We thank all the reviewers for their time, feedback, and insightful comments.

**Review #1.**

$\beta$-approximation: By $\beta$-approximation, we mean an algorithm that approximates the optimum within a multiplicative factor of $\beta$. We will add a formal definition in the revision.

Comparison to [14]: This manuscript appeared on arXiv contemporaneously with our submission to NeurIPS.

Compared to our results, the main takeaways from [14] are: (i) their algorithms are heuristic and therefore do not provably optimize for any objective such as value or revenue, and (ii) their algorithms are not provably fair under our fairness constraints. Thus, unlike ours, their work does not seem to offer any formal guarantees.

For an experimental comparison, we tried the following two routes during the rebuttal process:

(i) We reached out to the authors for a pointer to their code. The authors responded that their code is currently unavailable as it is being revised based on the feedback they received.

(ii) We implemented their algorithms from scratch. During this process, we realized that replicating their results is difficult due to the lack of a completely specified procedure for choosing the hyperparameters of their algorithm ($\alpha_0, \beta_0, \theta_1, \theta_2$). For these reasons, to ensure a fair comparison with their work, we postpone the experimental evaluation of their algorithm pending the release of the new implementation by authors.

**Review #2.**

Thank you for the positive comments. We will address your suggestions and include the additional citation in the revision.

**Review #3.**

Apologies for the typos and thank you for the careful reading.

Kleindessner et al. designed an algorithm for k-center with different type of fairness requirement. Instead of balancing different colors in each cluster, the goal is to pick centers (proportionally) from different colors. It is basically k-center under partition matroid.

Thank you for articulating this clearly; we will add this in the revision.

Both the bounded representation and the representations from Bera et al. and Bercea et al. generalize the problem of "balance". How do these different representations compare?

Bera et al. and Bercea et al. both introduce a generalized version of the constraint from [Ahmadian et al. 2019]. For each color, they specify the upper and lower bounds for their fractional representation in each cluster. Our results continue to hold even for this general constraints, as long as we can find a fairlet decomposition satisfying the conditions of Theorems 8 and 10. We will add this remark to the revision.

(Theorem 8): what is an upper bound on the value of $m_f$ in terms of $\alpha$ (for general values of $\alpha$)?

For two colors with general $\alpha = r/(b+r)$, where $r \geq b$, we have $m_f \leq b + r$ [Chierichetti et al. 2019]. For multiple colors with $\alpha = 1/t$, we have $m_f \leq 2t - 1$. Note that the bound for multiple colors for general alpha is an open question.

Algorithm 1 – step 2: how to find an initial solution?

For two colors with general $\alpha = r/(b+r)$, we use the fairlet decomposition method proposed in [Chierichetti et al. 2019]. For multiple colors with $\alpha = 1/t$, we use the method proposed in Lemma 24 in the Supplementary Material.

In the revision, will also add further justifications for studying fairness in a hierarchical clustering setting.