We thank the reviewers (**R1**, **R2**, **R3**, **R4**) for their thorough reading and constructive criticism. We will fix the mentioned typos and representation issues.

Curiosity about mapping points to group memberships for the two-color case and how the choice of the function affects the experiments (R1): This is indeed an interesting question which we thought about as well. We believe the experiments in Appendix F.1 (Figure 10) provide some information. There, the mapping function is varied from being completely random to deterministic. As the colors become more deterministically assigned, the price of fairness becomes larger.

Comparisons to previous fair clustering algorithms (R1,R4): Since previous fair clustering algorithms do not take probabilistic color assignments as an input, they are not directly applicable. However, as **R1** mentions, it is possible to assign points to their most likely color and then apply a fair clustering algorithm. Due to the page limit, this is done in Appendix F.1 (Figure 11). We also mention that it may not always be feasible to apply this method (see footnote in the appendix on page 22). We would be happy to run further suggested experiments of this kind and include a brief discussion of this point (and the preceding one from **R1**) in the main paper.

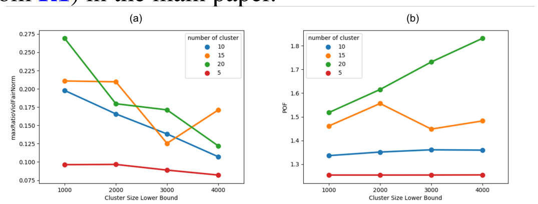

Figure 1: Results on the ***Census1990*** *dataset for different values of k. We see a reasonable degradation in the violation (a) and POF (b) for larger values of k.*

Large Cluster Assumption Issues (R1,R2,R3,R4): Reviewers have expressed concerns for the large cluster assumption & Algorithm 2. We acknowledge that the large cluster assumption is a limitation in our work. We consider our full solution to the two color case to be significant and this special case is commonly studied in the fair ML literature. Further, although the idea of sampling and then applying a fair clustering algorithm is simple, we show that arguably more elaborate methods do not seem to lead to better results under this assumption. Please see Appendix D where we present another solution based on dependent rounding for the k-center problem under the large cluster assumption and end up with similar guarantees. At the request of **R3**, we ran the large cluster experiment for different values of k; Fig. 1 shows a reasonable (and expected) degradation in quality. We will add these new experiments to the paper.

Motivation for specific case of metric membership (R1): One natural example is the case where the membership is income. Clustering could be performed and different outcomes assigned to different clusters. Forcing each cluster to have an average income around the global average prevents the possibility of having low or high income individuals from being over or under represented in a good or bad outcome. We will add this as motivation, and welcome other examples to include!

What's the current best possible for gamma (R2): As mentioned in Theorem 4.2, it is R for metric membership and 1 for the two-color probabilistic case as mentioned at the end of the theorem. We will be more clear about these values. Theorem 4.5 is for the large cluster multiple color setting; the lower and upper bounds are relaxed by ϵ, the factor of 2 in 2ϵ is a typo, it is in fact just ϵ.

Metric membership connection to the two-color case and a clear definition of the price of fairness (R2): Indeed we see that the reviewers agree on a more detailed explanation of metric membership, which we will add. As you mentioned (comment about Figure 3) we will indicate explicitly that the price of fairness is with respect to the output of the approximation algorithm not the true optimal since the true optimal is prohibitive to obtain (since the problem is NP-hard).

Presentation issues (R1): The suggested improvements regarding Appendix C, the figures, and the addition of proof sketches are all valid and would highly improve the presentation. We will try to incorporate them as much possible in the camera ready version within the given page limit.

Presentation issues and including further references (R4): We thank the reviewer for bringing up these issues and believe that fixing them in the camera ready version will significantly improve the presentation. We will add a high level description of our algorithms as well as a smoother introduction to metric membership and different notation. We were not aware of the paper by Ding and Xu, we will include it in the related work section.