



# UNIVERSITÉ DE GENÈVE

FACULTÉ DE MÉDECINE

DEPARTMENT OF GENETIC MEDICINE & DEVELOPMENT

CMU - Rue Michel Servet 1 / CH-1211 Genève 4  
phone. +41(0)22 379 54 83- Fax +41(0)22 379 57 06  
[www.unige.ch/medecine](http://www.unige.ch/medecine)

**Professor Emmanouil (Manolis) Dermitzakis**  
emmanouil.dermitzakis@unige.ch

RE: Nature manuscript 2012-12-15883

February 7, 2013

Dear Magdalena,

Thank you for the decision letter, even though the decision was not what we had hoped for. All the three reviewers provided valuable feedback, and we are happy that they all appreciated many aspects of the manuscript. Given the amount of positive comments, the fact that some key points of criticism were based on misunderstandings, and that we feel we can address the remain go the key comments, we would like to ask *Nature* to reconsider the decision to reject the manuscript.

The questions that the reviewers raise about individual analyses are straightforward to address (see below). However, there were two major concerns raised by the reviewers: (1) the novelty value of the study, and (2) technical quality and replicability of RNAseq data processing. We would like to respond to these main lines of criticism in the following:

- 1) We must disagree with the Reviewers 1 and 3 that question the novelty of our conclusions. Even though this is not the first population-level RNAseq dataset, this is the first one with multiple populations. Furthermore, we claim that this is actually the first really mature RNAseq study of human populations, where the high quality and quantity of raw data as well as developed analysis methods allows us to get beyond technical issues to really discover novel biological phenomena. Many of our results are novel and far from obvious: e.g. the special role of splicing in population differences, the abundance of regulatory variation that was never discovered with even large array studies, the independence of genetic variants affecting transcript levels and transcript usage, miRNA-mRNA interactions in populations, and the location of eQTLs with respect to recently reported comprehensive functional annotations. Furthermore, our characterization of causal variants and validation of loss-of-function variants contributes directly to some of the hottest topics in the field and has real practical implications for medical genetics studies, and serves as a model for combining other types of phenotypic data with genome sequencing data.
- 2) Reviewers 2 and 3 expressed general concern about robustness of the results regarding the complex RNAseq data processing steps, which is a reasonable general concern in all studies with big data. However, at the same time the quality of the analysis is complimented, and the few specific concerns we address below. We feel that it is unfair to our study to judge it based on general technical concerns that potentially might exist in the absence of specific indications of any. While it is impossible to share every line of our data processing scripts, we have made substantial effort to be transparent and share data, resources and documentation. This will allow our analyses to be replicated by others, which was a major concern of Reviewer 2 who mistakenly thought that data access is password protected. Specifically, the raw data has been freely accessible since November 2012 (<http://www.geuvadis.org/web/geuvadis/RNAseq-data-release> with links to ENA), our mapping pipeline is available online (<https://github.com/gemtools/gemtools>), and the project wiki has additional analysis results and discussion. The wiki will be open once the paper is published (which is more than most projects do; the reviewers may not have noticed this in the data sharing section), and we can provide immediate wiki access to the reviewers for the review process. Our methodological choices are naturally supported by much deeper analysis than what was included in the Supplement due to length constraints, but we will make additional effort to improve documentation of data processing in the Supplement.

The reviewers as well as the Editor expressed concern that the manuscript might be too dense, with analysis of a wide spectrum of topics that are difficult to cover in adequate depth, rather than focusing on a few key points. We agree with this to some extent, and we suggest omitting the following topics from the manuscript in order to improve its focus:

- Gene discovery as a function of sequenced individuals (pp. 4-5, Fig S10)
- Subtle splicing (pp. 5, Fig S13)
- Fusion transcripts (pp.5-6, Fig S14, Table S4)
- RNA editing (pp. 8, Fig S20)

Additionally, we would like to strengthen our most important theme of causal regulatory variants by adding the following novel analyses, which we have essentially completed:

- mapping rare putative regulatory variants using allele-specific expression data
- as per the advise of reviewers we have explored the functional validation of our eQTL putative causal variants in two published datasets. We have found strong enrichment of allele-specific binding of CTCF, as well as strong enrichment in DnaseI QTLs (dsQTLs). We will present these results in the revised version.

With these changes, the manuscript would have the following key themes:

- Quantitative versus qualitative variation in human populations: analyses comparing expression levels and transcript structure in (1) the transcriptome alone, (2) QTLs, and (3) allele-specific effects
- Micro-RNA variation in human populations and its effect on the transcriptome
- Functional mapping of putative causal regulatory variants using data from eQTLs, allele-specific expression, and loss-of-function variants, as well as characterizing the properties of these variants

We would like to emphasize that this is not a narrow hypothesis-based study addressing a single important question, but rather a reference study of the use of transcriptome sequencing and genome sequencing to understand genome function variability. We are confident that the biological insights, the methodological and technical advances, and the easily accessible data will have a large impact in functional and medical genomics in a similar manner to the 1000 genomes and the HapMap papers. In line with this, the main paper will be followed by a series of companion papers that analyze various topics in further depth: e.g. technical aspects of RNAseq (included in the main paper submission, now submitted to Nature Biotechnology), population variation in splicing, splicing QTL mapping methods, loss-of-function analysis, population genetics of regulatory variation, and noncoding RNAs. We preferred to prepare the main paper first rather than keep it waiting for the companions, so that the community will have access to the data and the main results without delay.

Additionally, we respond to the specific comments in below, which we will update to a detailed response letter with additional analysis results if we are given the chance to revise the manuscript:

We hope that this response is satisfactory in order to allow for a revised manuscript to be considered.  
Sincerely,

Tuuli Lappalainen  
Emmanouil Dermitzakis

### **SPECIFIC RESPONSE:**

#### **REVIEWER #1**

**The authors report an RNA-seq study of 465 cell lines from the HapMap/1000 Genomes collection. This is the largest RNA-seq study to date (although comparable in size to published microarray studies in many of the same samples). The paper describes the functional context of eQTL variants, splicing variants and so forth. Overall, the paper is well written and clearly presented.**

**That said, unfortunately, I feel that we do not learn enough here to justify an article in Nature. While this data set is bigger than earlier RNA-seq data sets, most of the take-home points in the paper have been considered elsewhere, especially in papers by Lee et al (2009), Montgomery et al (2010), Pickrell**

et al (2010), MacArthur et al (2012) and Gaffney et al (2012). I think that the paper would be better suited to a specialty journal.

RE: We discuss the general novelty issues of our study on page 1. Many of our results are completely novel, and we feel that early indications of some of the results in previous studies should not preclude high impact publication of the definitive result.

**Figure 1c: "transcript structure variation appears to contribute disproportionately to continental differences, suggesting a special role of splicing in human adaptation." The authors report an enormous difference between within-continent comparisons and between-continent comparisons. It's hard to conceive of a plausible model--either neutral or adaptive--that is consistent with this enormous variation in the relative fractions of gene-level vs transcript level changes, leading me to think that this is most likely artifactual. Perhaps this observation may be due to differences in power at the two levels?**

RE: In this analysis, we compare the relative contributions of expression level differences and transcript usage differences between populations, and we observe that is radically different for intercontinental population pairs. We have several reasons to be convinced that our result is driven by real biology. Figure 3c is based on analysis of differentially expressed genes/transcripts, but the same trend is seen using another method (quantitative analysis in Figure S12). Furthermore, two recent papers in *Science* (PMID 23258891 and 23258890) describe adaptive evolution through splicing changes in interspecies variation, which is completely independent analysis but perfectly consistent with our result from human populations. We do not think that the difference YRI-EUR vs EUR-EUR population pairs is too large to believe – the reason why it hasn't been found before is that this is the first data set where this can be properly analyzed. Finally, the sample sizes of different populations are similar, the sample processing for sequencing was randomized, and the effect is relatively consistent across cell line collections of different age. We rigorously explored the data and analysis and could not find any technical reason why YRI-EUR population pair comparisons would be systematically different from EUR-EUR population pairs.

**Several of the figure panels are relatively uninformative, including most of figure 1 as well as figure 3b. The analysis of NMD displayed in figure 4 is a simple update of a corresponding figure in MacArthur et al.**

RE: These figures show distributions of different phenomena (rather than barplot-type plots that capture numerical summaries of the data), and we feel that together with the summary statistics given in the text, a visualization of the shape of these distributions is important and informative. The loss-of-function analysis in Figure 4 shows inter-individual variation in NMD unlike MacArthur et al., and the quantification of NMD described in the text is more advanced and definitive with many more sites interrogated.

## **REVIEWER #2**

**This study represents some of the very best in contemporary genomics research, but also one of its darker aspects. The best is that you cannot fail to be impressed by the magnitude of the study, quality of the analysis, comprehensiveness of treatment, and attention to ensuring that the data becomes a crucial community resource. The dark is that it is really 20 papers condensed into one, with some very interesting topics given just a sentence with a deeply buried supplementary figure (eg RNA editing SNPs), but more troubling is that it also means that several of the major findings are not given the depth of coverage to convince at least this reviewer beyond reasonable doubt.**

**With over 50 pages of supplementary material, it is unlikely that my review comments will jibe with those of other reviewers, and I have to confess that I am concerned that the most appropriate reviewers should be post-docs or students who know their way around the details of the bioinformatics tools. I had some of my people read it and our consensus was that there is really no way of being sure of detailed claims without replicating the studies with each group's favorite pipelines. No doubt the big picture is largely correct, but my major comment is that data access is going to be crucial.**

RE: We would like to thank the reviewer for the enthusiastic support and for raising some caveats, which we are happy to respond to. After discussions with the Editor, we are willing to drop some analysis (see page 2) in order to improve the focus of the paper and make it less dense to read. We discuss the questions of data access and companion papers above (page 1-2), and we will indicate this better also in the manuscript. While we fully appreciate and will strive to provide data and information for other researchers to replicate our

findings, the review process in a journal should rely on whether the methodologies described are reasonable and not whether they can be replicated during the review process. Having said this, we would be happy to provide any information or intermediate data files for the reviewer to test and execute some of our analysis. We truly hope that many other research groups will use the data to replicate our findings, continue on analyses that we did not fully cover, and explore completely new questions.

**The authors have promised to make the data available through ArrayExpress and the European Nucleotide Archive, but the latter is password protected. Please clarify what the nature of the Data Access Committee will be, and what steps potential whole genome analysts will need to take to gain access.**

RE: We are extremely surprised if the reviewer has been prompted for a password when trying to access the data – this should not happen, the raw data and bam files have been freely accessible to everyone since November 2012 through ENA, without any passwords or data access committees. We will clarify the open access policy better in the Data Access section, and will provide a detailed guide of the different data files in the Supplement as well as in the Geuvadis project web page.

**The second is the thousands of candidate gene surveyors who are more likely to simply want access to the Geuvadis server, which is a terrific resource, but I think it is essential that users also gain access to the single gene-relevant data upon which the graphical output is based. Please clarify whether users will be able to download the sequences (genome, RNASeq, miRNASeq (or at least inferred abundance in tab format) for single genes. Graphical output is nice, but since it only hints at the causal variant, sophisticated users will want the actual data without having to access the entire dataset.**

RE: We agree with the reviewer that this is a valuable element for the Geuvadis server, and we can enable download of quantifications by gene or region. We will do our best to make additional data types available, with the understanding, of course, that our project was not funded to produce a full-blown browser so some limitations may exist.

**1. That most of the variation in exon-level read counts is due to transcript ratios rather than whole gene expression, and that a significant proportion of this is among populations. This conclusion is probably quite sensitive to the metric used to assess the relative proportions of tr and ge. Readers have to dig all the way through the Suppl Methods to find that it follows Gonzales-Porta et al (ref 46, not 47), the philosophy of which needs to be explained briefly to readers in the text. My concern is that it is sensitive to the contribution of relatively rare classes, and that if there was an exclusion of rarer transcript types the inference would be quite different.**

RE: We will clarify the methodology in the main text. We have analyzed the contribution of lowly expressed transcripts to these metrics, and the results are consistent; we will include these results in the Supplement.

**What may be very confusing to readers is the juxtaposition of this inference with the conclusion that there are several thousand eQTL for gene abundance but only a few hundred for transcript ratio abundance. This is explained away as a power issue, but needs to be dissected more deeply: is it that power is lower because the abundance of contributing classes is low?**

RE: Read counts of exon levels as well as whole gene levels are observed quantities that are well measured given the depth of the read data. On the other hand, transcript ratios are estimated based on models of read pair distribution and known annotation and this estimation is noisier, and thus eQTL analysis of exons and whole genes is more powerful than in transcript ratios. It is also possible that that transcript usage is truly less affected by genetic variants than expression levels, but given these technical constraints, our data does not permit such a conclusion. We will clarify this in the text and methods.

**Another aspect relevant here is the statement (which I completely missed on first reading) that most of the transcript ratio variance is not exon-skipping, but in 5' and 3' terminal exons, implying that the tr variation does not contribute greatly to protein variation.**

RE: We agree with the reviewer that this is an important point, and we will emphasize it in the text.

## **2. That RNA-Seq helps in the annotation of causal variants.**

This section could also be presented more clearly. It starts with the observation that there is enrichment for ENCODE-type features that is particularly notable for the "best" eQTL, then builds into a somewhat ad hoc estimate that 57% of the best eQTL are the true eSNP based on a somewhat arbitrary  $NLP > 1.5$  cutoff (the plateau is more like a gentle slope), and concludes that there is enrichment for GWAS hits. I think that Figure 2D is really the critical one, and suggest turning the discussion around, beginning by asking the question directly "how does RNA-Seq improve disease variant annotation", presenting a case study, and leading the reader through the process. The fact is that in half the cases the "best" eSNP in a statistical sense is not going to be the causal one, which is to be expected (the strongest statistical SNP is likely to be a function of methodology and sampling anyway). New Bayesian and other methods are being developed that will help refine likely causal variants further, which places a premium on full data access as argued above. My point is that enrichments of enrichments are not overly convincing, so perhaps better just to walk readers through how to use the resource rather than generate difficult-to-interpret figures like 2A.

RE: We will reconsider the structure of that paragraph, and we will clarify that our quantification of causal variant discovery is meant to be a lower limit. If we understand correctly, the reviewer shares our view that developing novel sophisticated causal variant modeling methods is beyond the scope of this paper, but we agree that our data will be of great importance for this purpose in the future.

## **3. That there is substantial variation due to relatively rare variants that impact transcript ratios in particular.**

This is very believable, but I am just nervous about it because RNASeq analysis is subject to so many biases. This is not really a criticism, just a statement. In general the authors have been very open about their analyses, use state of the art methods, and are very careful. But, different methods of normalization, alignment, transcript inference, and association testing could yield different results. I am not asking for reanalysis, but do feel that a cautionary statement and call for verification with other methods is appropriate. For example, PEER is excellent for cis-eQTL analysis, but I can well imagine it upsets the trans-miRNA associations and maybe affects the rare exon variant distribution. Similarly, the Flux Capacitor approach is just one way to infer transcript structure from paired-end reads - different graph theoretic algorithms (eg Cufflinks) will likely give different results. Of course the authors have the right to choose their method, and others can replicate it, but less familiar readers should have a sense of how robust the major biological conclusion is to methodology. Just one example of why one might be skeptical is in Fig 1C: it makes little sense that the tr:ge percentages are so different for the GBR:TSI and FIN:TSI comparisons than the CEU:TSI and CEU:FIN ones.

RE: We agree with the reviewer that it is difficult to evaluate the robustness of the results when data processing is becoming increasingly complex; see our response on page 1. We can assure that the analyses and data processing methods have been evaluated to a much greater depth than what is presented even in the Supplement due to length constraints. We have made substantial effort to make our analysis transparent, be conservative, evaluate the robustness of our methodology, and provide diagnostics of the most important methods (e.g. transcript quantification in Fig S11 and PEER in Fig S7). As an additional confirmation of the robustness of the results, we compared our final set of exon eQTLs to those obtained from the very first preliminary analysis based on a different mapper, normalization method, and a smaller sample set, and 86% of the smaller set of preliminary eQTLs replicate in the final set. Furthermore, many results are supported by two methods (e.g. Fig 1c and Fig S12), and we refrain from biological interpretation of results when we are not confident that the result is not affected by technical bias. The differences between European population pairs in Fig 1c are a good example of this as it may be affected by cell line age – however the main conclusion that we reported, the difference between YRI and European populations, appears reasonably robust to this (see also response to Reviewer 1 on page 2).

### **Some minor points:**

**The inference that the differences between YRI and EUR expression profiles is due to genetic divergence should be softened, given that batch effects of LCL production are known to cause expression divergence (see the Cheung/Storey discussion a few years back)**

RE: We will soften the statement. See also response to Reviewer 1 on Page 2.

**Regarding the lab effect, Fig SF6 shows that it is actually quite substantial, and much stronger than population effects in the raw data. That PEER removes this is not surprising, and is good news for the analysis, but it does not necessarily mean it properly adjusts for it. I feel the authors make too strong a claim that lab differences are negligible.**

RE: We have analyzed and quantified this further in the accompanying paper ('t Hoen et al.) that was included in the submission. Importantly, the biological differences between our samples are very small and thus e.g. in Fig S6 the lab effect being the strongest trend in the data still does not imply that it is a very strong effect.

**The section on miRNA-mRNA covariance requires more scholarship - there is quite a literature claiming feedback etc in cancer samples to explain relatively low correspondence between suspected targets and miRNAs. Please add some context and references.**

RE: We will revise the text of this section to make it more accessible and up to date as indicated.

**Middle of p7 it is claimed that gene eQTL and tr eQTL are expected to be orthogonal, contrary to the observed 45% overlap. It is not clear to me why they are expected to be orthogonal, particularly if most of the tr effects are in the 3' and 5' ends**

RE: We agree that the wording of that particular sentence was not optimal. What we meant by orthogonal was that our measurement of gene expression levels and transcript ratios are independent – a biological change in only one does not change the statistical measure of the other, thus any correlation is likely to be biological rather than technical.

**In discussing the RTS analysis, we need an explicit definition of "best" eQTL. Is this just the one with the smallest p-value, or are you referring to the best one conditioned on other evidence such as TFBS, DHS, histone proximity?**

RE: The best eQTL is just the one with the best p-value, as we defined on p. 8.

**One potential source of technical error contributing to Fig 3B is read alignment artifacts. Were QC checks performed to ensure that all ASE and ASTR tests were done on SNPs that are present in both the RNASeq and WGS data for that individual? Ideally the alignments would be performed against the person's individual genome, not RefSeq - does this have an impact on rare ASE annotation?**

RE: We fully agree that allelic mapping bias is an important issue in RNAseq. We have actually done a substantial amount of work to address this issue by simulating mapping bias of all 1000 Genomes variants, and we will add a section about this in the Supplementary Methods. Based on these simulations, we exclude about 12% of variants that we consider unreliable due to mapping issues. In addition, when personalized mapping was performed we observed little to no difference in ASE estimates vs. mapping to the reference genome and correction based on simulations. The SNP genotype data accuracy is much higher for the 1000 Genomes Phase 1 data than for the pilot so genotyping errors are not a major issue (as seen in Fig S28), but to account for any remaining errors, we use only sites where both alleles are observed in RNAseq data. Furthermore, we are preparing a separate paper addressing the effect of mapping bias on eQTL discovery, and we have found that the effect is negligible and causes false eQTL discoveries only in about a dozen genes.

**Fig 3c shows that ASTS tend to be enriched in introns and 3' UTR whereas ASE are relatively enriched in splicing regions, contrary to what I would expect. Please discuss.**

RE: We first found this surprising as well, but this pattern actually makes sense when considering the fact that we are only assessing the reads that cover a particular site. At least 16 reads of coverage over fully intronic sites is rather unusual, and can be due to rare intron retention effects or unannotated exons, both with unusual splicing easily leading to an ASTS signal. The ASTS signal in the 3' end is less surprising knowing the widespread variation in 3' ends of transcripts. Splice site variants are an interesting case; for example, let's consider a nonreference allele that decreases splicing efficiency by 60%. In this situation, 60% of nonref transcripts would skip the exon or undergo NMD, leading to a strong ASE signal over the splice site – however, the remaining 40% of the reads over the splice site that we would assess for ASTS would have the

normal splicing pattern, without ASTS over this site.

**What is the average read depth of the 1000G data?**

RE: We assume that the reviewer is referring to the genome sequence data. For the low-coverage sequencing of the whole genome the average depth is 5x, and for exomes it is 80x.

**Which FDR method was used throughout - some figures refer to  $Pi_0$  estimates, suggesting q-value analysis, but I suspect Benjamini-Hochberg FDR has been used.**

RE: In QTL analysis we used permutations, and in the miRNA-mRNA correlations Holm's procedure was used. The  $Pi_0$  statistics e.g. in Figure 3a is used as a quantification of sharing of the signal and not to estimate FDR per se.

**This is a lot of critique, but there is a lot to review. I just want to close by complementing the authors, and emphasizing that the criticisms are intended to assist in accessibility for a wide audience. In writing this review, it became apparent to me that my major recommendation for the paper is to seriously consider concentrating on just the few main biological issues and provide more evidence and analysis, while dropping the references to half a dozen other topics that really belong in other papers. I don't feel strongly about this - it is all interesting - but there is some frustration with strong claims that are probably true but could be supported with a sense of robustness to methodology for a subset of the data.**

RE: we agree with the reviewer that some of the topics may be distracting the reader from the core messages so we are happy to remove some of the less important topics with the assistance of the Editor.

**REVIEWER #3:**

**This paper provides a comprehensive analysis of human genetic variability. The authors acquired deep coverage RNA-Seq data from 465 individual samples covering 5 populations of the 1000 Genomes project and carried out intensive transcriptome and genotype comparisons within and across populations. Based on the reported statistics of the individual samples, libraries, and biological replicates, we can see the overall data quality is very good. With these data, authors explored transcription variation, eQTL and trQTL, loss of function variants and identified some new loci as the causalities of phenotypes/diseases. Most methodologies adopted in this work were all recently developed. These methods tackled many known biases. However, there were still questions in need of clarifications.**

**Overall Verdict: Major Revisions are needed**

**My overall feeling on the manuscript is that this is certainly an interesting subject area connecting functional data to genome variation, and the paper represents a substantial dataset. I was quite excited to read it. I do believe it has merit. However I think it requires large revisions before being acceptable. In particular the authors have to address the question of what novelty their analysis brings to this and they have to address a lot of technical RNA-seq processing issues, which I find fairly sloppily dealt with here. Again these are areas that I think can be fixed but will require substantial recalculations. It is somewhat disappointing that the authors did not do this initially.**

**The fundamental issue with this paper as I see it is, what is new? Obviously one has about 10 times as much sequence RNA data as previously in the path-breaking Pickrell et al. and Montgomery et al. papers. A lot of data has been generated but what new results do we see? The authors have to make this clear. It cannot just be simply that we find more eQTLs. One really has to understand how all this data gives rise to new ideas. I think the authors need to spend a considerable amount of text talking about how their current work qualitatively, not quantitatively, goes beyond previously published analyses, if, indeed, it actually does, with clear demonstration that the earlier results would not be possible with the smaller datasets.**

RE: We discuss the question of novelty value on page 1. While we understand the reviewer's point, we believe we have added substantial biological insight and analyze the data far beyond merely reporting the numbers of eQTL discoveries. Our results also strengthen signals that previously were only suspected. We

did not feel the need to justify the novelty of our analysis with point-to-point comparisons to previous papers in the text, but we will make an effort to revise the text to address the Reviewer's concern.

**More to the point, the results are consistent with our understanding and expectations from human genetic variability studies. However, just because most results could be anticipated, it seems there is lack of thrilling discoveries from this wonderful data set. Also, some more work such as following up experiments to show the molecular/biochemical mechanisms of the discovered causal regulatory variants could add great value to this work.**

RE: While it is of great value to perform follow-up experiments, it is out of the scope of this study. Testing biochemical mechanisms requires a completely different set up, and it would be difficult to achieve reliable validation estimates beyond individual case examples that in our opinion would be of limited value. However, we have now used existing data sets from McDaniell et al. (Science 2010) and Degner et al. (Nature 2012) to investigate allelic effects of our eQTLs in CTCF binding as well as enrichment in DNase I hypersensitivity QTLs and have observed strong enrichment in allelic binding effects (CTCF) and overlap with dsQTLs (DNaseI). Furthermore, the samples are available from Coriell and used by many laboratories, so in the future there will hopefully be additional data sets that can be linked to ours.

**The paper discusses the relationship between variants and total gene expression level, the expression level of exons and also of transcripts. Obviously the biologically relevant unit here is that of transcripts and one would expect to get the strongest correlations and the clearer signals in relation to that. In the paper the converse is true, far fewer eQTLs are found for transcripts than for exons and genes. Perhaps this is because of deficiencies in transcript structure in some of the RNA-seq processing method used. This needs to be explored in considerably greater statistical detail to understand why one is getting such a counterintuitive and to some degree obviously wrong result.**

RE: There seems to be a misunderstanding here – the trQTLs are not for transcript expression *levels* but for transcript *ratios* with the aim to find variants that affect the choice of transcript that is expressed rather than the level of transcriptional activity. These are two different quantitative traits, and thus the results are not expected to be similar. See also response to Reviewer 2 on page 4 discussing why the numbers of eQTLs and transcript ratio eQTLs may differ.

**Also as I am sure the authors well know, there is a huge controversy in the world of RNA editing, a lot of it requiring substantial validation. The authors should allude to this controversy and either improve the content of the manuscript technically or remove this section. Their discussion of editing could have been improved by more validation related to these results.**

RE: The difficulty of reliable calling of editing sites is the reason why we did our analysis in *known* edited sites rather than inferred ones from RNAseq and, importantly, our primary interest is in individual differences in the degree of editing, which is much less sensitive to error than determining if a particular site is edited in general or not. However, in order to obtain better focus of the manuscript and after discussions with the Editor, we are willing to omit this analysis from this manuscript.

**3. I was unimpressed with the consideration of read mapping. The whole point of the paper is to look at variants in relation to RNA-seq and one's using essentially a mapping strategy does not take into account these variants, mapping everything directly onto the reference genome. There are many additional papers now that are starting to take into account this type of thing, either in relation to local assembly or other strategies. I really think that in this type of paper in 2013 one should see a more developed mapping strategy being utilized. This becomes particularly important for editing and for allele analysis as well as for the eQTLs. See <http://www.ncbi.nlm.nih.gov/pubmed/19808877> and <http://www.ncbi.nlm.nih.gov/pubmed/21935354> .**

RE: See response to Reviewer 2 on page 5-6.

**4. Another issue was the linear regression for quantifying eQTLs. There have been several papers discussing various algorithms to detect eQTLs (e.g. <http://www.ncbi.nlm.nih.gov/pubmed/19303049>), and the papers claimed significant inconsistencies of eQTLs resulting from different algorithms. Some suggested the linear regression model might overestimate the amount of eQTL (such as in the same paper). It is thus necessary to further examine eQTLs using some other models.**



RE: We agree that linear regression is not the most sophisticated eQTL method, and is known to generate false positives when outliers are present. However, linear models are useful when one wants to use covariates, and transformation of our data to standard normal and performing permutations makes us immune to outlier problems. In our long experience with eQTL analysis, the inconsistencies described by the reviewer, while they exist, they do not necessarily suggest incorrect analysis but rather differences in the type of discoveries. In fact, there is a lot of debate as to whether more complex models are better or relevant to eQTL analysis than the more traditional SNP by SNP tests and we did not feel that our paper is the place to address this debate.

**5. A final technical question is the normalization of mRNA/miRNA. A recent paper showed that c-Myc had a huge impact on the total amount of RNA, and if this effect was ignored, it was very likely to draw some unreliable conclusions. (<http://www.ncbi.nlm.nih.gov/pubmed/23101621>). We do not know the direct transcription variations of c-Myc across the 465 individuals. But as shown in the supplementary results (Figure S22, Figure S24), the gene c-Myc was among the identified eQTLs as one of outstanding example. If c-Myc transcription variation was indeed large, the authors will have to take this aspect into consideration and revise the results accordingly.**

RE: Our focus on cis-eQTLs makes our analysis much less vulnerable to such effects that would indeed be a serious concern for trans-eQTL analysis or transcriptome profiling. The figures S22 and S24 do not actually indicate that c-Myc has or is an eQTL, but that c-Myc binding sites are enriched for cis-eQTLs. This does not imply that c-Myc levels in trans could drive these eQTLs in cis. More generally, we are aware that both known and unknown factors may be influencing global levels of expression but it is hard and statistically risky to correct for a large numbers of covariates. Instead, our approach to correct for PEER factors is expected to eliminate the effect of the global factors with strong effects. We further verified that we are not losing signal in the miRNA-mRNA analysis that is searching for trans effects by using a normalization that corrects only technical covariates, and we found hardly any additional correlations absent from the outcome of the PEER-normalized data.

**Page 3: "challenge has been to analyze cellular phenotypes, such as gene expression, resulting in large catalogs of regulatory variants, known to affect many human diseases and traits"  
It would be better to use one concrete example to show the power of genomic analysis to reveal disease causality.**

RE: We prefer to keep the introduction without examples in order to keep the manuscript within reasonable length but we are happy to add examples if considered necessary.

**Page 5: "As expected, the vast majority of the total transcription variation is among individuals within populations, with only 3% explained by population differences. Yet, between population pairs, we detect 263-4379 genes with significant differences in expression levels and/or transcript ratios"  
First, please clarify if the 263-4379 genes constitute the 3% population difference? And how about the likelihood that the 3% across population differences were due to some statistical randomness?**

RE: The analysis of total transcription variation (the 3%) and the analysis of differentially expressed/transcribed genes are two separate methods to analyze the same question, the only thing in common being the same transcript quantifications. The 3% is an overall genome-wide estimate, and additional analysis permuting the population labels of the samples has confirmed that this is indeed far beyond what would be observed by chance and we will include this result in the revised manuscript.

**Page 7: "Even though these quantitative traits are expected to be biologically orthogonal, we find a significant enrichment of genes with both types of QTL (279 genes = 45% of trQTL genes = 2.15× enrichment,  $\chi^2 p < 2.2 \times 10^{-16}$ )."  
Mechanisms of these two types of QTL are different, but this doesn't necessarily say if a gene is a QTL, it has to be subject to only one type of regulation. I would rather expect if a gene is a QTL, there might be a higher chance the gene is involved in more than one mechanism of regulation.**

RE: This is exactly the analysis that we describe in this section. We will try to write this in a clearer manner.

**Page 9: "In 13% of the genes the strongest eQTL variant is an indel, which is 37% more than for the matched null variants (Fisher exact test  $p = 5.6 \times 10^{-6}$ ; Fig S21), suggesting that indels are more likely to have functional effects than SNPs"**

**How many of this type of eQTL are protein-coding genes?**

RE: All the eQTL analysis is done on only protein-coding genes, as indicated on page 52 of the Supplementary Material, but we will specify this in the main text as well.

**Section on "Characterization of regulatory variants"**

**In relation to how the authors connect the eQTL to functional sites I would urge them to look at the DNase data (<http://www.ncbi.nlm.nih.gov/pubmed/22307276>). There have been many recent papers on this and I think that not discussing this data is a major limitation of the paper.**

RE: DNase hypersensitive sites are included in our annotations, and the strong enrichment in them is described in Fig 2a, S22, S23, and we also refer to the Degner et al. paper in the text on page 9. We have added an analysis of the overlap of our eQTLs with their dsQTLs – showing a strong enrichment over the null – but we feel that a full replication of their excellent analysis is outside the scope of this paper.