

RESPONSE TO REVIEWERS FOR “ALLELE-SPECIFIC BINDING AND EXPRESSION: A UNIFORM SURVEY OVER THE 1000-GENOMES-PROJECT INDIVIDUALS”

RESPONSE LETTER

Reviewer #1

-- Ref1 – Endorsement for publication --

Reviewer Comment	This reviewer did not have formal comments to the authors as s/he found the revised paper to be satisfactory and endorses publication.
Author Response	We thank the reviewer for his/her thorough examination of our manuscript and endorsing our paper for publication.

Reviewer #2

-- Ref2.1 – General comment --

Reviewer Comment	The authors did not adequately address my two major concerns.
Author Response	We thank the reviewer for the thorough examination of our manuscript. We have provided additional analyses and responses.

-- Ref2.2 – mapping to the personal diploid genome --

Reviewer Comment	<p>My first comment was that mapping bias should be addressed. The authors replied by explaining that they excluded reads that map to more than one location. This is indeed a standard step in more alignment. Yet, the challenge when looking for ASE is not standard. Different alleles may have different mapping probabilities and this must be taken into account. Failing to do so results in a high number of falsely identified ASE.</p> <p>I must admit that it is a bit concerning to me that the authors interpreted my comment as a question regarding their standard alignment approach. In my mind, it points to a deep lack of familiarity with the ASE literature.</p>
Author Response	<p>We would like to thank the reviewer for pointing out the importance of allelic mapping bias, which includes the reference bias. In fact, reference bias has been widely regarded as the main source of allelic mapping bias, since the more standard alignment procedure is actually the alignment of reads to the human reference genome, not to the personal genomes [1,3,4,5]. Many publications have specifically cited the use of the personal genomes as a rigorous</p>

Deleted: point

Deleted: that the reference bias is not a separate issue from

Deleted: is the generic term to describe differential mapping probabilities of the alleles; the allelic mapping bias

Formatted: Font: Not Italic

Deleted: 2,

Deleted: 6

but computationally intensive procedure to correct for reference bias [1,3,4,5]. Thus, we are acutely aware of this primary issue in mapping bias, and have chosen to focus specifically on rectifying the reference bias by aligning reads to their corresponding diploid personal genome.

There is currently no single solution to totally eliminate allelic mapping bias [1]. Hence, while a small proportion of the mapping bias will still exist, we do expect the majority of the allelic bias to be accounted for, or at least alleviated, in the form of the reference bias by the use of the personal genomes.

Nonetheless, in this revision, we have further examined another bias within this small proportion of allelic mapping bias, which we termed the 'ambiguous mapping bias'. In the context of the personal genome, this can occur due to sequence homology in other regions (new Figure 1 in the manuscript), as described also by previous studies [1,5,6]. To date, the primary strategy to manage this bias has been via simulation of uniquely mapped reads and the identification and removal of sites in which >5% of the total number of reads exhibit such ambiguous mapping bias [1,5,7,8,9]. However, we found that site removal can be overly stringent, since many of these implicated sites are still detected as allele-specific under the beta-binomial test if we remove only the reads with ambiguous mapping bias (new Supplementary Table 5 in the manuscript). Hence, we adopted the ambiguous-read-removal strategy. Even though it is computationally more expensive (since we need to filter and re-process the original read pile), it provides the double advantage of being able to remove potential false positives and yet still able to retain those that are strongly allele-specific. Interestingly, while we were working on this submission, van de Geijn et al. published in Nature Methods a tool that also similarly removes reads, instead of sites, in order to account for allelic mapping bias [6].

SJS

So far, we have reprocessed all the datasets and analyses carefully twice, with each round taking approximately 3 months. We hope we have satisfied the reviewer by carefully implementing and accounting for not one, but two, main types of allelic mapping bias, in the context of the diploid personal genome. Additionally, our approach is already conservative, with multiple additional filters in place, such as quality control via the removal of highly over-dispersed datasets and using the beta-binomial test with an FDR of 5% for all datasets.

Deleted: [6]. A recent study by Panousis et al. found that the bias towards the reference allele contributes to the main bulk of the overall mapping bias in allele-specific expression [5].

Deleted: to

Deleted: from each individual's assay

Deleted: While

Deleted: do

Deleted: This

Deleted: situations where short reads that carry one allele may map perfectly to a reference genome but reads with the other allele map to multiple loci (due to

Deleted:) (


Deleted:)

Deleted: We termed this 'ambiguous mapping bias'.

Deleted: way

Moved (insertion) [1]

Deleted: [10]. There is currently no single 'solution' to perfectly eliminate all allelic mapping bias [1].



Deleted: Figure 1. Adapted from van de Geijn et al. showing allelic mapping bias in a personal genome due to sequence homology in other locations. Here, Read 1 uniquely maps to the haplotype 1, but Read 2 with the alternate allele maps to multiple locations in the other haplotype, and is therefore removed. ¶

Here, we investigated the effect of the ambiguous mapping bias on the detection of allele-specific SNVs, in the context of the diploid personal genome. ¶

Moved up [1]: Interestingly, while we were working on

Deleted: [6]. ¶

Deleted: implemented an 'ambiguous-read-removal'

Deleted: individual, we (1) first align the reads to the

Formatted: Font: Not Bold

Deleted:) We then map these simulated reads to the

Formatted: Font color: Red

Deleted: removing

Formatted: Font color: Red

Deleted: RNA-seq and 10% for ChIP-seq

Formatted: Font color: Red

Deleted: The personal genome is also able to handle

	<p><u>Finally, we have improved the manuscript by including results of the additional analyses for ambiguous mapping bias in the supplementary materials, a discussion in the 'Discussion' section and details of the new AlleleDB pipeline in the 'Results' and 'Methods' sections</u></p> <p>[1] Castel <i>et al.</i> (2015). <i>Genome Biol.</i>, 16(1):195 [2] Degner <i>et al.</i> (2009) <i>Bioinformatics</i>. 25(24) [3] Satya <i>et al.</i> (2012) <i>Nucleic Acids Res.</i> 40(16):e127 [4] Stevenson <i>et al.</i> (2013) <i>BMC Genomics</i>. 14:536 [5] Panousis <i>et al.</i> (2014). <i>Genome Biol.</i>, 15(9):467 [6] van de Geijn <i>et al.</i> (2015). <i>Nat Methods</i>, doi: 10.1038/nmeth.3582 [epub ahead of print] [7] Kilpinen <i>et al.</i> (2013). <i>Science</i>, 342(6159):744-7 [8] Lappalainen <i>et al.</i> (2013). <i>Nature</i>, 501(7468):506-11 [9] The GTEx Consortium (2015). <i>Science</i>, 348(6235):648-60 [10] Dixon <i>et al.</i> (2015). <i>Science</i>, 518(7539):331-6 [11] Rozowsky <i>et al.</i> (2011). <i>Mol Syst Biol.</i>, 7:522 [12] Sudmant <i>et al.</i> (2015). <i>Nature</i>, 526(7571):75:81</p>
Excerpt From Revised Manuscript	<p><u>Please refer to Supplementary Tables 1, 3 and 5 and their corresponding legends. Please also refer to the 'Results' section under 'AlleleDB Workflow' and 'Methods' section under 'Accounting for ambiguous mapping bias'.</u></p> <p><i>“(3) The third module filters reads that preferentially map to one allele over the other due to sequence homology (Figure 1), which we term ‘ambiguous mapping bias’. This bias occur when reads containing one allele maps to multiple locations and are thus removed, not because of worse alignment, but because of ambiguous alignment. For a uniquely mapped read that overlap at least one heterozygous SNV on one parental genome (‘original read’), we simulate reads that represent all possible haplotypes of that read, even though we found that most original reads overlap only 1 heterozygous SNV (typically >90%; Supplementary Table 3). We then align the simulated reads to the other parental genome. Original reads with simulated reads that map to multiple locations or do not map back to the same location on the other parental genome are removed. (Figure 1). We subsequently re-align the filtered read pile to the diploid personal genome (see ‘Methods’).”</i></p> <p><u>Please refer to the ‘Discussion’ section for more description.</u></p> <p><i>“The second allelic mapping bias stems from loci with sequence homology. We term this ‘ambiguous mapping bias’, because reads from one allele might align ambiguously to multiple locations, resulting in reads with the other allele being unduly favored (Figure 1).^{19,38,35} Several strategies have been implemented in dealing with the ambiguous mapping bias (Supplementary Table 1). ... We also show that ambiguous mapping bias seems to have a greater effect on ChIP-seq than RNA-seq datasets, even after accounting for reference bias by the personal genomes (Supplementary Table 5). Besides allelic differences, ambiguous mapping is also highly dependent on the length of the read, as also shown by Degner <i>et al.</i> that the bias decreases with increasing</i></p>

Deleted: ¶
We have included new sections in the 'Results', 'Discussion' and 'Methods' section about our new module on allelic mapping bias.

Formatted: Normal, Left

Formatted: Font: Arial, Font color: Auto

read length.¹⁹ We envision that ambiguous mapping bias will be further alleviated by long read technologies being employed in functional assays.”

-- Ref2.3 – Over-dispersion –

Reviewer Comment	<p>My second major concern was regarding the binomial test to identify ASE. The authors begin their response by citing other papers that used such a test. I am not sure what the argument presented here, especially since the authors proceed by acknowledging over-dispersion in their data. So, yes, other paper got it wrong in the past, but this is hardly a reason to perpetuate this mistake.</p> <p>As for their revised approach, estimating a global over-dispersion parameter is not effective. Removing some loci because of 'too much' over-dispersion is ad hoc and was not justified. But more importantly, there are at least 3 published methods now to identify ASE using models that estimate site-specific over-dispersion, account for mapping bias, and report p values based on permutation. Why not use one of those published methods?</p>
Author Response	<p>While we thank the reviewer for his/her comment, we want to clarify that the purpose of the references is not to make any claims on the 'correctness' of the methods, but to point to the broader reality that there is currently a diversity of methods in the field, where there is no firm consensus on the 'right' approach. The fact that these publications are recent and peer-reviewed at influential journals indicates the plurality of the methods accepted by the community, each with their own advantages and limitations. For example, van de Geijn <i>et al.</i> [1] is a very recent publication in <i>Nature Methods</i> that presented a software, which performs alignment to the human reference genome, accounts for mapping bias and uses the beta-binomial test to account for an individual-specific (not site-specific) global over-dispersion. However, it is not able to take into account indels and larger structural variants, which can be accommodated by the construction of personal genomes. Moreover, the estimation of a global over-dispersion has also been employed extensively in many recent and peer-reviewed software that detect allele-specific expression [1-5].</p> <p><u>Additionally, our</u> revised approach estimates over-dispersion at two levels. An over-dispersion <u>parameter</u> is estimated for each dataset to remove <i>entire datasets</i> (not loci) that are deemed too over-dispersed and that might result in higher number of false positives. After which, for each sample (for RNA-seq and each sample and transcription factor, TF, for ChIP-seq experiments), we pool the datasets and estimate the individual-specific global over-dispersion (for each sample for RNA-seq and also each sample and transcription factor for ChIP-seq) and apply this estimation to</p>

Deleted: In particular, we have utilized our approach in the 1000 Genomes Structural Variant group, whose manuscript has recently been peer-reviewed and published by *Nature*.

Deleted: Our

Excerpt From Revised Manuscript	<p>the beta-binomial test for each site in that individual (or TF). Hence, in this manner, the estimation of the over-dispersion can accommodate user-defined site-specific estimation of over-dispersion if necessary. Our R code is provided on our website for modifications and more customized analyses by the user.</p> <p>We further point out that our two-step serial procedure is novel. <u>By removing datasets that are too over-dispersed at the outset, this first step serves as a quality control to homogenize the pooling of datasets before the second overdispersion calculation.</u> This fits very well into our pipeline as it facilitates the harmonization and uniform processing of large amounts of data and alleviates an ascertainment bias in which more positives might stem from these highly over-dispersed datasets if they are not removed.</p> <p>Hence, we have retained our estimation and use of a global over-dispersion for detecting allele-specific variants.</p> <p>[1] van de Geijn <i>et al.</i> (2015). <i>Nat Methods</i>, doi: 10.1038/nmeth.3582 [epub ahead of print] [2] Sun (2012). <i>Biometrics</i>. 68(1):1-11 [3] Mayba <i>et al.</i> (2014). <i>Genome Biology</i>. 15(8):405 [4] Crowley <i>et al.</i> (2015). <i>Nature Genetics</i>. 47(4):353-60 [5] Harvey <i>et al.</i> (2015). <i>Bioinformatics</i>. 31(8):1235-42</p>
---------------------------------	--

Deleted: and is introduced to homogenize the pooling of datasets, by

Reviewer #3

-- Ref3.1 – Endorsement for publication --

Reviewer Comment	The manuscript is much improved and the authors have sufficiently addressed the majority of my concerns. I have the following minor comments:
Author Response	We thank the reviewer for the thorough examination of the manuscript and we are pleased that the reviewer finds our improved manuscript satisfactory.

-- Ref3.2 – Include additional references --

Reviewer Comment	<p>1) Imprinting discussion should reference recent imprinting paper from GTEX. Lappalainen in Genome Research.</p> <p>2) Heritability analyses of ASE should reference Li, AJHG, 2014.</p>
------------------	---

Author Response	We have included the references in the respective sections of the manuscript.
Excerpt From Revised Manuscript	<p>Please refer to the 'Discussion' section and also the 'Results' section under "ASB and ASE Inheritance analyses using CEU trio".</p> <p>"It could also be a result of other epigenetic effects such as genomic imprinting where no variants are causal⁴¹", where reference⁴¹ is by the GTEx consortium and Baran <i>et al.</i> published in <i>Genome Research</i>.</p> <p>"The CEU trio is a well-studied family and with multiple ChIP-seq studies performed on different TFs. Previous studies have also presented allele-specific inheritance.^{10,15,21}", where reference 21 is by Li <i>et al.</i> published in <i>American Journal of Human Genetics</i>.</p>

Deleted: 35

Deleted: 35