

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

RESPONSE LETTER

Overall comment

We want to thank the reviewers for recognizing the importance of our study and offering insightful comments. We have significantly revised the manuscript to address their comments. In particular, we have performed a major overhaul of the pipeline and corresponding analyses. Per reviewers #2 and #3's suggestions to include the newer advances in allele-specific variant detection, we now use a beta-binomial test to account for the overdispersion properties of RNA-seq and ChIP-seq datasets in order to call ASE and ASB variants. Additionally, we implemented a first-pass filter in which we compute the overdispersion parameter for each of the 1,280 ChIP-seq and RNA-seq datasets prior to the pipeline to identify and exclude those that are highly overdispersed. Per reviewer #1's suggestion, we have uniformly re-called the peaks for 83 ChIP-seq datasets with a common peak caller. The peak calling was performed using the personal genomes for each of the 14 individuals with matching ChIP-seq data. To further build AlleleDB as a resource, we developed novel formalisms to call allele-specific genes and genomic elements. We also included new analyses and figures to illustrate the advantages of having results from allele-specific analyses obtained from a large number of genomes.

The specific reviewers' comments are further addressed below.

Reviewer #1

-- Ref1.1 – General positive comment --

Reviewer Comment	This manuscript by Chen, et al. entitled "Allele-specific binding and expression: a uniform survey over many individuals and assays". This study is an exploration of the effects of genomic variation on expression of one of alleles and on transcription factor binding using previously generated RNAseq and ChIPseq datasets. The importance of this topic is timely and potentially significant. The manuscript is written in a relatively clear manner. While there is much to recommend this manuscript, several areas and questions need to be addressed to assist the reader to better understand or accept the findings. The major issues include:
Author Response	We thank the reviewer for acknowledging the importance and timeliness of our study and for his/her thorough examination of our

	manuscript. We have made significant changes to specifically address the technical and analytical aspects of our manuscript for the reader.
--	---

-- Ref1.2 – DNA- versus RNA-aligners --

Reviewer Comment	It's not clear why Bowtie1 (DNA aligner) was used for aligning RNA-seq reads. This significantly reduces the number of mappable reads, and also can miss the allele-specific splicing. The use of a RNA (spliced) aligner might better be used for RNA-seq mapping.																													
Author Response	<p>One of the main thrusts of our study is uniformity. Since a RNA aligner is not typically used in ChIP-seq alignment, we chose a consistent aligner in order to uniformly process both ChIP-seq and RNA-seq datasets. That said, while we do agree with the reviewer that using a RNA (spliced) aligner improves the number of mappable reads, we have performed some additional analyses and found that the resultant AS calls between using a DNA and RNA aligner are very comparable (figure shown below). In panel a, we compare the number of mapped reads between using Bowtie1 (DNA aligner) and STAR (RNA aligner) at various coverage, sub-sampled from a RNA-seq dataset from NA12878 (the final data point is the entire dataset). On average, the amount of increment in mappability by STAR (with respect to Bowtie1) is < 16%. In panel b, we then detected ASE sites based on the entire dataset using Bowtie1 and STAR and found that the number of ASE calls between the two aligners are very similar (>90%). Thus, despite the increase in mappable reads, the effect of using a RNA aligner is relatively modest.</p> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>a</p> <table border="1"> <caption>Data for Panel a: Aligned reads vs Input reads</caption> <thead> <tr> <th>Number of input reads (x10⁶)</th> <th>Bowtie1 Aligned reads (x10⁶)</th> <th>STAR Aligned reads (x10⁶)</th> </tr> </thead> <tbody> <tr><td>0</td><td>0</td><td>0</td></tr> <tr><td>100</td><td>~0.7</td><td>~0.8</td></tr> <tr><td>200</td><td>~1.4</td><td>~1.6</td></tr> <tr><td>300</td><td>~2.1</td><td>~2.4</td></tr> <tr><td>400</td><td>~2.8</td><td>~3.2</td></tr> <tr><td>500</td><td>~3.5</td><td>~4.0</td></tr> </tbody> </table> </div> <div style="text-align: center;"> <p>b</p> <table border="1"> <caption>Data for Panel b: ASE SNVs overlap</caption> <thead> <tr> <th>Category</th> <th>Count</th> </tr> </thead> <tbody> <tr><td>Bowtie1 only</td><td>677</td></tr> <tr><td>Intersection</td><td>8234</td></tr> <tr><td>STAR only</td><td>296</td></tr> </tbody> </table> <p>Number of ASE SNVs that are called in Bowtie and STAR for the same set of reads (NA12878 RNA-seq data)</p> </div> </div>	Number of input reads (x10 ⁶)	Bowtie1 Aligned reads (x10 ⁶)	STAR Aligned reads (x10 ⁶)	0	0	0	100	~0.7	~0.8	200	~1.4	~1.6	300	~2.1	~2.4	400	~2.8	~3.2	500	~3.5	~4.0	Category	Count	Bowtie1 only	677	Intersection	8234	STAR only	296
Number of input reads (x10 ⁶)	Bowtie1 Aligned reads (x10 ⁶)	STAR Aligned reads (x10 ⁶)																												
0	0	0																												
100	~0.7	~0.8																												
200	~1.4	~1.6																												
300	~2.1	~2.4																												
400	~2.8	~3.2																												
500	~3.5	~4.0																												
Category	Count																													
Bowtie1 only	677																													
Intersection	8234																													
STAR only	296																													

-- Ref1.3 – ASE of the genes --

Reviewer Comment	While calculating ASE for each of the SNV is straightforward, it's not clear how the ASEs of the genes
------------------	--

	are calculated. This would require combining ASE from multiple SNVs (and isoforms) of the same gene
Author Response	We agree with the reviewer and indeed, the use of multiple SNVs of the same gene is the central premise of our enrichment analyses in determining the amount of allele-specific behavior for a gene. Perhaps our description in the method section was not sufficiently clear; we have clarified this in the ‘Methods’ section of the revised manuscript.
Excerpt From Revised Manuscript	Please refer to the ‘Methods’ section, under ‘Enrichment Analyses’. <i>“Enrichment analyses were performed in two ways: ‘collapsed’ and ‘expanded’ (Figure 4b). In both cases, we aggregate ASB and ASE SNVs within a specific genomic element, such as a gene or an enhancer. We then use the Fisher’s exact test to calculate the odds ratio and the hypergeometric p value, to test for the enrichment of allele-specific SNVs compared to ‘control’ SNVs, which are non-allele-specific ‘accessible’ SNVs.”</i>

-- Ref1.4 – Uniform peak calling --

Reviewer Comment	AS SNVs are counted if they fall within the called ChIPseq peaks. However, the peaks used were the ones used for each of the datasets studied with the exception of the McVicker's set. The lack of a uniform peak called for the calling of peaks will lead to significant variability due to the disparity in the results derived from various algorithms (i.e. some peaks cover more for the genome). In turn this can potentially inflate or diminish the number of sites evaluated. At least some evidence should be presented that the use of various peak callers will not significantly alter the number of variant with allele-specific phenotypes.
Author Response	We agree with the reviewer that uniformity in peak calling is vital and have taken major steps in the revision to alleviate the concern. We addressed this by re-aligning reads from all the 276 ChIP-seq datasets. This is additionally performed in the context of each of the 14 pairs of personal haplotypes of each dataset, re-calling the peaks for each haplotype using a common peak-caller, PeakSeq, and then re-combining the peaks per dataset and for each personal diploid genome (haplotype pair). We have also made the uniformly-processed peaks available as a resource on the website.
Excerpt From Revised Manuscript	Please refer to the ‘Methods’ section, under ‘Allele-specific SNV detection’. <i>“Peak regions are determined by first performing PeakSeq⁶⁰ for each of the personal haploid genome. Only a single read per strand per position is kept and duplicates removed. The fragment length is set to 200 bps. Peak calling is performed with default parameters and the final peak set for each transcription factor is identified at a false discovery rate of 5%. Finally, the coordinates of the peaks (based on the respective personal haploid genomes) are mapped to the reference genome and then finally being merged between the haploid genomes. We also make the uniformly-processed peaks available as a resource on the AlleleDB website.”</i>

-- Ref1.5 – The advantages of many genomes --

Reviewer Comment	<p>While the study has as one of its strengths the development of a pipeline that can handle many (380) genomes with low coverage, it is unclear what the biological insights on this tour de force are other than the identification of 144K and 169K unvalidated ASEs and ASBs, respectively.</p>
Author Response	<p>The reviewer's criticism is that he/she believes that the sole utility of detecting allelic events in many genomes is merely the identification of large numbers of ASE and ASB SNVs. We contend that it is precisely because of the identification of large numbers of ASE and ASB SNVs using multiple genomes, when appropriately processed, that more biological insights and uses can be developed.</p> <p>Our downstream analyses provide a window into some of these possibilities when many genomes are available. For instance, the enrichment analyses will not be feasible without a large number of ASE and ASB SNVs. It is important to appreciate that many SNVs are rare, thus the abundance and detection of rare allele-specific variants increase with many genomes. Previous studies mostly focus on a very small number of genomes. Hence, it is difficult to perform allele-specific analyses on rare variants from a single study. Yet, having large number of rare variants is important, especially when we want to define allele-specificity across a genomic region, as this requires combining information from multiple SNVs across the region, as already alluded to in the second comment of the same reviewer. In this case, further annotation and biological insights can then be provided for regions that seem to be more attuned to allele-specific behavior. In particular, these allelic regions form a useful genome annotation. Additionally, the aggregation of many genomes enables a more confident identification of common SNVs that have corroborating allele-specific evidence across multiple individuals, which in itself can serve as both a validation and a biological observation. Our study shows the value of combining many genomes across multiple studies in providing biological insights.</p> <p>In the revised manuscript, we have included a discussion of how having many genomes, when appropriately processed, can be useful, for instance, in enrichment analyses that aggregate rare variants in gene- or element-centric analyses. We have also included a new analysis to quantify the effects of having common allele-specific variants across multiple individuals. We now have two new figures: Figure 4, to illustrate the advantage of visualizing</p>

	and having many genomes in validating common variants, and Figure 5, to capitalize on common variants in performing a population-aware enrichment analysis.
Excerpt From Revised Manuscript	<p>Please refer to Figures 4 and 5 and their corresponding legends. Please also refer to the ‘Results’ section under ‘AS variants and enrichment analyses’. Please refer to the ‘Discussion’ section for more description.</p> <p><i>“When SNURF is queried in our database, we can see clearly that the allele-specificity is supported not only by evidence from 61 ASE loci across the gene but a number of variants are shown to be also allele-specific over multiple individuals, one variant even up to 169 individuals.”</i></p> <p><i>“However, such element-centric analyses will not be feasible without a large number of ASE and ASB SNVs. It is important to appreciate that a significant portion of SNVs are rare, thus the abundance of and detection of rare AS variants increase with many genomes. Previous studies mostly focus on a very small number of genomes. Hence, it is difficult to perform AS analyses on rare variants from a single study. Yet, having a large number of rare variants is important, especially in quantifying allelic activity in elements, as this requires aggregating information from multiple SNVs across the a genomic region.”</i></p>

-- Ref1.6 – Definitions for acronyms --

Reviewer Comment	Minor point: No definition for CEU (Northern Europeans from Utah) RPB2, PAX5, etc.
Author Response	All the definitions of the various human populations used for the 1000 Genomes Project were, in fact, already included in the original Methods section of the manuscript. They are intentionally omitted from the main text to enable readability. In the current revision, we have added a sentence in the main text, referring the reader to the Methods section. We have also included the full names of the transcription factors (or short descriptions if no full names) such as RPB2 and PAX5, in the main text, at instances where we first mentioned them.
Excerpt From Revised Manuscript	<p>Please refer to the main text for all definitions, but specifically the ‘Methods’ section for the population definitions, and for the following excerpts, please refer to the ‘Results’ under ‘Rare variants and purifying selection in AS SNVs’ and ‘AS variants and enrichment analyses’ respectively:</p> <p><i>“... about two folds higher in the YRI than the other European sub-populations of comparable (CEU, FIN) or larger (TSI) population sizes (see Methods for full explanation of population abbreviations)...”</i></p> <p><i>“Our visualization shows ASB loci from POL2 (RNA polymerase II largest subunit), RPB2 (RNA polymerase II second largest subunit) and MYC (also c-Myc, or v-myc avian myelocytomatosis viral oncogene homolog)...”</i></p>

Reviewer #2

-- Ref2.1 – General comment --

Reviewer Comment	This is an exceptionally naïve analysis of ASE and ASB patterns. The analysis to identify the ASE/ASB patterns is flawed, the statistical modeling is too basic, and the enrichment analysis is crude.
Author Response	We thank the reviewer for the thorough examination of our manuscript. AlleleDB is, in fact, intended as a resource for ASB and ASE. Nonetheless, we have taken into account advances in the field and implemented more sophisticated changes to the statistical underpinnings of our pipeline. We have also included more rigorous analyses in our revised manuscript.

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

Reviewer Comment	<p>I have two concerns that, in my mind, are fatal flaws of the current analysis:</p> <p>First, mapping to a personal diploid genome indeed reduces the reference bias, but it does not eliminate the error associated with differences in mappability between the two alleles. In other words, the bias is gone, but the inflated variance due to mappability issues still persists. The only solution to date has been to map each allele separately and only retain reads that map uniquely at each allele, before the counting is done. This is a crucial aspect of the analysis presented in this paper and it must be addressed.</p>
Author Response	We agree with the reviewer that, in addition to building a personal diploid genome, mapping only unique reads to the individual haplotype or allele before the counting process is important. Our approach does encompass this and we have emphasized this point in the revised manuscript to better reflect this.
Excerpt From Revised Manuscript	<p>Please refer to the ‘Methods’ section, under ‘Read alignment and estimation of ρ’.</p> <p><i>“Reads are aligned against each of the derived haploid genome (maternal/paternal genome for trio) using Bowtie 1.⁵⁶ When a read is aligned to the same locus, we only pick the alignment that map better to a haplotype. Otherwise, if a read is tied in alignment to both haplotypes, we discard it. No multi-mapping is allowed and only a maximum of 2 mismatches per alignment is permitted.”</i></p>

-- Ref2.3 – Over-dispersion --

Reviewer Comment	Second, the ASE analysis was performed using a simple binomial test. This leads to a large number of falsely identified ASE patterns because of over dispersion in the data. Over dispersion in both RNA-seq and ChIP-seq data sets has been documented and commented on in a large number of papers. The correct analysis must use some
------------------	--

	<p>strategy to estimate the over dispersion parameter and take it into account when testing for ASE.</p>
<p>Author Response</p>	<p>While we thank the reviewer for his/her suggestion, we also note that many very recent publications have also used a binomial test in their detection of ASE and ASB SNVs. We list some of them here:</p> <p>Lappalainen, T. et al. (2013). Nature. 501(7468):506-11 Kilpinen H. et al. (2013). Science. 342(6159):744-7 Ding, Z. et al. (2014). PLoS Genet. 10(11):e1004798 Dixon, JR. et al. (2015). Nature. 518(7539):331-6 The GTEx Consortium. (2015). Science. 348(6235):648-60</p> <p>Nonetheless, we agree with and have taken to heart the reviewer’s comment in order to provide a repository with ‘cleaner’ sets of ASE and ASB SNVs. As a result, we have significantly revamped our pipeline in terms of its statistical underpinnings and also re-processed all the 1,280 ChIP-seq and RNA-seq datasets. We now use a beta-binomial distribution to estimate the overdispersion behavior of each dataset and then use this as a filtering step to exclude those that exhibit a greater overdispersion estimated from the allelic ratio distribution. As pointed out by the reviewer, overdispersion leads to a large number of false positives. This first step thus acts as a first pass in identifying datasets that are too overdispersed to start off with. Subsequently, we then use a beta-binomial test to estimate and account for overdispersion during our ASE and ASB detection.</p> <p>The new Figure 1, the Results and Methods sections give descriptions of our revised pipeline.</p>
<p>Excerpt From Revised Manuscript</p>	<p>Please refer to the new Figure 1 and its corresponding legend and the ‘Results’ section under ‘AlleleDB Workflow’ for a summary of the revised pipeline. Please also refer to the ‘Methods’ section for more detailed explanations, under ‘Read alignment and estimation of ρ’ and ‘Allele-specific SNV detection’.</p> <p><i>“...In the first round, the alignment is performed for each of 276 ChIP-seq and 987 RNA-seq datasets to calculate a measure of overdispersion (with respect to an expected binomial distribution), ρ...”</i></p> <p><i>“To estimate ρ, we adopt a three-step approach. We first obtain the empirical histogram for the allelic ratios of all heterozygous SNVs with read counts ≥ 6. Next, we calculate the expected null distribution (where there is no allelic imbalance) using the probability density function (pdf) of the beta-binomial distribution using the R package, VGAM⁵⁹:</i></p> $P_{\text{betabin}}(X = k n, a, b) = \binom{n}{k} \frac{B(k + a, n - k + b)}{B(a, b)}$

	<p>where n represents the total number of reads at a particular locus, $B(x,y)$ represents the beta function with variables x and y, a and b represent the shape parameters of the beta distribution. For computational efficiency, if $n \geq 1000$, we set it to a maximum of 1000, but retain the allelic ratio at the SNV. The VGAM beta-binomial routines require the input of the overdispersion parameter, ρ, and probability of success (also the mean of the beta distribution), which we fix at $p=0.5$ since the null hypothesis assumes no allelic imbalance. We then obtain the expected beta-binomial distributions for $\rho=0$ to $\rho=1$ with increment of 0.1, and choose ρ that minimizes the least sum of squared errors (LSSE) between the empirical and the expected distributions. Lastly, to further refine our estimate, we iterate a bisection method to arrive at a LSSE (R pseudo-code available in Supp file 5).”</p> <p>“...Here, a beta-binomial p-value is derived based on the VGAM R package as described in the previous section...”</p>
--	--

-- Ref2.4 – Control SNVs --

Reviewer Comment	<p>I have a few other major concerns:</p> <p>It is not entirely clear to me how the 'control SNVs' were defined. Are these simply cases where ASE was not detected? This seems a bit naïve to me; is the probability of including as a control a case where the null is rejected with a marginal p value is the same as a case where the null is rejected at, for example, $P > 0.8$? Also, I don't understand what it means to match the controls to the test cases by 'accessibility for statistical significance'. The terminology used is strange to me; is this a complicated way to say that you matched the power? If so, how was it done? If a cutoff for power was used, this would not result in true matching because the controls would probably be biased towards the lower threshold. More details on this analysis are needed.</p>
Author Response	<p>We thank the reviewer for pointing out our definition of 'control SNVs'. Perhaps we were not sufficiently clear in our descriptions</p> <p>Our 'control SNVs' are not simply the complement of the allele-specific set nor are they defined by a single cutoff for power. We describe our control SNVs as “non-allele-specific accessible SNVs”. Each set of accessible SNVs is defined by a read-depth cutoff. This cutoff is the minimum read depth that a heterozygous variant needs to possess to be able to even reach the significant p value threshold (accessible). Since the significant p value threshold is dataset-specific, there is a minimum read depth cutoff for each dataset such that each set of control SNVs generated is matched, in terms of the minimum number of reads, to each individual dataset's AS SNVs. In addition, our significant p values are also corrected using FDR estimated from an explicit computational simulation, which takes into account the number of reads at each heterozygous SNV when correcting the p values.</p>

	<p>This would also attenuate the bias towards having SNVs with lower read depth.</p> <p>We have provided more detailed explanations in both the main text and the Methods section to further clarify the definition.</p>
Excerpt From Revised Manuscript	<p>Please refer to the ‘Methods’ section, under ‘Enrichment Analyses’</p> <p><i>“We define the set of accessible SNVs as all heterozygous SNVs that exceed the minimum number of reads required in order for SNVs to be significantly detectable by the beta-binomial test for each dataset; this includes both allele-specific and non-allele-specific SNVs. This is an additional, more stringent criterion imposed beyond the minimum threshold of 6 reads. Given a fixed FDR cutoff, for a larger dataset, the beta-binomial p-value threshold is typically lower, making the minimum number of reads (N) that will produce the corresponding p-value, larger. This alleviates a bias in the enrichment test for including SNVs that do not have sufficient reads in the first place. Considering an extreme allelic imbalance case where all the reads are found on one allele (all successes or all failures, i.e. allelic ratio is 0 or 1), this minimum N can be obtained from a table of expected two-tailed beta-binomial probability density function, such that accessible SNVs are all SNVs with a minimum number of reads, $n \geq \max(6, N)$. The minimum number of reads thus varies with the pooled size (coverage) of the ChIP-seq or RNA-seq dataset. Thus, the accessible SNVs are dataset-specific; they are determined for each pooled ChIP-seq (grouped by individual and TF, not by study) or RNA-seq dataset (grouped by individual). By considering only the cases with the largest effect size, we underestimate the number of accessible SNVs and this provides a conservative approximation of the statistical significance of the enrichment (or depletion). ‘Control’ SNVs are subsequently derived from accessible SNVs that are non-allele-specific, i.e. they are the set of accessible SNVs that has excluded the respective ASB or ASE SNVs for each dataset.”</i></p>

-- Ref2.5 – Heritability analysis --

Reviewer Comment	<p>The heritability analysis (using a single trio...) is confusing to me. It is, in a sense, a corrupted version of what is typically considered heritability analysis. The comment that analysis was performed separately for single parent - child pairs in order to 'maximize statistics' is entirely unclear, and in general, the entire analysis seems ad hoc and does not result in what we typically consider a measure of heritability.</p>
Author Response	<p>We thank for the reviewer for pointing out that this is not the measure of heritability of what we typically conceive. In our original manuscript, we did recognize that this is an adapted version of the conventional heritability analysis in population genetics. Even though we do not have a population of trios, there is still information about heritability that can be gleaned from a single trio. In fact, many other studies have also adapted their analyses in similar ways to show inheritance; we provide the citations for some of them:</p> <p>McDaniell, R. et al. (2010). Science. 328(5975):235-9 Kasowski, M. et al. (2013). Science. 342(6159):750-2</p>

	Kilpinen, H. et al. (2013). Science. 342(6159):744-7 We have re-worded it to better reflect this.
Excerpt From Revised Manuscript	Please refer to ‘Methods’ section under ‘Allele-specific inheritance analyses’. <i>“The conventional measure of ‘heritability’ allows the estimation of (additive) genetic contribution to a certain trait. The population genetics definition of ‘heritability’ in a parent-offspring setting is described by the slope, β, of a regression ($Y=\beta X + \alpha$), with the dependent variable being the child’s trait value (Y) and the independent variable (X) being the average trait values of the father and the mother (‘midparent’).⁶⁰ This is a population-based measure typically performed on a large set of trios for a particular trait (e.g. height) and β is not necessarily bound between 0 and 1... Given that we have only a single trio, we adapt the typical definition of ‘heritability’ to quantify allele-specific inheritance for each TF.”</i>

-- Ref2.6 – Causality of the AS variants --

Reviewer Comment	The analysis of functional annotation of the identified SNVs implicitly makes the assumption that these are causal variants. This is not the case, especially for the ASE, where the typed SNV is most likely in LD with the causal regulatory locus. As has been previously shown, the causal assumption is a poor one when ASB is considered as well.
Author Response	We are in total agreement with the reviewer that these AS SNVs are not necessarily causal and have never intended to imply causality in our writing. We will re-word the manuscript to better reflect this.
Excerpt From Revised Manuscript	Please refer to the ‘Discussion’ section. <i>“...it is important to note that the AS SNVs detected are not causal. The resultant allelic difference in gene expression and binding can be due to another undetected causal variant that has a strong linkage disequilibrium with the detected variant or, it could be due to a group of variants that act collectively to give the resultant allelic expression or binding.³⁷”</i>

-- Ref2.7 – Background description --

Reviewer Comment	Minor comments: From the intro: "AS variants can be detected regardless of their population allele frequencies." - This is actually not true in practice. ASE in intermediate frequency alleles are still easier to detect in the entire population because one can estimate the over dispersion parameter more precisely.
Author Response	The sentence was not meant to refer to the ease in detection, but rather the range, of allele frequencies that can be detected in allele-specific variants. We will re-word the manuscript to better reflect our intention.
Excerpt From Revised Manuscript	Please refer to the ‘Introduction’ section.

	<i>“Using each allele in a diploid genome as a perfectly matched control for the other allele, allele-specific variants can be detected even at low population allele frequencies.”</i>
--	---

-- Ref2.8 – Computational competence --

Reviewer Comment	I applaud the author's computational competence, but is the sentence, in the Results, on the amount of CPU time needed for the analysis really adds to the narrative? I think such details should be reported in the methods section.
Author Response	We agree with the reviewer and have moved the sentence from the main text to the Methods section.

Reviewer #3

-- Ref3.1 – General positive comment --

Reviewer Comment	This manuscript provides analysis of allele-specific binding (ASB) and expression (ASE) data for many individuals and assays. The authors compile this information in a database and further focus on describing the properties of transcription factors and genes which are enriched or depleted in ASB and ASE. I have the following major comments:
Author Response	We thank the reviewer for the thorough examination of the manuscript.

-- Ref3.2 – Using binomial test --

Reviewer Comment	The methods do not take into account the known statistical challenges of calling ASB or ASE and current advances in this area. The authors simply resort to a simple binomial test setting a minimum depth of 6 reads. It should be obvious that even with an FDR of 5 or 10% that low depth sites are going to be enriched in significant sites. They should plot depth by percentage of significant sites as a sanity check.
Author Response	<p>We agree with the reviewer's comment that using a binomial test, even with multiple hypothesis correction, significant sites will be enriched with low read depth sites. Hence, we have built a more sophisticated statistical model using a beta-binomial test to account for overdispersion. As mentioned also in response to reviewer #2's comment, we have now re-processed all the 1,280 ChIP-seq and RNA-seq datasets by first estimating the overdispersion parameter of each dataset. We then filter out those that exhibit a greater overdispersion (overdispersion parameter, $\theta > 0.3$ for ChIP-seq datasets and $\theta > 0.125$ for RNA-seq datasets) estimated from the allelic ratio distribution. Finally, we use a beta-binomial test to estimate and account for overdispersion during our ASE and ASB detection. The new Figure 1 and Methods section give a summary of our revised pipeline.</p> <p>Additionally, we provide a new supplementary figure (Supplementary Figure 2) to show that the percentage of AS SNVs called by our pipeline, with respect to the accessible SNVs is consistent across all read depths.</p>
Excerpt From Revised Manuscript	Please refer to Supplementary Figure 2 and the corresponding legend.

-- Ref3.3 – Compare AS calls with existing studies --

Reviewer Comment	Many of the datasets they use have ASB and ASE already called on them. The advantage of their approach to existing data is not compared. How different are there ASB/ASE calls to gEUVADIS or McVicker et al? One would expect that even if these are not online, the methods should minimally be compared as there will be an expectation that the AlleleDB database would yield similar quality results.																																			
Author Response	<p>As also noted by the first reviewer, it is important to appreciate that there is a fair amount of heterogeneity in the parameters and tools used in other studies, for instance the peak callers, aligners, detection strategies and reference genomes. Hence, there will naturally be great disparities and variability when comparing AlleleDB with results from the other studies. In fact, AlleleDB is motivated by the need to harmonize and uniformly reprocess all the datasets for allele-specific detection instead of simply combining the results from these various studies. We have included a new supplementary table (Supplementary Table 1) to show some of the differences in methods used across various studies.</p> <p>In response to the reviewer, we also did attempt to compare our AS SNVs and those called, if any, in the studies that are used in AlleleDB. We found the following:</p> <ol style="list-style-type: none"> 1) Lalonde et al. does not call allele-specific variants. 2) We were unable to find SNV lists for the study by Kasowski et al. 3) The ASB and ASE analyses performed by the ENCODE consortium were done with our pipeline but using the original binomial test. 4) The foci of Montgomery et al., Pickrell et. al, the gEUVADIS consortium and McVicker et al. were not AS analyses. They primarily utilized ASE variants for refinement of their detection of eQTLs or candidate regulatory variants. Unfortunately, they did not provide any ASE SNV lists online for comparison. 5) The only study we were able to make at least a minimal comparison was the study by Kilpinen et al., who provided the per-individual mean number for ASE, ASB and accessible sites that were called in their RNA-seq and some ChIP-seq datasets. Their analyses were for eight unrelated CEU individuals. We then computed the per-individual means for our corresponding datasets. <table border="1" data-bbox="448 1705 1408 1877"> <thead> <tr> <th></th> <th colspan="3">Kilpinen</th> <th colspan="3">AlleleDB</th> </tr> <tr> <th></th> <th>ACC</th> <th>ASB/ASE</th> <th>Proportion</th> <th>ACC</th> <th>ASB/ASE</th> <th>Proportion</th> </tr> </thead> <tbody> <tr> <td>POL2</td> <td>9,254</td> <td>525</td> <td>0.06</td> <td>2,817</td> <td>199</td> <td>0.07</td> </tr> <tr> <td>PU.1</td> <td>930</td> <td>154</td> <td>0.17</td> <td>83</td> <td>3</td> <td>0.04</td> </tr> <tr> <td>ASE</td> <td>4,061</td> <td>190</td> <td>0.05</td> <td>14,923</td> <td>1,266</td> <td>0.08</td> </tr> </tbody> </table>		Kilpinen			AlleleDB				ACC	ASB/ASE	Proportion	ACC	ASB/ASE	Proportion	POL2	9,254	525	0.06	2,817	199	0.07	PU.1	930	154	0.17	83	3	0.04	ASE	4,061	190	0.05	14,923	1,266	0.08
	Kilpinen			AlleleDB																																
	ACC	ASB/ASE	Proportion	ACC	ASB/ASE	Proportion																														
POL2	9,254	525	0.06	2,817	199	0.07																														
PU.1	930	154	0.17	83	3	0.04																														
ASE	4,061	190	0.05	14,923	1,266	0.08																														

	We note that even though the proportions of accessible SNVs being AS SNVs are very similar, they have more ASB but less ASE calls. The differences are probably due to the discrepancies in parameters in the detection pipelines.
--	--

-- Ref3.4 – AS inheritance at binding sites --

Reviewer Comment	How is AS inheritance at binding sites not a universal phenomenon of TFs like MYC or RPB2? This seems like a pretty bold assertion. Isn't it more likely that there is something wrong with your method for these sites? Low read depth, poor antibody efficiencies, non-specificity of binding profiles, etc. Why make a biological claim before you have exhausted technical sources of error.
Author Response	We agree with the reviewer and have removed the assertion from the revised manuscript.

-- Ref3.5 – Causality of AS variants --

Reviewer Comment	The phrasing of the paper suggests that ASE sites are actually causal. For instance, the relative numbers of sites with ASE and ASB are compared. What does this even mean? How are these even directly comparable? An ASE effect suggests the genes are imbalanced. This imbalance could be due to multiple causal ASB events. Furthermore, the ASE site is not causal. It is only indicating the potential presence of a causal regulatory variant.
Author Response	We are in total agreement with the reviewer that these sites are not necessarily causal and have never intended to imply causality. We have re-worded the manuscript to make this point explicit. The AS sites are used as markers for potential allele-specificity in the genome. By visualizing ASE and ASB SNVs side by side, we had meant to provide some context and possibly set the stage for some biological insights. We have included a new figure, Figure 4, to show how concurrent visualization of ASB and ASE SNVs in a population of individuals for the gene <i>ZNF331</i> can potentially provide some biological insights into the allele-specific properties of the gene and its specific sub-regions (exons).
Excerpt From Revised Manuscript	Please refer to the 'Discussion' section. <i>"However, even with lower number of false positives, it is important to note that the AS SNVs detected are not causal. The resultant allelic difference in gene expression and binding can be due to another undetected causal variant that has a strong linkage disequilibrium with the detected variant or, it could be due to a group of variants that act collectively to give the resultant allelic expression or binding.³⁷ It could also be due to other epigenetic effects such as genomic imprinting where no variants are causal..."</i>

-- Ref3.6 – Allele-specific behavior in genes --

Reviewer Comment	The authors don't seem to understand why a gene would be depleted in allele-specific behavior. Is there expectation that allele-specific behavior should influence all genes equally? Furthermore, I worry that depth might be more deterministic of which genes are enriched or depleted.
Author Response	We agree with the reviewer that we do not expect allele-specific behavior to influence all genes, or even sub-regions within a single gene, equally. Indeed, genomic regions with lower read depth will not have enough power to detect AS behavior. However, for regions with sufficient read depth, our implementation of an explicit FDR simulation takes into account the respective read depth of each heterozygous SNV in order to determine whether a SNV is allele-specific or not. Additionally, to determine whether a gene (or genomic element) is enriched or depleted in allele-specific behavior, we incorporated the use of 'control' SNVs that are matched in power to the allele-specific SNVs.

-- Ref3.7 – Selection in ASB and ASE sites --

Reviewer Comment	Why would ASB be under less selective constraint than ASE SNVs? This probably only has to do with the background of being in a gene versus being in a non-coding region. Again ASE SNVs are not causal, so what is selection acting on. Figure 4 makes no sense. Beyond this, I don't even see a difference between the ASB +/- sites at low frequencies.
Author Response	<p>We had intended to report the rare variant load in ASE and ASB sites only as a general observation that might be suggestive, and not indicative, of less natural selection in ASE sites. We have re-worked the main text to better reflect this.</p> <p>We agree with the reviewer that it is reasonable that the background of being in a gene versus a non-coding region can contribute to the higher enrichment of rare variants in ASE SNVs than ASB SNVs. Hence, we have removed the sentence that makes this ASE-ASB comparison.</p> <p>Furthermore, inspired by the reviewer's comment, we have constrained our analysis to only the coding DNA sequences and the transcription factor binding motifs and re-calculate the rare variant loads in ASE, accessible non-ASE, ASB and accessible non-ASB sites. Indeed, we do not find any significant enrichment of ASB sites, compared to accessible non-ASB sites.</p>
Excerpt From Revised Manuscript	<p>Please refer to 'Results' under 'Rare variants and purifying selection in AS SNVs'.</p> <p><i>"Our results in Figure 6 show a statistically significant lower enrichment of rare variants in ASE SNVs as compared to non-ASE SNVs (Fisher's exact test odds</i></p>

	<i>ratio=0.2, $p < 2.2e-16$) but statistically insignificant higher enrichment of rare variants in non-ASB SNVs than ASB SNVs (Fisher's exact test odds ratio=1.4, $p=0.08$). This observation suggests that ASE variants may be under weaker selection than non-ASE variants."</i>
--	--

-- Ref3.8 – Replicates --

Reviewer Comment	Do the authors have any insight into how well their calls replicate and then their replication at various depths.
Author Response	<p>In response to the reviewer's suggestion, we have included two new analyses and corresponding figures as supplementary materials.</p> <p>The first analysis investigates the replication of AS calls between two technical replicates by comparing their AS calls. We obtained two equally-sized subsets of 245M ('M' denotes 'million of reads'), which are randomly sampled from a pooled RNA-seq dataset of NA12878 without replacement, i.e. these two sets are mutually exclusive. We show that our calls replicate very well (>70% overlap) (Supplementary Figure 3).</p> <p>The second analysis investigates the replication of AS calls at increasing read depths randomly subsampled from the same pooled RNA-seq dataset of NA12878, namely subsets of 100M, 200M, 300M, 400M and 490M. Each subsample is a direct subset of the other larger subsamples, e.g. 100M is a direct subset of all the other sets. We then re-align each pool of reads to NA12878 personal genome and re-call the AS sites for each subset. Supplementary figure 4 shows that the majority of the AS sites (>77%) are consistent in at least 2 subsets. There is also only a small proportion of sites that are unique to each set (< 6% for each set, except for the entire set of 490M, which has the proportion at ~17%).</p>
Excerpt From Revised Manuscript	Please refer to Supplementary Figures 3 and 4 and their corresponding legends.