyo-Martínez A, Da Fonseca RR, Halldórsdóttir K, Árnason E, Mailund T, et al. 2019. Identifying loci under positive selection in complex population histories. *Genome Res*. 29:1506–1520.
- Ruiz-Perez VL, Ide SE, Strom TM, Lorenz B, Wilson D, et al. 2000. Mutations in a new gene in Ellis-van Creveld syndrome and Weyers acrodermal dysostosis. *Nat Genet*. 24:283–286.
- Sachdeva H, Barton NH. 2018. Replicability of introgression under linked, polygenic selection. *Genetics*. 210:1411–1427.
- Samanta S, Li Y-J, Weir BS. 2009. Drawing inferences about the coancestry coefficient. *Theor Popul Biol*. 75:312–319.
- Sams AJ, Dumaine A, Nédélec Y, Yotova V, Alfieri C, et al. 2016. Adaptively introgressed Neandertal haplotype at the OAS locus functionally impacts innate immune responses in humans. *Genome Biol*. 17:15.
- Sankararaman S, Mallick S, Dannemann M, Prüfer K, Kelso J, et al. 2014. The genomic landscape of Neandertal ancestry in present-day humans. *Nature*. 507:354–357.
- Schraiber JG, Evans SN, Slatkin M. 2016. Bayesian inference of natural selection from allele frequency time series. *Genetics*. 203:493–511.
- Schumer M, Xu C, Powell DL, Durvasula A, Skov L, et al. 2018. Natural selection interacts with recombination to shape the evolution of hybrid genomes. *Science*. 360:656–660.
- Setter D, Mousset S, Cheng X, Nielsen R, DeGiorgio M, et al. 2020. VolcanoFinder: Genomic scans for adaptive introgression. *PLoS Genet*. 16:e1008867.
- Shchur V, Svedberg J, Medina P, Corbett-Detig R, Nielsen R. 2020. On the distribution of tract lengths during adaptive introgression. *G3*. 10:3663–3673.
- Sikora M, Seguin-Orlando A, Sousa VC, Albrechtsen A, Korneliussen T, et al. 2017. Ancient genomes show social and reproductive behavior of early Upper Paleolithic foragers. *Science*. 358:659–662.
- Skoglund P, Malmström H, Raghavan M, Storå J, Hall P, et al. 2012. Origins and genetic legacy of neolithic farmers and hunter-gatherers in Europe. *Science*. 336:466–469.
- Steinrücken M, Spence JP, Kamm JA, Wieczorek E, Song YS. 2018. Model-based detection and analysis of introgressed Neandertal ancestry in modern humans. *Mol Ecol*. 27:3873–3888.
- Suzuki K, Okuno T, Yamamoto M, Pasterkamp RJ, Takegahara N, et al. 2007. *Semaphorin 7A* initiates T-cell-mediated inflammatory responses through $\alpha 1\beta 1$ integrin. *Nature*. 446:680–684.
- Tanabe T, Beam KG, Adams BA, Niidome T, Numa S. 1990. Regions of the skeletal muscle dihydropyridine receptor critical for excitation-contraction coupling. *Nature*. 346:567–569.
- Telis N, Aguilar R, Harris K. 2020. Selection against archaic hominin genetic variation in regulatory regions. *Nat Ecol Evol*. 4:1558–1566.
- The 1000 Genomes Project Consortium 2015. A global reference for human genetic variation. *Nature*. 526:68–74.
- The Heliconius Genome Consortium 2012. Butterfly genome reveals promiscuous exchange of mimicry adaptations among species. *Nature*. 487:94–98.

- The International HapMap Consortium 2007. A second generation human haplotype map of over 3.1 million SNPs. *Nature*. 449:851–861.
- Uecker H, Setter D, Hermisson J. 2015. Adaptive gene introgression after secondary contact. *J Math Biol*. 70:1523–1580.
- Vernot B, Akey JM. 2014. Resurrecting surviving neandertal lineages from modern human genomes. *Science*. 343:1017–1021.
- Weir BS, Carolina N, Hill WG. 2002. Estimating F statistics. *Annu Rev Genet*. 36:721–750.
- Whitney KD, Randell RA, Rieseberg LH. 2006. Adaptive introgression of herbivore resistance traits in the weedy sunflower *Helianthus annuus*. *Am Nat*. 167:794–807.
- Yoshikawa M, Asaba K. 2020. Single-nucleotide polymorphism rs4664308 in PLA2R1 gene is associated with the risk of idiopathic membranous nephropathy: a meta-analysis. *Sci Rep*. 10:8.
- Yu Y, Chen S, Lu GF, Wu Y, Mo L, et al. 2014. Alfvbeta6 is required in maintaining the intestinal epithelial barrier function. *Cell Biol Int*. 38:777–781.
- Yu-Kemp HC, Kemp JP, Briehner WM. 2017. CRMP-1 enhances EVL-mediated actin elongation to build lamellipodia and the actin cortex. *J Cell Biol*. 216:2463–2479.
- Zeberg H, Pääbo S. 2021. A genomic region associated with protection against severe COVID-19 is inherited from Neandertals. *Proc Natl Acad Sci USA*. 118:e2026309118.
- Zhang X, Smits AH, van Tilburg GB, Jansen PW, Makowski MM, et al. 2017. An interaction landscape of ubiquitin signaling. *Mol Cell*. 65:941–955.
- Zhang X, Witt K, Ko A, Yuan K, Xu S, et al. 2020. The history and evolution of the Denisovan-EPAS1 haplotype in Tibetans. *bioRxiv*. 1–21.
- Zheng Y, Peng M, Wang Z, Asara JM, Tyner AL. 2010. Protein tyrosine kinase 6 directly phosphorylates Akt and promotes Akt activation in response to epidermal growth factor. *Mol Cell Biol*. 30:4280–4292.
- Zhou S, Butler-Laporte G, Nakanishi T, Morrison D, Afilalo J, et al. 2020. A Neanderthal OAS1 isoform protects against COVID-19 susceptibility and severity: results from Mendelian randomization and case-control studies. *medRxiv*. 2020.10.13.20212092.

Appendix A

A.1 Simulation details

A.1.1 Selection simulations

Keeping demography constant, we simulated different combinations of t_b , s , t_s , and set of selected populations (Supplementary Tables S4 and S5 in Supplementary File S1). After the divergence between the ancestors of Neanderthals and modern humans, we added a mutation with $s=0.025$ to a single Neanderthal haplotype at the selected site. It became neutral before admixture with modern humans. We restarted any runs in which the selected mutation did not reach fixation in the Neanderthal population prior to admixture, was not segregating in all selected populations at the onset of selection, was lost from any selected population during the sweep, was not at greater than 20% frequency in at least one selected population at the sweep finish, or was not at greater than 20% frequency in either the European or East Asian population sampled in the present day. The 20% frequency cutoff was motivated by the same criteria with which adaptive introgression candidates were identified (Racimo et al. 2017).

As frequency trajectories in SLiM are stochastic, we chose combinations of s and t_s that would target different final frequencies (x_s) under a deterministic trajectory. For each s and target x_s , we calculated t_s according to Equation (4) and assuming that the frequency at the sweep start is the admixture proportion we used in

simulations ($g=0.02$). However, because our simulations condition on the selected mutation segregating in all selected populations at the selection onset, the average starting frequency in our simulations tended to be greater than g , and increased the later selection began. Therefore, our chosen t_s on average led to rather high x_s in cases of late selection, even when we targeted $x_s=0.2$. Alternative approaches to target a certain low x_s in our simulations would have led to many simulations where a target x_s was reached due to genetic drift rather than natural selection. Those approaches are inappropriate for us, as we focus on understanding the history of natural selection, rather than distinguishing between genetic drift and natural selection.

Of all polymorphic sites that were not filtered, we randomly down-sampled the data set to 12,000 sites, which is close to the number of sites analyzed for data in our application with loci of the same size. As described in the main text, we used a different sampling approach for simulations generated to identify the $CLR_{t_b>0}$ boundary for the regions we analyzed in our application.

A.1.2 Neutral simulations

We generated simulations under neutrality using the same demographic history as the selection simulations so that we could estimate the “genome-wide” neutral F . We simulated 2000 independent trees by separately recording the tree sequences of 1 bp loci and subsequently overlaying multiple mutations. From these 2000 independent “loci” we estimated the neutral F from 19,323 sites after filtering.

A.2 Estimating probabilities of coalescing from allele frequencies

Using allele frequencies, we obtain unbiased estimates of F , our genome-wide neutral probabilities of coalescing. Rearranging the equation providing the covariance of the change in population allele frequencies from the ancestral population, the probability a pair of ancestral lineages coalesce before the root is

$$f_{ij} = \frac{\text{Cov}[\Delta x_i, \Delta x_j]}{x_a(1-x_a)}. \quad (\text{A.1})$$

Since $\Delta x_i = x_i - x_a$, the above covariance equals $\mathbb{E}_i[x_{il}x_{jl}] - \mathbb{E}_i[x_{al}^2]$. In our scripts, we estimate f_{ij} using an equivalent expression that directly averages over the two possible alleles at a locus,

$$f_{ij} = 1 - \frac{\mathbb{E}_i[x_{il}(1-x_{jl}) + x_{jl}(1-x_{il})]}{\mathbb{E}_i[2x_{al}(1-x_{al})]}. \quad (\text{A.2})$$

The numerator of the fraction represents the probability of sampling different alleles from population i and population j , which we estimate by taking its average over all loci. When i and j correspond to the populations that root the tree, this average serves as the unbiased estimator of the denominator, $2x_a(1-x_a)$. When $i=j$, such that the numerator represents the expected heterozygosity within a population, we must account for how allele frequencies change when we sample without replacement, i.e., we must remove the finite sampling bias. If n_{il} represents the sample size of population i at locus l , then when $i=j$,

$$\mathbb{E}_i[x_{il}(1-x_{jl}) + x_{jl}(1-x_{il})] = \frac{1}{L} \sum_{l=1}^L \frac{n_{il}}{n_{il}-1} (2x_{il}(1-x_{il})). \quad (\text{A.3})$$

When validating that our model predictions match the probabilities of coalescing in our selection simulations, such as in Figure 4, we repeated the above procedure for sets of loci binned by genetic distance to the selected site.

A.3 Inference details

When we calculate likelihoods of population allele frequencies, we need to account for the finite sampling bias in our estimates. Previously, we described the variance in the true population allele frequency due to genetic drift, however there is an additional variance in our estimates due to sampling. Since the count of a certain allele type in a sample is binomially distributed according to the true allele frequency, then with sample size n_{il} the variance in allele frequency due to sampling is expected to be approximately $x_a(1-x_a)\frac{1}{n_{il}}$. Thus, the total variance in the population allele frequency is $x_a(1-x_a)(f_{il} + \frac{1}{n_{il}})$. Therefore, when we calculate the probability of population allele frequencies at a locus, we first modify $\mathbf{F}^{(S)}$ by adding $\frac{1}{n_{il}}$ along the diagonal for each population i .

Because we do not know the ancestral allele frequency at each locus, we approximate it with the mean across our sampled population allele frequencies,

$$\bar{x}_l = \frac{1}{k_l} \sum_{i=1}^{k_l} x_{il}, \quad (\text{A.4})$$

where k_l is the number of populations with allele frequency data at locus l . We acknowledge that in our case \bar{x}_l will be biased toward European allele frequencies. With allele frequencies now distributed relative to each other, we accordingly modify the population allele frequencies and covariance matrix by mean-centering them. The new population allele frequencies therefore

represent deviations from their mean, and a negative covariance among pairs of populations implies that we predict their allele frequencies to be on opposite sides of this mean. As \bar{x}_l uses information from all populations, we lose a degree of freedom and thus drop information from a single population. Our resulting mean-centered allele frequencies are

$$\vec{x}_l' = \mathbf{T}_l \bar{x}_l \quad (\text{A.5})$$

and our mean-centered covariance matrix is

$$\mathbf{F}^{(S)'} = \mathbf{T}_l \mathbf{F}^{(S)} \mathbf{T}_l^T, \quad (\text{A.6})$$

where \mathbf{T}_l is the $k_l - 1$ by k_l matrix with $k_l - \frac{1}{k_l}$ on the diagonal and $-\frac{1}{k_l}$ elsewhere. Therefore, we calculate the probability of mean-centered allele frequencies at a locus as

$$\begin{aligned} P(\vec{x}_l' | r_l, \mathbf{F}, A_G, s, t_b, x_s) \\ \approx \mathcal{N}(\vec{x}_l' | 0, \bar{x}_l(1-\bar{x}_l) \mathbf{F}^{(S)'}(r_l, \mathbf{F}, A_G, s, t_b, x_s)). \end{aligned} \quad (\text{A.7})$$

We take the same approach when calculating probabilities under the null model.

To allow for computational efficiency, we first bin each site's absolute genetic distances to the selected site into 10^{-3} cM intervals. Each bin's midpoint value is then used as the representative recombination rate to predict $\mathbf{F}^{(N)}$ and $\mathbf{F}^{(S)}$ for each parameter combination. When we calculate the probability of allele frequencies at each site, we use its bin's representative \mathbf{F} and remove the rows and columns corresponding to unsampled populations at the site prior to sample size correction and mean centering.