We thank the reviewers for the thorough reviews and helpful comments.

- Concerning the generalization to a larger number of communities (**Rev1, 2, 4**):

As common in the majority of research on the SBM, we decided to focus on $k = 2$ as it encapsulates the core-challenges and the main technical contributions. We expect the overall results to generalize to the case of $k > 2$. In the course of the work leading to the current paper, we formally proved the following theorem statement for $k > 2$ for a "parallel" variant of Louvain where all swaps are done "in parallel" (moving at each step every vertex to its best community, disregarding the choices of other vertices): *Let* $\varepsilon > 0$*, not necessarily constant. Consider a graph* $G \sim SBM(k, n, p, q)$ *with* $k = O(1)$*. Then, with probability at least* $2/3$*, Parallel Louvain recovers the partition* $\{V_1, V_2, \ldots, V_k\}$ *after* $O(\varepsilon^{-1})$ *rounds, if* $p - q \geq \frac{\log^5 n}{n^{1/4-\varepsilon}}$. We expect this statement to extend naturally to standard (sequential) Louvain.

We did not attempt to include such a statement in the submission due to readability concerns (the proofs for $k > 2$ do become longer, but not much more interesting) and due to space limits. Indeed, the purpose of our paper is to prove the first theoretical guarantees on a non-trivial setting for the Louvain algorithm, together with understanding its limitations. We believe that understanding the behavior of Louvain on simple, yet non-trivial, clusterable graphs is a first step towards getting better algorithms. Hence, in our opinion focusing on two communities helps make our argument clearly.

Given the Reviewers' interest, we propose to add further comments on the case of $k > 2$ to the closing remarks of our revised paper.

- Concerning the experimental section (**Rev1**):

We would like to stress that the goal of the paper is to explain the success of the Louvain algorithm – which is widely used, in particular due to its efficiency and easiness to implement – and expose the strength and weaknesses of the modularity objective function. It is well known that semi-definite programming approaches achieve optimal recovery bound in the stochastic block model, and would naturally outperform Louvain in terms of recovery. However, the running time is significantly worse rendering them impractical for even medium-sized graphs. Moreover, spectral approaches are known to be very brittle and often fail in practice when noise is present, and in general we believe Louvain is much more robust. However, before considering such sophisticated clustering scenarios, it is benefitial to understand the standard case (i.e., the stochastic model) and it turns out that this already presents a plethora of challenges that need powerful machinery to be overcome.

Our goal was not to provide an optimal algorithm for the stochastic block model, but to establish the first provable bounds for Louvain, taking one step at a time and laying the groundwork for subsequent extensions.

- Concerning the presentation (**Rev1, 2**):

We acknowledge the typos and presentation comments, and we will address them. Note that $O(n^2 p)$ does indeed signify linear running time (our running time is linear w.r.t. to the input size, which is linear in $n^2 p$).

- Concerning the warm start (**Rev3**):

The point of Thm 1.1 is to show that a 'good' initialization (with high Delta), obtained using e.g. some other algorithm or additional knowledge, allows Louvain to converge to the ground-truth community in a near-optimal regime of $p - q$. We agree that for a random initialization, Thm 1.1 is not useful, in which case we appeal to Thm 1.2 which handles the random initialization scenario but for a stronger assumption on $p - q$. We will make this explicit in the paper.

- Concerning balanceness condition (**Rev4**):

In fact, we actually prove that the balancing step is not necessary. This is done in suppl. mat., Section F: "Allowing Size Imbalances". Note that we did not include this (important) improvement in the main proof for readability (since it seems already fairly sophisticated) and in view of the page limit. In the final version, we'll make sure to emphasize that we do not need the balancing.