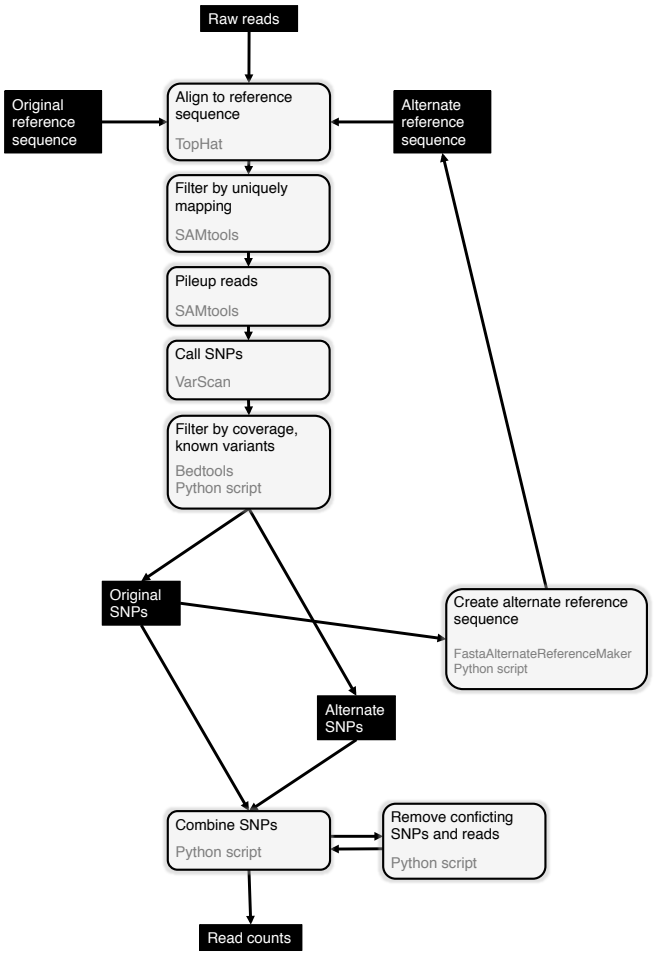

Fig. 1. Pipeline for generating allele-specific expression estimates. Raw reads are aligned to the original reference sequence using TopHat, retained if they are uniquely mapping, and piled up using SAMtools. SNPs are called using VarScan and filtered according to a fixed and variable coverage cutoff, as well as by a set of known SNPs with Bedtools intersect. These SNPs are used to create an alternate reference sequence. The previous steps are then repeated, this time with alignment to the alternate sequence. Finally, unique reads from both alignments are combined to obtain measures of ASE. If the data are phased, conflicting SNPs and reads are removed during this last stage.

Table 1. The number of sequence reads that are discarded from the analysis of ASE when SNPs are removed because they generate conflicts. A conflict occurs when two SNPs assign the same read to different parental genotypes, and the conflict threshold is the number of conflicting reads that are required before a SNP is discarded.

Conflict Threshold	Reads Discarded
1	98,775
2	51,837
3	38,308
4	21,907
5	21,367
6	19,944
7	18,712
8	18,274
9	17,827
10	17,617

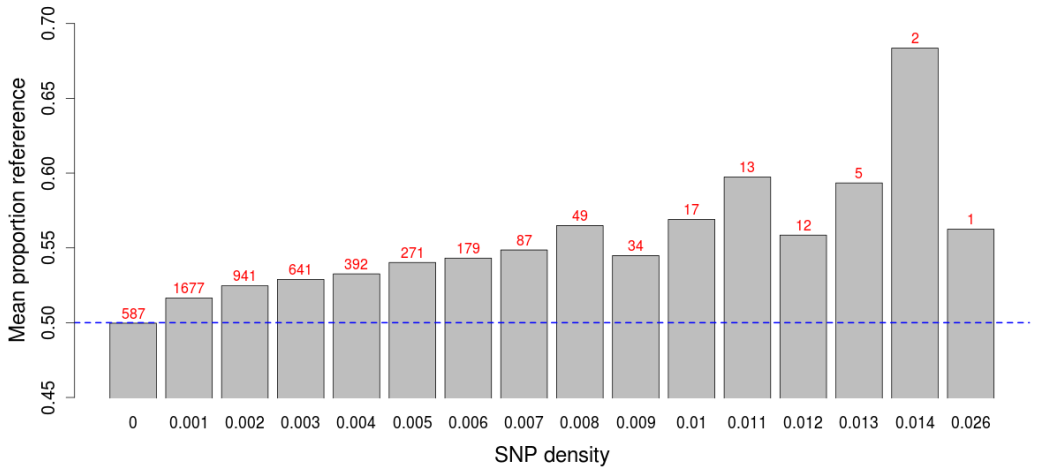

Fig. 2. The bias towards the reference allele is correlated with the density of SNPs (Spearman-Rank Correlation: P=2.2x10^{-16}). The reads were aligned to a single reference genome with two mismatches allowed between reads and the reference genome. SNP density was estimated as the number of SNPs observed per gene divided by the total coding gene length and was rounded to the nearest thousandth. Mean proportion reference was the average proportion reference for all genes in SNP density bin. The number of genes in each bin is shown in red.

Fig. 3. The proportion of the reference allele in SNPs when data were aligned with TopHat with a varied number of read-mismatches allowed (TopHat −N parameter). In all instances, indel length was set to two. The dashed horizontal line represents the expected average value. Boxes represent the interquartile range and whiskers extend to the most extreme data point that is no more than twice the interquartile range.

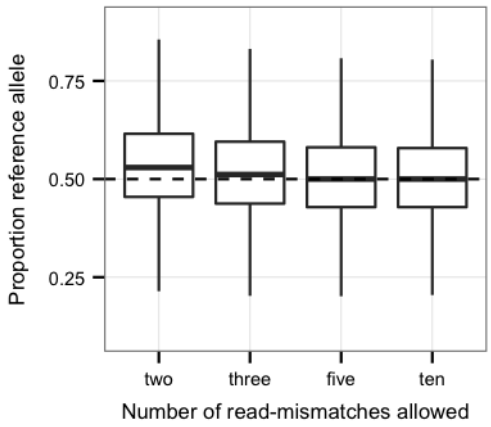

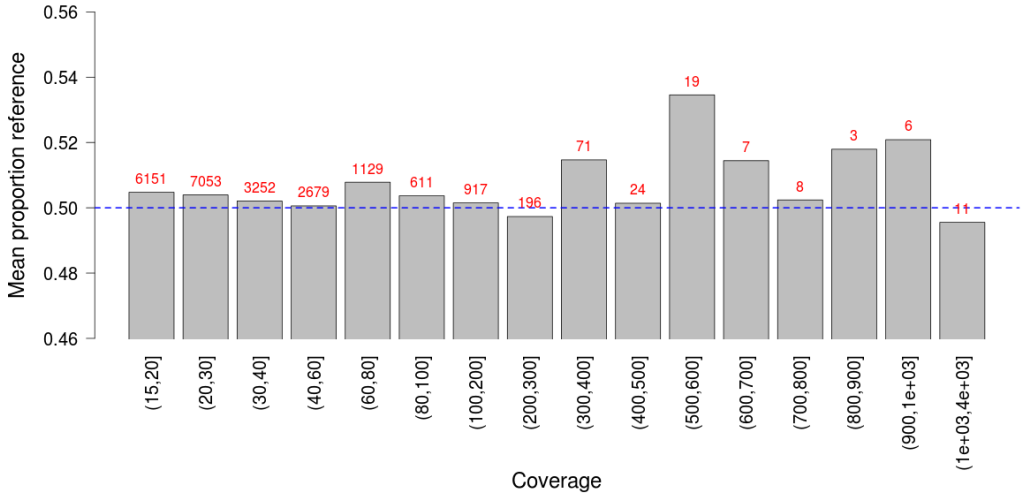

Fig. 4. There is no correlation between the mean depth of coverage and the bias towards reads carrying the reference allele when SNPs are called from the RNAseq data (Spearman's Rank Correlation: rho=0.15, $p = 0.68$). The data plotted is the same as in Figure 6B, main text.