We would like to thank all reviewers for their thoughtful comments. Below we address their concerns.

For Reviewer 2: First, we address the major comment of Reviewer 2 regarding where ϵ is used in the two algorithms. **Indeed, L', which was defined in Line 149 of page 4, depends on ϵ.** So both algorithms depend on ϵ. We are sorry for the confusion. In the final version, we shall define L' in the pseudo-codes to emphasize the dependence on ϵ.

To give more details, L' is the lowest level ℓ in T satisfying $\lambda^{\ell} \geq \epsilon f$. Note that in both algorithms, we are marking all vertices on level L' and this means there cannot be any minimal-marked vertices above level L'. Thus, the open facilities can only be at level L' or at lower levels. In the analysis of the total facility cost obtained in Algorithm 1, which is the proof of lemma 1, we obtain an upper bound $\frac{1}{\epsilon} \cdot$ opt for the cost of open facilities which are at level L', using $f \leq \frac{1}{\epsilon}\lambda^{L'}$. Combining with the facilities at lower levels, we upper bound the total facility cost with $(1 + \frac{1}{\epsilon})$opt. Similarly, in the proof of lemma 4, in finding an upper bound for the total expected facility cost of Algorithm 2, we are using $\lambda^{L'} \geq \epsilon f$ property of L'. This is how the privacy parameter ϵ comes out in the analysis of our algorithms.

Explain in the related work section why previous methods cannot be used.

As discussed in our paper, the techniques of Gupta et al. [8] can give an ϵ-DP algorithm for UFL under our model, but with much worse approximation ratio. We have many new ingredients in our paper that led to the improved approximation ratio.

What are “LDP algorithms" (see beginning of Section 3)

This is a typo. LDP means "locally differentially private". We derived an LDP algorithm for the problem, but did not include it in the paper.

We will fix typos and add additional explanations in the final version. Thanks for pointing them out.

For Reviewer 3: Why frame this problem as the super output setting instead of framing it as an allocation problem subject to Joint Differential Privacy (JDP), since that is how the solution is framed.

We used the super-set output setting because this describe more precisely how the solution is presented. In our algorithm, we are not revealing each user the location to which they connect, instead we are revealing a super-set of the open facilities and each client connects to the closest facility in the set. Of course this is very related to JDP as discussed in Section A of the supplementary material. But using the super-set is just one way to guarantee the JDP property. We will discuss our terminology more thoroughly in the final version.