Dear Robin,

Thank you for taking the time and effort to comment on our manuscript preprint. We highly value your feedback and appreciate the opportunity to discuss this controversial topic openly. Please find our point-by-point responses to the queries and criticisms you raise below, in which we hope to have clarified some ambiguities in our original text and explain the analysis decisions that we have made. We will follow-up with a revised manuscript subsequently.

In their preprint, Young et al. draw two major conclusions:

1. That "*bidirectional transcription initiation from accessible chromatin is not sufficient for, nor specific to, enhancer activity*"

2. That the majority of eRNAs are likely not functional, based upon lack of evolutionary constraints and no evidence of purifying selection in eRNA exons

In this post I mainly focus on the first claim. The second claim very much agrees with our own previous analyses on evolutionary conservation of enhancer transcription units (TUs) and also with the fact that most eRNAs are rapidly degraded in the nucleus [2] [3]. (For an overview of how the fate of these RNAs are determined and how the cell deals with pervasive transcription see these recent reviews [4] [5].) Just as a minor comment I want to point out that the focus on multi-exonic eRNAs bias the analysis to a very small fraction of eRNAs (see our previous work on this [2] [3]) and therefore does not support general conclusions about eRNAs as a group. However, assuming functional eRNAs are more stable, the fraction of eRNAs with potential function likely belong to this minority group as splicing is strongly related to RNA stability [3].

We agree that our focus on multi-exonic eRNAs biases our sample to the more highly-expressed and relatively stable eRNAs. Of all eRNAs this subset may be reasonably considered to be the most likely to have biological function and with the relative stability of the transcripts the most likely for that function to be encoded by the processed transcript. However, this selection was also for the practical reason so that a comparison could be made between sequences contributing to the mature transcript (eRNA exons) and those processed from it (eRNA introns). As enhancers are themselves known to be evolutionarily conserved (Pennacchio et al. 2006) we set out to test whether this sequence constraint was concentrated within the transcribed regions of the enhancer and found that, in agreement with others (Marques et al. 2013), it wasn't. The important new insight this provides is that the selection was measured within the human population (using a derived allele frequency test). This is pertinent because it is known that functional regulatory elements can be quickly gained and lost through evolution (Meader, Ponting, and Lunter 2010; Villar et al. 2015; Young et al. 2015) and the absence of a signature of selective constraint between species could be explained away by invoking arguments of regulatory sequence turnover. This is directly addressed in our presented analysis where there is no evidence for purifying selection in the spliced eRNA transcripts within the human population whereas the same tests do readily detect purifying selection in other categories of processed transcripts.

The current preprint joins the debate on how to best predict active enhancers in mammalian genomes. During the last ten years the standard practice has mainly been to predict activity based on histone modifications and/or TF or co-activator binding (see my previous review [6]

for an overview of some approaches). Several recent studies (e.g. [2] and [7]) have, however, suggested that such predictions contain a large fraction of false positives. Despite this, the field is today highly dogmatized on what constitutes a proper enhancer. This is apparent in the manuscript by Young et al. in their trust in chromatin state predictions of types of regulatory regions and that these states represent functionally distinct units.

This was in fact the original purpose of our study – to test whether the traditional enhancer definition using histone modifications was outperformed by transcriptomic approaches (Andersson et al. 2014). By intersecting chromatin state maps with evidence for bidirectional and unidirectional transcription, we showed that the chromatin state was a better predictor of enhancer activity both in reporter assays and by *in vivo* correlations of enhancer transcription with that of putative target promoters. So rather than dogmatically accepting histone modifications, bidirectional transcription or DNase hypersensitivity as markers of enhancer activity we have tested their performance against measures of transcriptional output from a target promoter. We accept that there are limitations to both *in vitro* enhancer assays and assignment of enhancer-to-regulation-target by spatial proximity. However, both approaches are widely used (including in (Andersson et al. 2014)) and differ markedly in their perceived limitations, yet both approaches give us equivalent answers: chromatin state categorisation is a better predictor of enhancer activity than bidirectional (or unidirectional) transcription.

While core promoter regions are quite straightforward to identify with current accurate sequencing approaches, e.g. CAGE (5' end sequencing of cap-selected RNAs) and GRO-cap (5' end sequencing of cap-selected nascent RNAs), enhancer identification is harder. Our approach is based on the identification of loci with divergent transcription initiation producing balanced amount of eRNAs on both strands. This is challenging because eRNAs are often low abundant and rapidly degraded in the nucleus by the 3'-5' exoribonucleolytic exosome complex. I therefore welcome studies in which our approach is challenged and may point us in directions on how to improve. However, we first have to assess the reliability of methods and claims by Young et al.
Young et al. claim that bidirectional transcription is not specific to enhancers. I strongly agree with this claim and would even propose that the vast majority of human transcription initiation events are coupled with proximal upstream transcription initiation in the other direction (we have previously examined this property of human gene promoters [8]). So where do our results disagree? It boils down to whether non-gene promoter transcription initiation events are limited to enhancers or whether all open chromatin loci show transcription initiation.

Young et al, utilized CAGE, GRO-cap, GRO-seq, and PRO-seq data to determine transcription initiation levels and directionality at regulatory regions predicted by chromatin state segmentations. Predictions of stability of transcripts were inspired by a recent study by Core et al. [9]. In the present study, the authors state that absence of CAGE signal but presence of GRO-cap, PRO-seq or GRO-seq signal identifies unstable RNAs, while presence of CAGE and any of the other identifies stable RNAs. At comparable sequencing depths and comparable genomic positional distributions, the identification of unstable transcripts by presence of GRO-cap and absence of CAGE reads is rather elegant. In its original form, Core et al. compared TSS-focused reads (GRO-cap) with CAGE, while Young et al. also include TSS-unspecific reads (GRO-seq and PRO-seq). This may easily affect results when intragenic reads from the latter two sequencing methods (but not GRO-cap) overlap putative regulatory regions. For instance, CTCF-promoted RNA polymerase II pausing (e.g. [10]) in introns will likely generate local accumulation of non-TSS GRO-seq or PRO-seq reads. Secondly, presence of CAGE reads and reads from nascent RNA

sequencing approaches (even if TSS specific) does not necessarily mean that the RNA is stable. Note that the CAGE data used by Young et al. is the very same that was used by us to identify enhancers, whose produced RNAs are mostly unstable. I would recommend the authors to restrict their TSS analyses to GRO-cap and CAGE and for stability measures use data in which an RNA decay pathway has been perturbed (like the HeLa-S3 exosome knockdown data [3] produced by us).

These comments seem to have arisen from a lack of clarity in the methods section of our manuscript, for which we apologise and will amend in a revised version of this manuscript. Figure 1 shows a pattern of bidirectional, unstable transcription as measured by the absence of CAGE and scored only using GRO-cap reads, exactly as suggested above. This pattern is seen for all chromatin state annotations as CAGE in the top row (and Supplementary Figure 1). Furthermore, we see the same relationship between the level of accessible chromatin and the frequency of bidirectional transcription at essentially all chromatin state annotations using both CAGE alone (Supplementary Figure 3) and GRO-cap only in the absence of CAGE (Supplementary Figure 4). As stated above, the results for GRO-seq and PRO-seq show a less focused, but still apparent, bidirectional peak (Supplementary Figure 2). Importantly for our conclusions, this is still seen at CTCF and 'Other' marked DHSs and is not specific to enhancers.

While we agree that not all transcripts detected by CAGE will be stable, we still show that this and all other published technologies lack specificity in predicting enhancers as this bidirectional transcriptional pattern can be detected at all types of open DNA regardless of the state of that chromatin.

So, when their analysis shows that bidirectional transcription is not specific to chromatin state predicted enhancers and that bidirectional transcription is for instance present also at CTCF-bound loci I don't really trust the results. Transcription initiation at CTCF-bound regions could simply be due to intragenic PRO-seq or GRO-seq signal, i.e. not at real TSSs as mentioned above.

This comment again arises from a misunderstanding of our methods. As stated above, the GRO-seq and PRO-seq data is only used in the bidirectional plots shown in Supplementary Figure 2.

In addition, other studies (e.g. [11]) suggest that most CTCF sites are in fact untranscribed. Inaccurate classification of regulatory regions from chromatin data or the fact that the regulatory regions considered by Young et al. were defined in a hierarchical manner (i.e. enhancers could be predicted to be promoters as well) could also bias the results. If these predicted regulatory regions are really functionally different could we then really trust the results?

The hierarchical definitions were only considered for the correlation analysis presented in Figure 3 in which it was necessary to have consistency of annotation across cell types. Thus a region marked as enhancer in any of the cell types was considered in the enhancer category for correlation analysis. We similarly annotated promoter regions, then transcribed, then CTCF and then repressed regions. Bidirectionally transcribed regions shown in these same plots were calculated independently from chromatin state annotation. If the changing of chromatin states between cell types diminishes the correlation it would be specific to the chromatin states annotations but not the bidirectional transcription defined sites which were defined solely on the basis of transcription.

The discrepancy with the results of (Danko et al. 2015) regarding the transcription initiation proximal to CTCF binding sites is an interesting one. The approach of Danko was to identify transcriptional regulatory elements through a machine learning approach using GRO-seq or PRO-seq data. We agree that CTCF regions often show lower average rates of transcription initiation than other regions (evident in Figure 1, supplemental figures 1,2,3,4). However the pattern of transcription initiation where it is observed is the same as for other DHSs.

Furthermore, Young et al. only required a single CAGE read on each strand at a regulatory locus to call it bidirectionally transcribed. Although CAGE (and GRO-cap) is a quite excellent method, all sequencing techniques are error prone and studies using such data therefore need to carefully assess the signal with respect to sequencing library noise levels. A single CAGE tag is not sufficient to reliably distinguish true signal from noise. What is the expected sequencing noise level in the CAGE (and GRO-cap, GRO-seq, PRO-seq) libraries investigated?

Our approach to analysing the published set of reporter assays in K562 cells (Kwasnieski et al. 2014) and our own in HepG2 cells was to be conservative in detecting the absence of transcription, i.e. we wanted to be confident that our 'untranscribed' set convincingly had no evidence of being transcribed. We further filtered the transcribed set and, at enhancers, found that even a single read on one strand corresponded to increased enhancer activity relative to this untranscribed set.

None of the transcribed enhancer categories show significantly different activities from each other (Figure 2; Mann-Whitney, $p > 0.05$) but all showed significantly greater activity than the repressed set of reporter assays suggesting that, although there is inevitably noise in this dataset, it does not affect our way of approaching these categories. While we could consider only bidirectionally transcribed enhancers with multiple reads on each strand this would only function to limit sample size and does not alter our conclusions.

Young et al. further claim that bidirectional transcription is "*a by-product of an opening of chromatin at all types of regulatory regions*". This claim is based on two observations. First they observe that the fraction of DHSs with detected transcription is proportional to the strength of the DHS signal. Is this really unexpected? DNase-seq signal is in individual cells binary (indicating open or close) and the signal over a population of cells measures the fraction of cells in agreement, i.e. DNase-seq signal reflects sample cell population heterogeneity in chromatin openness: the fraction of cells in a sample with a locus having open chromatin and being transcribed will determine both the expression and DNase-seq signal of that locus in a measured cell population.

We agree with this interpretation, but we were particularly interested to note that this relationship was not observed only at enhancers, but could be detected at all types of chromatin. While it is certainly interesting to speculate as to the exact relationship between accessible chromatin and transcription, this does not contradict our central hypothesis in this paper that bidirectional transcription appears to be quantitatively related to the accessibility of all underlying chromatin states. Even selecting the most highly-transcribed regions would not enrich for genuine enhancers, but just the most accessible chromatin which can have multiple functions.

They later state that DHS signal shows clear discrimination between chromatin states in their correlation with (most proximal) genic transcription but that enhancer transcription level does

not. How can the DHS signal be strongly correlated with local transcription levels and with genic expression levels when its local transcription level is not correlated with genic expression level?

Some other apparent issues are

1) the now well-established fact that the most proximal promoter cannot be assumed to be regulated by an enhancer (ideally for such kind of an analysis chromatin interaction data is needed);

Even though it is an imperfect measure, we still expect to see some enrichment of genuine signal through using closest candidate-enhancer:promoter pairs. This approach appears to work well for DNase hypersensitivity measurements (Figure 3c) but less well for transcription (Figure 3b). An almost identical approach was used in earlier work: *"Uniquely, FANTOM5 CAGE allows for direct comparison between transcriptional activity of the enhancer and of putative target gene TSSs across a diverse set of human cells. Based on pairwise expression correlation, nearly half (40%) of the inferred TSS-associated enhancers (Methods) were linked with the nearest TSS, and 64% of enhancers have at least one correlated TSS within 500 kilobases"* (Andersson et al. 2014). We are currently re-processing the data to perform the corresponding 500 kb correlation analysis and will incorporate those results into a revised manuscript.

2) far from all DHSs are transcribed;

It is our contention that "the pervasive, low-level initiation of transcription associated with all categories of highly accessible chromatin represents a form of biological noise". In essence RNA polymerase is stochastically sampling the genome in a manner biased by the accessibility of the DNA and the availability of RNA polymerase in the local nuclear environment (see below). Active enhancers clearly contain DHS and are also expected to be located in or at least close to transcriptionally active regions, hence under this model it is not surprising that DHS in enhancers are initiating transcription. For other categories of DHS they are by definition accessible but may or may not be located in transcription factories, so under this model are expected to show corresponding variance in the rate of observed transcription initiation. However we clearly do find DHSs that exhibit stable or unstable bidirectional transcription which do not exhibit enhancer activity when tested (Figure 2) and consistent with our model of stochastic sampling, transcription initiation from all categories of chromatin state or regions defined purely on the basis of bidirectional transcription exhibit similar levels of correlation with the transcriptional activity of the nearest genic promoter (Figure 3b).

3) RNA polymerase II is not uniformly distributed in the nucleus (tend to co-localize in transcription factories together with multiple active regulatory elements).

We fully agree that RNA pol II is not uniformly distributed in the nucleus and already acknowledge this to some degree "DNA accessibility is not the sole determinant of transcription initiation, as nuclear position and the presence of specific transcription factor transactivation domains can dramatically influence transcription output [28]." However, considering your thought-provoking comments and our responses above this is an issue that we will further address in the discussion of an updated manuscript. Taking extremes we consider it likely that two equally accessible regions of DNA (as measured by DNase hypersensitivity) would differ in their rate of transcription initiation if one were physically

located within a transcription factory and one towards the nuclear periphery and lamin associated. Such nuclear localisation differences and sequestration away from RNA pol II may help explain why CTCF sites for example tend to exhibit low yet still detectable levels of transcription initiation relative to enhancers and why apparently transcriptionally silent enhancers exhibit less reporter activity than enhancers with detectable transcription initiation.

Furthermore, the order of events (transcription -> opening of DHSs) is not experimentally shown by Young et al., making this an unsupported claim.

We have not made any specific claim as to the order of events in the manuscript originally posted. However, the model we have advocated in this response implicitly assumes opening of DHSs → transcription (stochastically and very infrequently). We are not aware of any good temporarily resolved data or approaches to generate such data that will unambiguously resolve the ordering of these events. We also appreciate that there may be feedback between DNA accessibility and transcription initiation that could make any such clean separation impossible. In our next revision we will make clear that any such inferences are clearly identifiable as models and hypotheses rather than claims.

Now, let's get back to the burning point of disagreement: can enhancer RNAs be used to reliably predict enhancer activities? Young et al. state "*transcription initiation provides positive predictive value for accessible DNA but no power to discriminate enhancer from non-enhancer*". Since in vitro enhancer reporter assays can only investigate enhancer potential in a quite artificial manner (enhancer close to promoter), regions that positively validate in such assays can be inactive in vivo. Furthermore, a negative validation result doesn't necessarily mean that the region doesn't have enhancer potential, due to putative enhancer-promoter preference (but only one minimal promoter is tested for multiple candidate enhancers) and the high background observed at many minimal promoters, which could mask a weak enhancement and result in a false negative. Therefore, statistics (in particular sensitivity) of any approach will be hard to assess when based on in vitro reporter assays. It is also important to point out that candidate sequences need to be selected randomly to not bias the validation results. The practice of testing a list of top candidates will therefore likely favor the validation results, making a comparison of statistics between studies close to meaningless.

Enhancer reporter assays are a widely used technique for validating enhancer function (Noonan and McCallion 2010) and it is a little disingenuous to criticise this assay when it was used at the primary means to validate the initial approach of detecting enhancers with CAGE (Andersson et al. 2014). In our analyses we do not score results as validated or not, but instead make use of all experiments to study this quantitative measure of enhancer activity. As shown in Figure 2, these measures clearly segregate on the basis of chromatin state and not transcription.

Rather than test candidate sequences randomly, we considered the complete set (n = 1,499) of reporter experiments carried out in K562 cells (Kwasnieski et al. 2014). This data set is over an order of magnitude large than other published studies of enhancer discovery (Ernst et al. 2011; Andersson et al. 2014) and, as such, represents perhaps the best test that histone modifications can reliably detect enhancers. In HepG2 cells, we were limited by the number of bidirectionally-transcribed, repressor regions that could be successfully cloned and similar numbers were selected for the other groups tested. While lacking the strong

statistical support of the Kwasnieski et al. data set, these results again confirm the lack of bidirectional transcription alone in predicted DNA sequences which can act as enhancers.

A technical note on their own assays: if I understand the method section correctly the candidate enhancer is inserted just upstream of the minimal promoter, i.e. there is no polyA site in between that could hinder transcription originating from the candidate enhancer. How will this affect the results?

We did not add a polyA site to our assays as we wanted to closely mimic the approach of Kwasnieski et al., albeit on a much smaller scale. Regardless, the experiment is well controlled as regions with enhancer chromatin marks and with repressive (non-enhancer, non-transcribed, non-promoter) marks were treated identically in the enhancer assays, regardless of observed transcriptional activity, and show clear separation based on chromatin marks but not on transcription initiation (Figure 2b).

Despite the issues listed above, one can still test the potential of predicted positives (positive predicted value (PPR): the fraction of predicted positives that validates positively) with enhancer reporter assays. We tested this previously [2] on randomly selected candidate enhancers and found that, as Young et al. report, around 70% of candidate enhancers predicted from bidirectional transcription validated. In contrast, we found that only around 25% of non-transcribed enhancers predicted from chromatin data did. Hence, their finding that the fraction of chromatin-defined enhancers without detected transcription initiation show significantly lower reporter activity than transcribed candidate enhancers agrees with FANTOM.

They further find that bidirectional transcription also occurs at other regulatory regions that do not validate in reporter assays. This would indicate that bidirectional transcription as a marker of enhancers has a high false discovery rate (FDR). Can we trust that the bidirectional transcription initiation events considered in this study are OK (and not noise or intragenic non-TSS reads as discussed above)? I am a bit hesitant to trust the results mainly because of their use of GRO-seq and PRO-seq data to determine TSSs and our previous results [3] showing that repressive chromatin states tend to be untranscribed (quite a large fraction of repressed regions are in contrast found to be transcribed in the study by Young et al.).

As noted above, you were mistaken about our use of GRO-seq and PRO-seq though we accept that we could have made the description of the relevant analysis clearer. We agree that certainly rates of transcription (e.g. RNA-seq, GRO-seq) are lower in repressive chromatin regions and also that transcription initiation is also less frequent. However, it still shows the same characteristic pattern of bidirectional initiation near to DHS and transcribing away from the DHS (Figure 1, Supplementary Figures 1, 2, 3, 4).

I would recommend that the authors provided PPR and FDR statistics for the different approaches. Regardless of chromatin-state prediction, what are the statistics of predictions based on bidirectional transcription alone and according to varying expression cutoffs or signal to noise ratios? In addition, I would recommend that they reproduce their analyses using alternative massive reporter assay data (e.g. [7]). Will their claims still hold?

As stated above, we chose not to determine a threshold for ‘validated’ enhancer activity but instead considered the distribution of activities across all constructs tested. Our results are therefore robust to any arbitrary choice of activity required to ‘validate’ an enhancer. We

show that, on the much wider scale afforded by the Kwasnieski data set, chromatin states clearly separate DNA sequences that can behave more effectively as an enhancer while bidirectional transcription cannot.

With this approach, even only considering the stably transcribed regions using CAGE, we see a clear separation in reporter assay activity between those regions marked as enhancers and those marked as repressed. Our 'unstable' measures of transcription, which here did not require support from GRO-seq or PRO-seq data, , only confirm these results.

The paper suggested to confirm these results only examines the activity of regions within the enhancer chromatin state (Kheradpour et al. 2013). Crucially, this dataset cannot test our central conclusion that bidirectional transcription does not predict enhancers and will only confirm the results published by several groups that transcribed enhancers are more active untranscribed enhancers.

Thank you for taking the time and effort to help with the open peer review of this manuscript. We have undertaken to revise the manuscript as outlined above to address specific points you raised and we think doing so will improve the clarity and rigour of the manuscript. One of your main criticisms was the use of GRO-seq and PRO-seq which we think should have been addressed in full, primarily by us clarifying exactly when these data were and were not used. That eRNA transcripts are not selectively constrained is not contested by you and supported by the work of others as cited, though hopefully our responses have now clarified the choices we made for the analyses performed and highlight the importance of demonstrating the lack of sequence constraint within the human population. The main bone of contention is that you conclude in your previous work that bidirectional transcription marks active enhancers (Andersson et al. 2014), and we don't contest that. We do however extend those observations by showing that additional non-enhancer DHSs exhibit the same signature. In your paper you looked for bidirectional transcription at enhancers and found it. Many of them (~70%) validated in enhancer assays. What you did not do, was to ask if that same signature of bidirectional transcription occurred at other, clearly non-enhancer sites and this is where we extend your observations. While we detect the signature of bidirectional transcription to varying degrees at all classes of accessible chromatin. We consider this to extend the observations of (Andersson et al. 2014) rather than to invalidate them. As we have emphasised in our manuscript, it is highly beneficial to measure both open chromatin and gene expression in the same sample, and our results support the notion that many of the bidirectionally initiated regions in the open chromatin are likely to be active enhancers. However, we think it important to highlight the widespread existence of this transcription initiation signature outside enhancers such that the genomics community does not develop a blinkered view as to its specificity or function. We hope this work will provoke further investigation of our increasingly intricate understanding of the relationship between chromatin, transcription and gene regulation.

Robert Young and Martin Taylor
MRC Human Genetics Unit
MRC Institute of Genetics and Molecular Medicine
University of Edinburgh

References cited in response:

Andersson, Robin, Claudia Gebhard, Irene Miguel-Escalada, Ilka Hoof, Jette Bornholdt, Mette Boyd, Yun Chen, et al. 2014. “An Atlas of Active Enhancers across Human Cell Types and Tissues.” *Nature* 507 (7493): 455–61. doi:10.1038/nature12787.

Danko, Charles G., Stephanie L. Hyland, Leighton J. Core, Andre L. Martins, Colin T. Waters, Hyung Won Lee, Vivian G. Cheung, W. Lee Kraus, John T. Lis, and Adam Siepel. 2015. “Identification of Active Transcriptional Regulatory Elements from GRO-Seq Data.” *Nature Methods* 12 (5): 433–38. doi:10.1038/nmeth.3329.

Ernst, Jason, Pouya Kheradpour, Tarjei S. Mikkelsen, Noam Shoresh, Lucas D. Ward, Charles B. Epstein, Xiaolan Zhang, et al. 2011. “Mapping and Analysis of Chromatin State Dynamics in Nine Human Cell Types.” *Nature* 473 (7345): 43–49. doi:10.1038/nature09906.

Kheradpour, Pouya, Jason Ernst, Alexandre Melnikov, Peter Rogov, Li Wang, Xiaolan Zhang, Jessica Alston, Tarjei S. Mikkelsen, and Manolis Kellis. 2013. “Systematic Dissection of Regulatory Motifs in 2000 Predicted Human Enhancers Using a Massively Parallel Reporter Assay.” *Genome Research* 23 (5): 800–811. doi:10.1101/gr.144899.112.

Kwasnieski, Jamie C., Christopher Fiore, Hemangi G. Chaudhari, and Barak A. Cohen. 2014. “High-Throughput Functional Testing of ENCODE Segmentation Predictions.” *Genome Research* 24 (10): 1595–1602. doi:10.1101/gr.173518.114.

Marques, Ana C., Jim Hughes, Bryony Graham, Monika S. Kowalczyk, Doug R. Higgs, and Chris P. Ponting. 2013. “Chromatin Signatures at Transcriptional Start Sites Separate Two Equally Populated yet Distinct Classes of Intergenic Long Noncoding RNAs.” *Genome Biology* 14 (11): R131. doi:10.1186/gb-2013-14-11-r131.

Meader, Stephen, Chris P. Ponting, and Gerton Lunter. 2010. “Massive Turnover of Functional Sequence in Human and Other Mammalian Genomes.” *Genome Research* 20 (10): 1335–43. doi:10.1101/gr.108795.110.

Noonan, James P., and Andrew S. McCallion. 2010. “Genomics of Long-Range Regulatory Elements.” *Annual Review of Genomics and Human Genetics* 11: 1–23. doi:10.1146/annurev-genom-082509-141651.

Pennacchio, Len A., Nadav Ahituv, Alan M. Moses, Shyam Prabhakar, Marcelo A. Nobrega, Malak Shoukry, Simon Minovitsky, et al. 2006. “In Vivo Enhancer Analysis of Human Conserved Non-Coding Sequences.” *Nature* 444 (7118): 499–502. doi:10.1038/nature05295.

Villar, Diego, Camille Berthelot, Sarah Aldridge, Tim F. Rayner, Margus Lukk, Miguel Pignatelli, Thomas J. Park, et al. 2015. “Enhancer Evolution across 20 Mammalian Species.” *Cell* 160 (3): 554–66. doi:10.1016/j.cell.2015.01.006.

Young, Robert S., Yoshihide Hayashizaki, Robin Andersson, Albin Sandelin, Hideya Kawaji, Masayoshi Itoh, Timo Lassmann, et al. 2015. “The Frequent Evolutionary Birth and Death of Functional Promoters in Mouse and Human.” *Genome Research* 25 (10): 1546–57. doi:10.1101/gr.190546.115.