

**The effect of genomic inversions on estimation of population genetic
parameters from SNP data**

Nafisa-Katrin Seich al Basatena^{*}, Clive J. Hoggart^{**},

Lachlan J. Coin^{***}, Paul F. O'Reilly^{±§}

^{*} Department of Immunology, Imperial College London, W2 1PG, United Kingdom

^{**} Department of Pediatrics, Imperial College London, W2 1PG, United Kingdom

^{***} Department of Genomics of Common Disease, Imperial College London, W12
0NN, United Kingdom

[±] Department of Epidemiology and Biostatistics, Imperial College London, W2 1PG,
United Kingdom

[§] Medical Research Council-Health Protection Agency Centre for Environment and
Health, Imperial College London, W2 1PG, United Kingdom.

Short Title: The effect of inversions on SNP data inference

Corresponding author:

Paul F. O'Reilly

Department of Epidemiology and Biostatistics

Imperial College London

St Mary's Campus

London, W2 1PG, United Kingdom

Email: paul.oreilly@imperial.ac.uk

Keywords: Inversions; Population genetics; Principle Component Analysis;
Recombination; SNP data inference

ABSTRACT

In recent years it has emerged that structural variants have a substantial impact on genomic variation. Inversion polymorphisms represent a significant class of structural variant, and despite the challenges in their detection, data on inversions in the human genome are increasing rapidly. Statistical methods for inferring parameters such as the recombination rate and selection coefficient have generally been developed without accounting for the presence of inversions. Here we exploit new software for simulating inversions in population genetic data, invertFREGENE, to assess the potential impact of inversions on such methods. Using data simulated by invertFREGENE, as well as real data from several sources, we test whether large inversions have a disruptive effect on widely applied population genetics methods for inferring recombination rates, for detecting selection, and for controlling for population structure in genome-wide association studies (GWAS). We find that recombination rates estimated by LDhat are biased downward at inversion loci relative to the true contemporary recombination rates at the loci but that recombination hotspots are not falsely inferred at inversion breakpoints as may have been expected. We find that the iHS method for detecting selection appears robust to the presence of inversions. Finally, we observe a strong bias on the genome-wide results of principle components analysis (PCA), used to control for population structure in GWAS, in the presence of even a single large inversion, confirming the necessity to thin SNPs by LD at large physical distances in order to obtain unbiased results.

INTRODUCTION

Inversion polymorphisms are a copy neutral form of structural variant, which although largely uncatalogued in human populations, are likely a common feature in the human genome (KIDD *et al.* 2008). The study of inversions in human populations is in its infancy but early indications suggest that their influence on genetic variation (STEFANSSON *et al.* 2005) and phenotypic traits (TANTISIRA *et al.* 2008) may be significant. One of the first large inversions (~1.1 Mb) to have been discovered in the human genome is located at the *MAPT* locus on chromosome 17 (BOETTGER *et al.* 2012; STEINBERG *et al.* 2012). The inverted sequence was found to be associated with elevated fertility in women and an increased recombination rate genome-wide (STEFANSSON *et al.* 2005). Another study found that three of five detected inversions were located at sites of recurrent microdeletion associated with human disease (KIDD *et al.* 2008) adding to several studies linking inversions with susceptibility to disease (ANTONACCI *et al.* 2009; KOOLEN *et al.* 2006; LAKICH *et al.* 1993; OSBORNE *et al.* 2001; TANTISIRA *et al.* 2008). Figure S1 shows the distribution of known human inversions according to the TCAG database.

To complement the investigation of the biological impact of inversions, it is important to establish their effect on methods that exploit genetic variation data to infer population genetic parameters without accounting for the presence of inversions in the genome. While there have been indications that inversions may bias statistical methods for detecting selection (DENG *et al.* 2008; O'REILLY *et al.* 2008), measuring population structure (PRICE *et al.* 2008), and estimating recombination rates (WEGMANN *et al.* 2011), there has been no formal study to test this. Here we perform such a study. We simulate genome-wide SNP data with and without inversions using invertFREGENE (O'REILLY *et al.* 2010) and then, using

these data, estimate the effect of inversions on leading methods for inferring recombination rates, detecting recent selection, and for controlling for population stratification in genome-wide association studies (GWAS). We investigate whether there is support for our simulation findings in real data, using several resources including the HapMap (FRAZER *et al.* 2007) and Wellcome Trust Case Control Consortium (WTCCC 2007) data sets. While our focus is on humans, our results may generalise to other species for which these methods are applied and that have inversions with similar characteristics.

METHODS

To test the effect of inversions on estimates of population genetics parameters we simulate genetic data with and without inversions, using invertFREGENE (O'REILLY *et al.* 2010). We test the effect of inversions on LDhat for estimating recombination rates (MCVEAN *et al.* 2004; MYERS *et al.* 2005), the iHS statistic for detecting selection (VOIGHT *et al.* 2006), which extends the Extended Haplotype Homozygosity (EHH) method (SABETI *et al.* 2002), and on principal components analysis (PCA) for controlling for population structure in GWAS (PRICE *et al.* 2006). For each we simulate slightly different data sets. We first describe the data simulated to test LDhat, and then detail how the data simulated for the subsequent two studies differ. We also describe the use of real data in the investigation relating to PCA, while we describe the real data analyses relating to the investigations into recombination and selection in the Results section.

To test the effect of inversions on LDhat recombination rate estimates, data are simulated by invertFREGENE according to a neutrally evolving panmictic population, under a hotspot model of recombination developed to reflect the principal

features of human genome variation data (SCHAFFNER *et al.* 2005). The average recombination and mutation rates were set as 1.25×10^{-8} and 2.5×10^{-8} per site, per sequence, per generation, respectively, reflecting human genome-wide averages (JOBLING 2003; MATISE *et al.* 2007). Gene conversions are modelled with a fixed tract length of 500 bp and occur independently, across inverted and non-inverted sequence, at a constant rate of 2×10^{-8} per start site, per sequence, per generation (PADHUKASAHASRAM *et al.* 2004). We generated one hundred 2 Mb pseudo-chromosomes, sampled from a simulation with a population size of 20000 chromosomes ($N=10000$), with an inversion simulated to span a 500 kb locus at the centre of each 2 Mb region (0.75-1.25 Mb). According to the invertFREGENE model recombination events are not viable within the span of an inversion for 'individuals' that are heterozygous for inversion status (gene conversions are viable), but the recombination process continues normally directly outside the inversion. While the suppression of recombination may in fact drop-off gradually adjacent to breakpoints (NAVARRO *et al.* 1997), our discrete-breakpoint assumption has no effect on the conclusions that we reach (see Discussion).

invertFREGENE was first run, without modelling an inversion, until an equilibrium state was reached (50000 generations) so that the data have levels of heterozygosity and linkage disequilibrium reflecting that observed in humans (HOGGART *et al.* 2007) and then these data were used as input to a second invertFREGENE run that modelled inversions. An inversion mutation occurs at the specified locus in the first generation and segregates in the population of sequences until it is either lost or reaches the specified frequency; if the latter, then in a large population the final frequency will be very close to that specified. The age of the inversion is a function of its frequency and genetic drift, with the relationship the

same as for a single variant and thus predicted by standard population genetic theory (KIMURA 1985). If an inversion exceeds a frequency of 5% and is then lost in the population, the inversion simulations are started with a new seed to ensure negligible influence of historic inversions in the final data. Thus the final data reflect a locus with only a single inversion mutation in the history of the population. This simulation was repeated to produce inversion data with the inverted sequence at different frequencies, from 10% to 90% (with 10% increments). 96 individuals were then sampled from the final generation because the largest LDhat likelihood table available is based on 192 chromosomes. The data were then thinned from sequence-level density by first removing SNPs with $MAF < 1\%$, to reflect the common SNPs targeted by genotyping arrays and large-scale genotyping projects, and then further thinned at random to produce SNP data at a density reflecting that of HapMap project (Phase II+III) data (FRAZER *et al.* 2007). We then applied LDhat to each of the simulated data sets to obtain recombination rate estimates across the loci (see File S1). We compare the performance of LDhat on the data simulated without inversions (using data outside the central 500 kb inversion locus, unaffected by the inversion according to our model) and at inversion loci at the range of inverted sequence frequencies. We also test whether false recombination hotspots are inferred at the inversion breakpoints.

To examine the distribution of the iHS statistic, developed to detect loci subject to selection (VOIGHT *et al.* 2006; see File S2) at inversion loci, we used the simulated inversion data described above. In order to compare the iHS distribution calculated from these neutrally evolved inversion loci with that from loci subject to selection we also generated 100 sets of 1 Mb locus genotype data under different strengths of positive selection using SelSim (SPENCER and COOP 2004; see File S3).

We considered neutrality and five intensities of positive selection: (i) neutral selection, (ii) $s=0.375\%$, (iii) $s=1\%$, (iv) $s=2\%$, (v) $s=5\%$ and (vi) $s=10\%$, and set the frequency of the selected allele in the population to be 70%, at which the method has approximately greatest power (VOIGHT *et al.* 2006). We used the same recombination map in the model of both the invertFREGENE and SelSim simulations, and all SNP data was at HapMap (Phase II+III) density. We compared the distribution of the iHS statistic at loci containing inversions to those without inversions under neutrality and subject to selection, considering inversions of length 500 kb and 4 Mb.

We performed PCA (PRICE *et al.* 2006) on both simulated and real genome-wide SNP data to estimate the effect of inversions on inferring population structure. The real data comprised 1480 population-based (control) samples from the 1958 British Birth Cohort used in the Wellcome Trust Case-Control Consortium study (WTCCC 2007), genotyped using the 550k Affymetrix chip. While the real data contain many known inversion loci, allowing us to use these data alone to assess the impact of inversions on PCA, we also combined the data with a simulated inversion whose characteristics we could control. The inversion locus was simulated as described above to produce inversions of length 500 kb and 4 Mb (with 50 kb non-inverted locus flanking each end), and we sampled 1480 'individuals' from the simulated population to match the sample size of the WTCCC data. The data were simulated to produce replicates of an inversion with inverted sequence frequencies of 20%, 50% and 80%, and the SNP data were then thinned as above but to the SNP density of the 550k Affymetrix chip. With the simulated and real data matched in this way, we could then insert the simulated inversions into the real data and perform PCA on the entire data to assess the impact of the simulated inversions. We

also combined the simulated inversion data with simulated genome-wide non-inverted and unstructured data and subsequently performed PCA. We add a single inversion locus to each genome-wide simulated/real data set. The data on which we performed PCA were of four types: (i) unstructured, simulated (ii) unstructured, simulated, incorporating simulated inversions of different sizes/frequencies, (iii) real (WTCCC), structured and (iv) real, structured, incorporating simulated inversions of different sizes/frequencies. We assessed the impact of inversions on the results by inspecting the SNP loading contributions (JOLLIFFE 2002) to each of the first 6 principle components.

RESULTS

Recombination rate estimation: Patterns of linkage disequilibrium (LD) in samples of high-density SNP data can be used to infer recombination rates over the genome at a fine-scale (STUMPF and McVEAN 2003). Recombination events between two SNPs tend to reduce the association between the alleles of the SNPs and thus reduce LD. LDhat is an approximate-likelihood method that infers historical recombination according to the breakdown of LD between pairs of SNPs in a sample (McVEAN *et al.* 2004). LDhat has been applied to estimate recombination rates and detect recombination hotspots across the human genome, contrast recombination rates between humans and chimpanzees, while LDhat-derived recombination maps are used in most GWAS to verify the validity of association signals at loci (MYERS *et al.* 2005; NEWTON-CHEH *et al.* 2009; WINCKLER *et al.* 2005). Since LDhat exploits LD patterns, the presence of inversions, which have been found to be associated with strong long-range LD in humans (DENG *et al.* 2008; O'REILLY *et al.* 2010; PRICE *et al.* 2008) and *Drosophila* (MUNTÉ *et al.* 2005; SCHAEFFER *et al.* 2003), may affect the

accuracy of estimates. We note that because inversions suppress viable recombination between diploid chromosomes at loci that have inverted and non-inverted sequences, the true contemporary recombination rate is affected by inversions, so we expect the estimates to be affected accordingly. However, the recombination rate that is estimated by LDhat exploits LD that is formed over many generations, while the suppression of recombination is a function of the present inverted sequence frequency, therefore the estimated rate may not reflect the contemporary rate at inversion loci. A recent study (WEGMANN *et al.* 2011) found that many of the largest discrepancies between admixture-based and LDhat-based recombination rate estimates correspond to inversion loci, which may be explained by inaccurate LDhat estimates at inversions.

We now assess LDhat on the data simulated with an inversion, where the inverted sequence has reached a certain frequency (10-90%) in the sample. Figure 1 shows that LDhat tends to produce systematically lower recombination rate estimates at inversion loci compared to loci without inversions. Table 1 shows that there is a general trend of reducing average recombination rate estimates with increasing inverted sequence frequency from 10% to 80%, and a slight increase in estimates at 90%. For example, the recombination rate estimated by LDhat is a mean ~55% lower than the mean simulated rate within inversions of frequency 20%, while it is ~86% lower within inversions of frequency 80%. Table S1 shows the median reductions in rate estimates at the inversion loci, although we focus on the mean rates because they reflect the reduction over any physical distance and because the genetic distance over a region is usually of interest when considering recombination (see Table S1 caption for more explanation). We also compare the mean LD-based estimates with those that would be expected from pedigree-based

estimates, which should reflect contemporary recombination rates (Table 1). These results indicate that population-based recombination rates, estimated from LD, are biased downward relative to contemporary pedigree-based estimates. See the Discussion for further consideration of these findings.

To find out whether these results are supported by real data we first recorded the locations of all autosomal inversions ≥ 500 kb in the TCAG database (URL: <http://projects.tcag.ca/variation/>), which lists all reliably discovered human inversions. This provided 16 large real inversions to investigate. Next we downloaded the latest pedigree-based recombination map (KONG *et al.* 2010) and the LDhat-based recombination map from the HapMap project website (URL: <http://www.hapmap.org/>), and recorded the ratio of the two rates computed over 100 kb sliding windows (10 kb increments) across the entire (autosomal) genome. To assess whether the ratio of the two rates was unusual at the 16 large inversions we took the maximum and minimum values of these ratios within each inversion (multiple windows within each inversion) and compared them to an empirical null distribution formed by taking the maximum and minimum values from 1000 randomly selected loci of the same length as the inversion being tested; for example, the maximum and minimum ratios within a 700 kb inversion are compared to those from 1000 randomly selected 700 kb regions. We find strong evidence for a downward bias in LDhat recombination rates compared to pedigree-based rates ($P=0.0005$) and also some evidence ($P=0.04$) for an upward bias in LDhat rates (Table 2). It appears that there may be a downward bias at some inversion loci, an upward bias at others, and possibly a bias in both directions at some inversions given that both high and low values of the ratio of the rates were observed at one of the inversions (Table 2; chr16, 21.5-22.6 Mb). Therefore, the evidence from real data is highly

supportive of our simulation-based predications that LDhat produces a downward bias in recombination rate estimates, while there is also evidence that an upward bias may exist (see Discussion).

Next we test the hypothesis that recombination hotspots are falsely inferred by LDhat at inversion breakpoints. For large inversions, where the inverted sequence is at high frequency, there will likely have been many recombination events between SNPs spanning an inversion breakpoint and physically close according to the reference map of the genome, because their true separation will be large in all inverted sequences. Therefore, LD may be low between such SNPs despite their apparent proximity, so we may expect recombination hotspots to be inferred at inversion breakpoints irrespective of the true recombination intensity. We investigate this by examining the difference between simulated and estimated recombination rates between pairs of SNPs that span inversion breakpoints, and comparing those to the difference in rates between pairs of SNPs that are inside or outside inversion loci. Figure S2 shows that LDhat does not systematically overestimate recombination rates at the inversion breakpoints in our data, and therefore, according to these simulations, there is not an enrichment of falsely inferred recombination hotspots at inversion breakpoints as hypothesised.

Detecting selection: Since the availability of genome-wide SNP data, numerous scans to detect loci subject to recent positive selection across the human genome have been performed, using a variety of methods (AKEY 2009), in the hope of revealing mechanisms of adaptive change in human evolution and highlighting genes of medical importance. One such method is the Integrated Haplotype Score (iHS) test developed by VOIGHT *et al.* (2006), which extends the leading approach to

detecting recent positive selection based on the Extended Haplotype Homozygosity (EHH) (FRAZER *et al.* 2007; SABETI *et al.* 2002; SABETI *et al.* 2007). The assumption underlying this approach is that haplotypes at a locus that harbour a beneficial allele will have little breakdown given their frequency compared to the alternative haplotypes at the locus, due to reduced opportunity to undergo recombination as they are “swept” to high frequency in relatively few generations. However, there may be a similar disparity in the breakdown of haplotypes at a neutrally evolving inversion locus as a result of the suppression of recombination between sequences in the opposite orientation. Haplotypes in the inverted orientation may have been subject to distinctly fewer recombinations than non-inverted haplotypes because at low frequencies the inverted sequence will generally be present in a diploid pair with a non-inverted sequence, with which it cannot recombine, while most non-inverted sequence will be in homozygous form and will thus recombine as usual. Therefore, inversion loci may represent a potential source of false-positive findings for the EHH approach to detecting selection, which we test here in the form of the iHS statistic (VOIGHT *et al.* 2006).

First, we ran the iHS method on the simulated neutral inversion data described in the Methods, computing the standardised iHS statistic for all SNPs. We then applied the iHS to the data simulated under neutrality and five intensities of positive selection: (i) neutral selection, (ii) $s=0.375\%$, (iii) $s=1\%$, (iv) $s=2\%$, (v) $s=5\%$ and (vi) $s=10\%$, for selective sweeps where the beneficial allele has reached 70% in the population. Figure 2 displays the distributions of the iHS under neutrality and the different selection intensities, and that corresponding to inversion loci of 500 kb and 4 Mb where the inverted sequence has reached 60%, which is the inverted frequency showing greatest deviation from the iHS null (Figure S3). These simulation

results indicate that while there can be a noticeable deviation from the iHS null distribution, particularly in the lower tail, for large inversion loci, the deviation for even the smallest selection coefficient ($s=0.375\%$) tested here is considerably larger. Therefore, we would not expect the extreme tails of the iHS statistic computed genome-wide to be particularly enriched for inversion loci unless there were many more large inversions across the genome than beneficial alleles with selection coefficient $s \geq 0.375$. To investigate evidence for this in real data we compared the iHS values at the 16 large inversions (≥ 500 kb) from the TCAG database with those from an empirical null distribution of the iHS in the European HapMap population (CEU) , having retrieved genome-wide iHS values from <http://haplotter.uchicago.edu/>. To make this comparison we followed the same procedure as described for recombination rates in the previous section. We find no significant evidence for enrichment of extreme iHS values ($P = 0.58$) at the 16 inversion loci (Table 2), indicating that the iHS statistic is reasonably robust to the presence of inversion polymorphisms, as suggested by our simulation results.

Controlling for population structure: Systematic non-random mating in populations results in genetic stratification. A primary source of systematic mating in humans is geographic separation (NOVEMBRE *et al.* 2008). The resulting genetic stratification can be problematic for genome-wide association studies (GWAS) because the phenotype under study may have environmental risk factors stratified by geography, producing the potential for false-positive associations between SNP genotypes and the phenotype because the genotypes may be associated with the phenotype via the risk factors. A routinely used approach for controlling for such confounding is to perform Principle Components Analysis (PCA) on the genotype data and to subsequently include the top principle components (PCs) as covariates

in the GWAS analysis (PRICE *et al.* 2006). The PCs provide genome-wide summaries of the genotype data for each individual and are thus reflective of the genetic similarity between the individuals, which should correspond broadly to geographic separation. Each independent locus should contribute equally to this estimate because each provides equal evidence in terms of relatedness. However, some loci have a large number of SNPs in high LD, which up-weight their contribution and make the PCs less reliable. Using simulated and real data, O'REILLY *et al.* (2010) show that inversion loci can have a particularly high proportion of SNPs either perfectly or highly correlated with each other, and since inversions can be megabases (MBs) in size, then they may compromise the results from PCA. While PCA is often performed on genome-wide data that has been thinned to remove correlated SNPs, this is generally only performed locally using sliding windows of, say, 50 kb or 100 kb, and in some cases is not performed at all. We investigate the effect of the presence of large inversions (≥ 500 kb) on PCA when SNP data are not thinned.

We performed PCA on both simulated and real genome-wide SNP data to estimate the effect of inversions on inferring population structure. The data on which we performed PCA were of four types: (i) unstructured, simulated (ii) unstructured, simulated, incorporating simulated inversions of different sizes/frequencies, (iii) real (WTCCC), structured and (iv) real, structured, incorporating simulated inversions of different sizes/frequencies. We assessed the impact of inversions on the results by inspecting the SNP loading contributions (JOLLIFFE 2002) to each of the first 6 principle components.

When PCA is performed on the simulated unstructured data the SNP loadings contributing to each of the first 6 principal components are highly uniform across the

genome. This finding suggests that the correlation between SNPs modelled in these simulations, which generates regions of high and low LD (SCHAFFNER *et al.* 2005), is not sufficient to compromise the PCA. Next we assess the loadings when a single simulated inversion locus is added to these data. When a 500 kb inversion is added to the data, the loadings, again, have a similar distribution across the genome. However, the inclusion of a 4 Mb inversion has a substantial influence, with the 1st PC strongly affected by the data of the inversion locus (Figure S4), while all other loci for each of the first 6 PCs display approximately the same distribution as each other, as observed when an inversion is not present.

Next, we perform PCA on the WTCCC data with and without the inclusion of a simulated inversion locus. Figure 3A shows the genome-wide SNP loadings from the PCA on the data without simulated inversions, and reveals that the 3rd, 4th and 5th PCs are strongly influenced by localised sources of variation. The three genomic regions that have a marked impact on PCs 3, 4 and 5 are, respectively, located at: (i) the major histocompatibility complex (MHC), which is a 3.8 Mb immune-dense region found in chromosome 6p23, exhibiting particularly strong and extended LD blocks (KELLEY *et al.* 2008; YU *et al.* 2005), (ii) a 3.8 Mb inversion polymorphism on chromosome 8p23 (DENG *et al.* 2008) and (iii) a 1.1 Mb inversion polymorphism on chromosome 17q21 (STEFANSSON *et al.* 2005). Therefore, these results provide empirical evidence that inversions can disrupt PCA performed on genome-wide SNP data. When a 500 kb simulated inversion is added to the real data there is little apparent impact on the first 6 PCs as assessed by SNP loadings (data not shown), while the inclusion of a 4 Mb inversion can be a distinct source of variation in the top PCs in these data (Figure 3B and Figure S5). The frequency, as well as the size, of

the inversion appears to be important, such that the higher the frequency of the inverted sequence, the greater the tendency for the lower order PCs to be disrupted.

DISCUSSION

Statistical methods that use population genetic data to infer parameters of interest, such as intensity of recent natural selection, have been developed over previous decades. Through the use of more appropriate models, such as the coalescent model (HUDSON 1983), and greater knowledge of genetic variation data, such as the discovery that most recombination in the human genome occurs in recombination hotspots (JEFFREYS *et al.* 2001; MYERS *et al.* 2005), methods for inferring parameters of interest have evolved to better capture the underlying data (MARCHINI *et al.* 2007; PRICE *et al.* 2009). A significant finding in recent years has been the ubiquity of structural variation in the human genome (LI *et al.* 2011), with copy number variation regions spanning an estimated 12% of the human genome (REDON *et al.* 2006), yet most statistical genetics methods take no account of their presence. Here we tested the effect of one such type of structural variant, inversion polymorphisms, which are known to suppress recombination in individuals heterozygous for the variant (NAVARRO *et al.* 2000) and thus have a substantial influence on local genetic variation and therefore potentially on methods of inference that exploit such data.

We formally investigated the impact of inversions on three widely used statistical methods that use SNP data on samples of individuals to estimate (i) recombination rates, (ii) recent selection acting at loci, and (iii) population structure. We used simulated and real data to assess the effect of inversions on the three methods. We found that inversions can lead to systematic inaccuracy in estimating recombination rates using LDhat (MCVEAN *et al.* 2004) and in measuring population

structure using Principle Component Analysis, while the *iHS* statistic for detecting selective sweeps (VOIGHT *et al.* 2006) appears robust to the presence of inversions.

We illustrate that within large inversions (500 kb) LDhat underestimates the recombination rate, and by more than 80% for frequencies of the inverted sequence $\geq 50\%$ (Table 1). We expect the estimated recombination rate to be reduced due to the actual reduction in recombination in individuals heterozygous for the variant, but the level of underestimation that we observe is greater than expected. For instance, inversions with inverted sequence frequency 50% should reduce viable recombination by around half, yet LDhat estimates are 81.4% lower than the simulated rates for such inversions, and inversions with inverted sequence frequency 80% reduce viable recombination by an average 32%, yet LDhat estimates are an average 85.8% lower than the simulated rates. For inversions at low frequencies few individuals will have been heterozygous for the variant, thus little overall historical suppression of recombination, whereas those at high frequencies will have experienced greater suppression of recombination due to segregating for many generations with a high proportion of heterozygotes for inversion status. To check that the observed underestimation is not simply due to greater underestimation of recombination when the true recombination rate is lower than the genome-wide average, we assessed LDhat estimates at loci where the simulated rates were approximately half the genome-wide average. This should be a sufficient check because the maximum suppression of recombination at inversions is 50% (at inversion frequency 50%). We find that the average recombination rate is marginally overestimated at loci with simulated recombination rates between $5.75e-9$ and $6.75e-9$, and that there is a tendency for LDhat to overestimate recombination rates simulated with rates lower than the mean. Therefore the downward bias in estimates

that we observe at inversion loci cannot be explained by the behaviour of LDhat at loci with low recombination rates, and so our findings are conservative in this way.

While our real data analysis provides strong support for a downward bias in LDhat rates at inversion loci, it also indicates that LDhat estimates may be biased upward compared to contemporary rate estimates (pedigree-based rates here) at some inversion loci. Wegmann *et al.* (2011) found that many of the largest discrepancies between their admixture-based recombination rate estimates and LDhat-based estimates were at the sites of inversions. The authors found that LDhat estimates could be both substantially lower and substantially higher than the admixture-based estimates, which should reflect contemporary recombination rates, at inversion loci. However, the admixture-based estimates are based on a Hidden Markov Model (HMM), which may be affected by unusual patterns of LD and assumes that all sequences are in the same orientation, and so these estimates may also be biased at inversion loci. Furthermore, no formal testing was performed to establish whether the differences between the estimates were statistically significant. Considering the evidence from our simulation study and from real data, we conclude that we would expect to observe a downward bias in LDhat estimates at most inversion loci, but that there may be an upward bias in the estimates at some inversion loci. We consider the most likely reason for a downward bias in LDhat estimated recombination rates to be that particularly high LD is produced at inversion loci because recombination only occurs between homozygotes for inversion status (BOETTGER *et al.* 2012; O'REILLY *et al.* 2010), so that the high LD is not only due to the reduced recombination rate but also the stratification of recombination among the sequences. Regions of high LD are inferred as having low recombination rates by LDhat. If there is in fact an upward bias in LDhat estimates relative to contemporary

recombination rates at some inversion loci then there must be features of inversions, or of real genetic data, not fully captured by invertFREGENE as applied in our simulation study. The recombination model used by invertFREGENE generates data that correspond well to average levels of homozygosity and LD observed in real genetic data, but, for example, may not reflect the large ‘recombination deserts’ present in the human genome. LDhat tends to produce overestimates when very low recombination rates are simulated, so inversions in regions of particularly low recombination may have high LDhat estimates relative to contemporary rates. invertFREGENE assumes neutral evolution of sequences and so it may be that an upward bias in population-based recombination rates is likely under certain scenarios of selective pressures. Furthermore, invertFREGENE assumes panmictic mating, so may produce data that does not reflect highly stratified or admixed human populations well. Additional limitations of the invertFREGENE software, which may explain differences between simulated and real data results, are discussed below.

Our results on recombination suggest the need for further development of methods such as LDhat, or the re-evaluation of recombination rate estimates at known inversion loci, in order to account for the presence of inversions in the human genome. Further investigation is also required to establish whether there is an upward bias in LDhat at some inversion loci and, if so, what the causes are, which should form an intriguing area of research as our knowledge of genetic data and real inversions increase over time.

Our finding that recombination rate estimates are not systematically higher at inversion breakpoints, and thus do not falsely infer recombination hotspots, is reassuring in terms of hotspot inference using LDhat, but highlights a potential problem with recent methods developed to infer the location of inversion

polymorphisms from SNP data. An important assumption made by some methods for detecting inversions from SNP data (BANSAL *et al.* 2007; SINDI and RAPHAEL 2010) is that LD will be greatly reduced at inversion breakpoints because the separation of SNPs spanning the breakpoints can be large due to the inversion, despite being physically close according to their reference genome positions. However, inversions have been found to display strong LD between SNPs across their entire length (DENG *et al.* 2008; O'REILLY *et al.* 2010; PRICE *et al.* 2008), so the reduction in LD due to distance of separation may be far lower than expected at inversion loci, as indicated by our results, owing to the particular effect that inversions have on local LD via their impact on the recombination process. However, further testing is required to establish the effect of this on specific methods for detecting inversions.

Although there has been concern that inversions may disrupt methods for detecting selection (DENG *et al.* 2008; O'REILLY *et al.* 2008), we find that while the distribution of the iHS statistic for detecting selective sweeps is changed somewhat by the presence of inversions the statistic is sufficiently robust. Our real data analysis, in particular, suggests that any difference in haplotype breakdown between the inverted and non-inverted sequence caused by the suppression of recombination at inversion loci is insufficient to produce extreme values of the statistic that would be interpreted as recent positive selection.

While the disruptive effect of inversions on principle component analysis is evident from some empirical studies (DENG *et al.* 2008; JAKKULA *et al.* 2008), the problem has not been formally investigated. By applying PCA on both real (WTCCC 2007) and simulated genetic data, we find evidence that large inversion polymorphisms (> 500 kb) can dominate the leading principal components, even in a structured population. Our analyses of the real data showed that the data at known

inversions, on chromosomes 8p23 and 17p21, bias some of the top six PCs. In an unstructured simulated population the effect of a simulated 4 Mb inversion is even more pronounced and affects the first PC, while there was no noticeable effect of a 500 kb inversion. Using PCs as covariates has become a standard strategy for controlling for population structure in genome-wide association studies (GWAS). However, the correlation structure between SNPs across the genome is an inherent problem in using PCA to measure population structure, because each locus should contribute equally to the estimate of structure in a sample, yet loci with many SNPs in high LD will contribute more. We have found that the high LD between many SNPs at large inversions is sufficient to generate pronounced bias in the low order PCs. This suggests that it is critical to thin (or “prune”) SNP data according to LD over large physical distances prior to performing PCA in order to ensure its reliability. Such thinning should have little adverse impact on the results of PCA since $\sim 10^4$ SNPs are generally sufficient for the PCs to capture population structure effects (ASTLE and BALDING 2009).

The main limitation of this study is its reliance on one simulation software package, and thus specific model for inversion polymorphisms, in testing the population genetics methods on simulated data. invertFREGENE models inversions as having arisen from a single mutation event, after which neutral drift governs the frequency of the inverted sequence until a set frequency is reached. Real inversions may often be located at the sites of genomic instability and thus inversion loci may often be subject to several inversion, deletion or duplication mutations. They may also frequently reside within repetitive sequence, more prone to the 'errors' in recombination that can lead to inversions, where SNP coverage is typically reduced. It should also be noted that our simulated inversions are, on average, younger than

real inversions of the same frequency because our simulations stop as soon as the specified frequency is reached, although we do not expect this to change our general conclusions since we explored a range of frequencies and thus ages of inversions. Our model for inversions assumes discrete breakpoints, in that the recombination process outside and even adjacent to the inversion is entirely unaffected by the inversion. However, we would expect the conclusions drawn from our main simulation results to hold if recombination were suppressed beyond the breakpoints as well; we would expect the bias observed in LDhat recombination rate estimates to extend beyond the breakpoints, the iHS statistic to remain robust to inversions, and we would expect a slightly greater impact on PCA relative to inversions where suppression of recombination is only the length of the inversion. However, modelling discrete breakpoints could have made a critical difference to the results of our investigation into the hypothesis that recombination hotspots are falsely inferred at inversion breakpoints; discrete breakpoints should increase the probability of accepting this hypothesis, yet our results led to the rejection of the hypothesis and so we can conclude that 'soft breakpoints' would only further support our conclusion.

Although invertFREGENE models gene conversion events it does not model double crossover events. While some evidence for a double recombination event has recently been found within the inversion at the *MAPT* locus (STEINBERG *et al.* 2012), this event was predicted to span only 30 kb of the ~1.1 Mb inversion, and no other evidence for a double recombination event was found across the rest of the inversion after extensive investigation by two studies (BOETTGER *et al.* 2012; STEINBERG *et al.* 2012). To our knowledge, there is no other evidence for the occurrence of double crossovers within human inversions, and thus no other information from which to model their characteristics (such as size and rate of

occurrence); if they can potentially be large and occur frequently then the resulting pattern of variation data could be very different from considered here and so our predictions may not be accurate for such inversion loci. However, despite these reservations, there is no other software available for simulating inversions in population genetic data, and while the model used may not reflect the characteristics of all inversions, we believe that our results provide a good guide to the potential effects of inversions on these methods. The fact that our real data analyses were generally reflective of our simulation results in the three Results sections provides support for the reliability of the simulation model. However, further testing as more data on real inversions become available is required to assess how closely our simulation results compare to reality.

Another limitation of this study is the fact that we investigate only three population genetic methods, when there are many in use that do not account for the presence of inversions. Our findings on the effect of inversions on recombination rate estimates and principle component analysis suggest that large inversions can have a marked effect on local patterns of LD, and may therefore compromise any method influenced by LD that does not account for extreme deviations in LD. Methods for detecting selection that combine evidence from test statistics calculated independently at neighbouring SNPs (WEIR *et al.* 2005) and those that test 'runs of homozygosity' (ROH) (LENCZ *et al.* 2007) are examples of methods that may be biased by genomic regions with unusually strong and extended LD, such as inversions. Our focus has been on methods of inference commonly applied to human genetic variation data, and invertFREGENE aims to model human population genetic data; we suggest that researchers interested in the impact of inversions on inference

in other species should consider whether our conclusions are likely to generalise to their setting, and if not to use our study as a guide to investigations of their own.

Our study estimates the impact of inversion polymorphisms on three widely applied population genetic methods of inference, covering inference of recombination rates, selection intensity and population structure. The results highlight the importance of accounting for large-scale sources of variation in the genome, such as inversions, in population genetic methodology development and suggest the need to re-assess results from existing methods at known inversion loci.

ACKNOWLEDGEMENTS

Funding: Wellcome Trust (to N.-K.S.); ENGAGE consortium (grant P12892_DFHM to P.F.O.) and the National Institute for Health Research (NIHR) Imperial College Health Care NHS Trust Biomedical Research Centre (to P.F.O). Conflict of Interest: None declared.

REFERENCES

- AKEY, J. M., 2009 Constructing genomic maps of positive selection in humans: where do we go from here? *Genome Research* **19**: 711-722.
- ANTONACCI, F., J. M. KIDD, T. MARQUES-BONET, M. VENTURA, P. SISWARA *et al.*, 2009 Characterization of six human disease-associated inversion polymorphisms. *Human Molecular Genetics* **18**: 2555-2566.
- ASTLE, W., and D. J. BALDING, 2009 Population Structure and Cryptic Relatedness in Genetic Association Studies. *Statistical Science* **24**: 451-471.
- BANSAL, V., A. BASHIR and V. BAFNA, 2007 Evidence for large inversion polymorphisms in the human genome from HapMap data. *Genome Research* **17**: 219-230.
- BOETTGER, L. M., R. E. HANDSAKER, M. C. ZODY and S. A. MCCARROLL, 2012 Structural haplotypes and recent evolution of the human 17q21.31 region. *Nature Genetics* **44**: 881-885.
- DENG, L., Y. ZHANG, J. KANG, T. LIU, H. ZHAO *et al.*, 2008 An unusual haplotype structure on human chromosome 8p23 derived from the inversion polymorphism. *Human Mutation* **29**: 1209-1216.
- FRAZER, K. A., D. G. BALLINGER, D. R. COX, D. A. HINDS, L. L. STUVE *et al.*, 2007 A second generation human haplotype map of over 3.1 million SNPs. *Nature* **449**: 851-861.
- HOGGART, C. J., M. CHADEAU-HYAM, T. G. CLARK, R. LAMPARIELLO, J. C. WHITTAKER *et al.*, 2007 Sequence-level population simulations over large genomic regions. *Genetics* **177**: 1725-1731.
- HUDSON, R. R., 1983 Testing the constant-rate neutral allele model with protein sequence data. *Evolution* **37**: 203-217.
- JAKKULA, E., K. REHNSTRÖM, T. VARILO, O. P. H. PIETILÄINEN, T. PAUNIO *et al.*, 2008 The Genome-wide Patterns of Variation Expose Significant Substructure in a Founder Population. *The American Journal of Human Genetics* **83**: 787-794.
- JEFFREYS, A. J., L. KAUPPI and R. NEUMANN, 2001 Intensely punctate meiotic recombination in the class II region of the major histocompatibility complex. *Nature Genetics* **29**: 217-222.
- JOBLING, M., 2003 *Human Evolutionary Genetics: Origins, People and Disease*. {Garland Science/Taylor & Francis Group}.
- JOLLIFFE, I., 2002 *Principal component analysis*. Springer, New York.
- KELLEY, J. M., L. B. HUGHES, R. FENG, N. LIU, M. A. PADILLA *et al.*, 2008 Evaluating linkage disequilibrium and recombination provides a haplotype-tagging SNP panel of the major histocompatibility complex in African Americans. *Genes and Immunity* **9**: 271-273.
- KIDD, J. M., G. M. COOPER, W. F. DONAHUE, H. S. HAYDEN, N. SAMPAS *et al.*, 2008 Mapping and sequencing of structural variation from eight human genomes. *Nature* **453**: 56-64.
- KIMURA, M., 1985 *The Neutral Theory of Molecular Evolution*. Cambridge University Press.

- KONG, A., G. THORLEIFSSON, D. F. GUDBJARTSSON, G. MASSON, A. SIGURDSSON *et al.*, 2010 Fine-scale recombination rate differences between sexes, populations and individuals. *Nature* **467**: 1099-1103.
- KOOLEN, D. A., L. E. L. M. VISSERS, R. PFUNDT, N. DE LEEUW, S. J. L. KNIGHT *et al.*, 2006 A new chromosome 17q21.31 microdeletion syndrome associated with a common inversion polymorphism. *Nature Genetics* **38**: 999-1001.
- LAKICH, D., H. H. KAZAZIAN, S. E. ANTONARAKIS and J. GITSCHIER, 1993 Inversions disrupting the factor VIII gene are a common cause of severe haemophilia A. *Nat Genet* **5**: 236-241.
- LENCZ, T., C. LAMBERT, P. DEROSSE, K. E. BURDICK, T. V. MORGAN *et al.*, 2007 Runs of homozygosity reveal highly penetrant recessive loci in schizophrenia. *Proceedings of the National Academy of Sciences* **104**: 19942-19947.
- LI, Y., H. ZHENG, R. LUO, H. WU, H. ZHU *et al.*, 2011 Structural variation in two human genomes mapped at single-nucleotide resolution by whole genome de novo assembly. *Nat Biotech* **29**: 723-730.
- MARCHINI, J., B. HOWIE, S. MYERS, G. MCVEAN and P. DONNELLY, 2007 A new multipoint method for genome-wide association studies by imputation of genotypes. *Nat Genet* **39**: 906-913.
- MATISE, T. C., F. CHEN, W. CHEN, F. M. DE LA VEGA, M. HANSEN *et al.*, 2007 A second-generation combined linkage physical map of the human genome. *Genome Research* **17**: 1783-1786.
- MCVEAN, G. A. T., S. R. MYERS, S. HUNT, P. DELOUKAS, D. R. BENTLEY *et al.*, 2004 The fine-scale structure of recombination rate variation in the human genome. *Science (New York, N.Y.)* **304**: 581-584.
- MUNTÉ, A., J. ROZAS, M. AGUADÉ and C. SEGARRA, 2005 Chromosomal inversion polymorphism leads to extensive genetic structure: a multilocus survey in *Drosophila subobscura*. *Genetics* **169**: 1573-1581.
- MYERS, S., L. BOTTOLO, C. FREEMAN, G. MCVEAN and P. DONNELLY, 2005 A fine-scale map of recombination rates and hotspots across the human genome. *Science (New York, N.Y.)* **310**: 321-324.
- NAVARRO, A., A. BARBADILLA and A. RUIZ, 2000 Effect of inversion polymorphism on the neutral nucleotide variability of linked chromosomal regions in *Drosophila*. *Genetics* **155**: 685-698.
- NAVARRO, A., E. BETRÁN, A. BARBADILLA and A. RUIZ, 1997 Recombination and gene flux caused by gene conversion and crossing over in inversion heterokaryotypes. *Genetics* **146**: 695-709.
- NEWTON-CHEH, C., T. JOHNSON, V. GATEVA, M. D. TOBIN, M. BOCHUD *et al.*, 2009 Genome-wide association study identifies eight loci associated with blood pressure. *Nature Genetics* **41**: 666-676.
- NOVEMBRE, J., T. JOHNSON, K. BRYC, Z. KUTALIK, A. R. BOYKO *et al.*, 2008 Genes mirror geography within Europe. *Nature* **456**: 98-101.
- O'REILLY, P. F., E. BIRNEY and D. J. BALDING, 2008 Confounding between recombination and selection, and the Ped/Pop method for detecting selection. *Genome Research* **18**: 1304-1313.
- O'REILLY, P. F., L. J. M. COIN and C. J. HOGGART, 2010 invertFREGENE: software for simulating inversions in population genetic data. *Bioinformatics (Oxford, England)* **26**: 838-840.
- OSBORNE, L. R., M. LI, B. POBER, D. CHITAYAT, J. BODURTHA *et al.*, 2001 A 1.5 million-base pair inversion polymorphism in families with Williams-Beuren syndrome. *Nature Genetics* **29**: 321-325.
- PADHUKASAHASRAM, B., P. MARJORAM and M. NORDBORG, 2004 Estimating the Rate of Gene Conversion on Human Chromosome 21. *American Journal of Human Genetics* **75**: 386-397.
- PRICE, A. L., N. J. PATTERSON, R. M. PLENGE, M. E. WEINBLATT, N. A. SHADICK *et al.*, 2006 Principal components analysis corrects for stratification in genome-wide association studies. *Nat Genet* **38**: 904-909.
- PRICE, A. L., A. TANDON, N. PATTERSON, K. C. BARNES, N. RAFAELS *et al.*, 2009 Sensitive Detection of Chromosomal Segments of Distinct Ancestry in Admixed Populations. *PLoS Genet* **5**: e1000519-e1000519.

- PRICE, A. L., M. E. WEALE, N. PATTERSON, S. R. MYERS, A. C. NEED *et al.*, 2008 Long-range LD can confound genome scans in admixed populations. *American Journal of Human Genetics* **83**: 132-135; author reply 135-139-132-135; author reply 135-139.
- REDON, R., S. ISHIKAWA, K. R. FITCH, L. FEUK, G. H. PERRY *et al.*, 2006 Global variation in copy number in the human genome. *Nature* **444**: 444-454.
- SABETI, P. C., D. E. REICH, J. M. HIGGINS, H. Z. P. LEVINE, D. J. RICHTER *et al.*, 2002 Detecting recent positive selection in the human genome from haplotype structure. *Nature* **419**: 832-837.
- SABETI, P. C., P. VARILLY, B. FRY, J. LOHMUELLER, E. HOSTETTER *et al.*, 2007 Genome-wide detection and characterization of positive selection in human populations. *Nature* **449**: 913-918.
- SCHAEFFER, S. W., M. P. GOETTING-MINESKY, M. KOVACEVIC, J. R. PEOPLES, J. L. GRAYBILL *et al.*, 2003 Evolutionary genomics of inversions in *Drosophila pseudoobscura*: evidence for epistasis. *Proceedings of the National Academy of Sciences of the United States of America* **100**: 8319-8324.
- SCHAFFNER, S. F., C. FOO, S. GABRIEL, D. REICH, M. J. DALY *et al.*, 2005 Calibrating a coalescent simulation of human genome sequence variation. *Genome Research* **15**: 1576-1583.
- SINDI, S. S., and B. J. RAPHAEL, 2010 Identification and frequency estimation of inversion polymorphisms from haplotype data. *Journal of Computational Biology: A Journal of Computational Molecular Cell Biology* **17**: 517-531.
- SPENCER, C. C. A., and G. COOP, 2004 SelSim: a program to simulate population genetic data with natural selection and recombination. *Bioinformatics (Oxford, England)* **20**: 3673-3675.
- STEFANSSON, H., A. HELGASON, G. THORLEIFSSON, V. STEINTHORSDDOTTIR, G. MASSON *et al.*, 2005 A common inversion under selection in Europeans. *Nature Genetics* **37**: 129-137.
- STEINBERG, K. M., F. ANTONACCI, P. H. SUDMANT, J. M. KIDD, C. D. CAMPBELL *et al.*, 2012 Structural diversity and African origin of the 17q21.31 inversion polymorphism. *Nature Genetics* **44**: 872-880.
- STUMPF, M. P. H., and G. A. T. MCVEAN, 2003 Estimating recombination rates from population-genetic data. *Nature Reviews. Genetics* **4**: 959-968.
- TANTISIRA, K. G., R. LAZARUS, A. A. LITONJUA, B. KLANDERMAN and S. T. WEISS, 2008 Chromosome 17: association of a large inversion polymorphism with corticosteroid response in asthma. *Pharmacogenetics and Genomics* **18**: 733-737.
- VOIGHT, B. F., S. KUDARAVALLI, X. WEN and J. K. PRITCHARD, 2006 A map of recent positive selection in the human genome. *PLoS Biology* **4**: e72-e72.
- WEGMANN, D., D. E. KESSNER, K. R. VEERAMAH, R. A. MATHIAS, D. L. NICOLAE *et al.*, 2011 Recombination rates in admixed individuals identified by ancestry-based inference. *Nature Genetics* **43**: 847-853.
- WEIR, B. S., L. R. CARDON, A. D. ANDERSON, D. M. NIELSEN and W. G. HILL, 2005 Measures of human population structure show heterogeneity among genomic regions. *Genome Research* **15**: 1468-1476.
- WINCKLER, W., S. R. MYERS, D. J. RICHTER, R. C. ONOFRIO, G. J. McDONALD *et al.*, 2005 Comparison of fine-scale recombination rates in humans and chimpanzees. *Science (New York, N.Y.)* **308**: 107-111.
- WTCCC, 2007 Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature* **447**: 661-678.
- YU, H. X., J.-M. CHIA, G. BOURQUE, M. V. WONG, S. H. CHAN *et al.*, 2005 A population-based LD map of the human chromosome 6p. *Immunogenetics* **57**: 559-565.

Inversion frequency (%)	Mean % reduction in LDhat rec. rates	Expected % reduction in pedigree rec. rates	Expected mean ped/pop ratio of rec. rates
10	26.9	18	1.12
20	55.0	32	1.51
30	68.0	42	1.81
40	75.3	48	2.11
50	81.4	50	2.69
60	84.2	48	3.29
70	85.7	42	4.06
80	85.8	32	4.79
90	84.1	18	5.16

TABLE 1: Reduction in LDhat recombination rate estimates relative to the simulated rates and expected pedigree-based rates. The mean percentage reduction in recombination rates estimated by LDhat compared to the simulated rates between adjacent pairs of SNPs within inversion loci is shown across the range of inversion frequencies, for the simulation study described in the text and illustrated in Figure 1. The reductions in observed recombination rate estimates are contrasted with those that would be expected for inversions at these frequencies from pedigree-based recombination rate estimates, based on the expected proportion of heterozygotes for the inversion (assuming Hardy-Weinberg Equilibrium and no viable recombination events for heterozygotes for the inversion).

Inversion location (chr: start-end in Mb)	Inversion length (in Mb)	% of empirical dist. > max. ped/pop ratio	% of empirical dist. < min. ped/pop ratio	% of empirical dist. > max. abs(iHS)
8: 8.135-12.369	4.23	86.6	8.46	39.5
8: 126.776-129.803	3.03	87.2	19.9	71.5
15: 28.524-30.602	2.08	21.4	17.2	1.91
9: 85.702-87.641	1.94	52.1	74.9	90.2
17: 16.659-18.263	1.60	85.3	11.8	9.40
17: 31.888-33.393	1.51	23.9	86.3	81.5
15: 72.151-73.356	1.21	0.991	88.7	80.1
8: 6.913-8.103	1.19	1.16	100	100
16: 21.485-22.595	1.11	2.67	2.35	56.5
17: 40.900-41.989	1.09	0.847	44.6	86.8
13: 24.477-25.542	1.07	56.0	82.5	39.6
3: 12.128-12.876	0.75	60.6	40.0	38.5
4: 8.759-9.462	0.70	0.108	100	37.7
16: 32.061-32.743	0.68	73.6	2.98	58.0
16: 32.907-33.539	0.63	100	1.27	78.3
13: 99.738-100.312	0.57	18.6	86.6	92.5

TABLE 2. Recombination rate estimates and the iHS statistic at the 16 largest known human inversions (≥ 500 kb). Pedigree-based and population-based (LDhat) recombination rate estimates from human genetic data are first compared by taking the ratio of the rates in 100 kb sliding windows genome-wide, as described in the text. Bias in these rates at inversion loci is investigated by computing the proportion of the empirical distribution that is greater and smaller, respectively, than the maximum and minimum ‘ped/pop ratios’ observed within each of the 16 largest known human inversions (after omitting shorter inversions within, or overlapping, any of the 16, and any without markers between the breakpoints). The enrichment of values in the top 5% of the empirical maximum ped/pop distribution among the inversions has $P=0.000857$ (5 of 16), and in the top 1% has $P=0.000508$ (3 of 16), assuming a Binomial distribution, while the enrichment of values in the top 5% of the empirical minimum ped/pop distribution has $P=0.0429$ (3 of 16) and there are no minimum ped/pop values in the top 1% of the empirical distribution at the inversion

loci. Values of the *iHS* statistic were calculated genome-wide and investigated at the inversion loci similarly, with no enrichment of maximum absolute *iHS* values observed at the inversion loci ($P=0.560$, 1 of 16 in the top 5% of the empirical distribution).

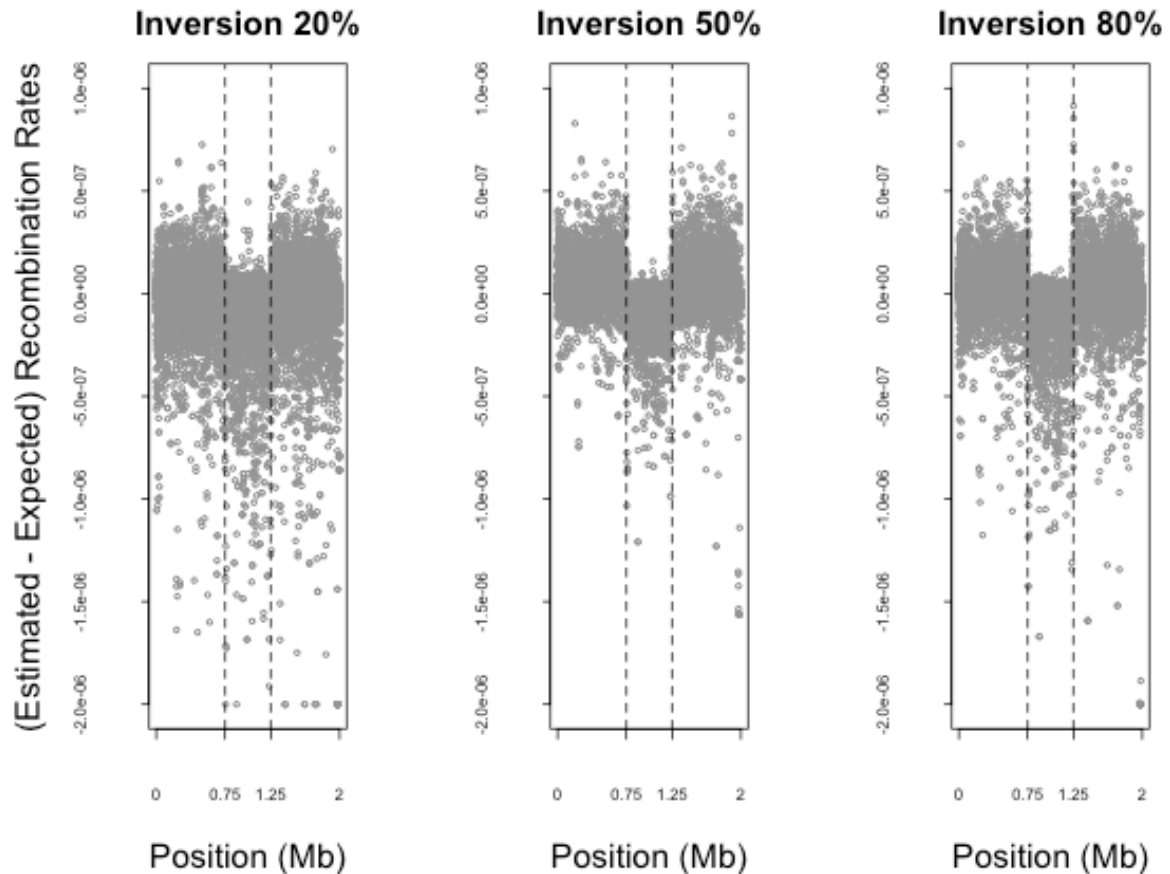


FIGURE 1. Difference between estimated and expected recombination rates for different frequencies of inversion. The difference between the recombination rates estimated from LDhat and those expected according to the simulated rates between each pair of SNPs and the expected reduction in those rates due to the proportion of heterozygotes for inversion status (see Table 1), described in the simulation study in the text are shown for the simulations relating to inversions at frequencies of 20%, 50% and 80%. Each point shows the difference (on the Y-axis) between the estimated and simulated recombination rates between a pair of adjacent SNPs, and (on the X-axis) the position in the 2 Mb simulated region of the first of the two adjacent SNPs. A small number (<15) of outliers are truncated at a difference in rates of $-2e^{-6}$. The inverted region is located between 0.75 Mb and 1.25 Mb (dashed vertical lines highlight inversion breakpoints). A downward bias in estimated, compared to expected, recombination rates is visible within the inversion region. For mean reductions in rates see Table 1.


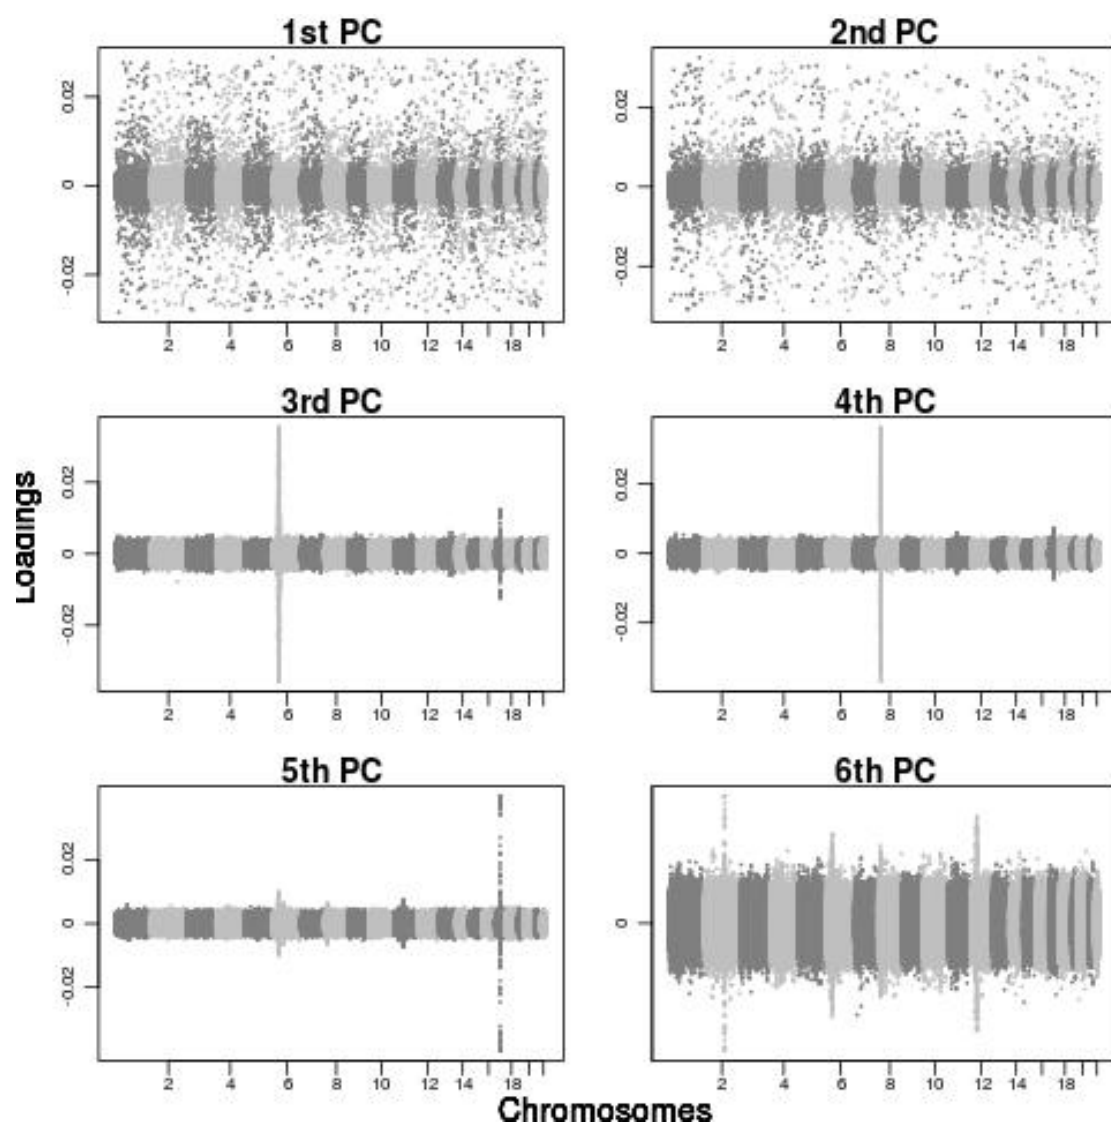


FIGURE 2. The probability density function of the standardised iHS statistic under different intensities of selection and at an inversion of frequency 60%. The probability density function (PDF) of the standardised iHS statistic computed from data from a 500 kb inversion with inverted frequency 60% under neutrality (dashed purple), a 4 Mb inversion with inverted frequency 60% under neutrality (dotted purple), data simulated under several intensities of positive selection with no inversion present, compared to the standard normal distribution (yellow), representing theoretical data under neutrality with no inversions, and invertFREGENE simulation data under neutrality with no inversions (light blue).

A



B




FIGURE 3. SNP loadings from Principle Component Analysis corresponding to the top 6 PCs, on WTCCC data with and without the inclusion of a simulated inversion. The loadings of the SNPs on the top six PCs across the 22 human autosomal chromosomes of the WTCCC data are shown. (A): No simulated inversion included in the data. (B): a simulated 4 Mb inversion at frequency 50% included in the data (highlighted in orange and at position 'A'), showing high localised variance in SNP loadings in PCs 3-6. See Figure S5 for more results at inversion frequencies of 20%, 50% and 80%.