We thank the reviewers for their time, valuable feedback, and recommendations for improving the manuscript. All the reviewers seem to agree that our contributions are valid and interesting. *In support of this assessment, we would like to reiterate that IPR is a completely novel PR method for seed-set expansion that defies and disproves the validity of common methods that use decreasing weights for landing probabilities [20].* In our subsequent response, we focus on further highlighting the differences between IPR and PPR (**Rev1**) and IPR and spectral clustering (**Rev2**). We also discuss a condition in our theoretical results questioned by **Rev2**, and address **Rev3**'s concern about future work.

***IPR vs PPR with parameter 0.99. (Rev1)** The reviewer's intuition that PPR with a parameter close to 1 has a performance similar to IPR is correct. However, this special case does not imply that IPR is merely a simple modification of PPR. **Rev1** seemed to overlook the key new insight motivating the IPR method, demonstrated both by new theoretical results and experiments described in the manuscript: The discriminative power of large-step LPs does not decrease or decrease as fast as previously expected based on a mean-field analysis alone [20]. Our finding has truly far-reaching consequences as it shows that near-optimal seed-expansion requires choosing geometrically increasing rather than geometrically decreasing weights; similar results may be derived for recommender systems/link prediction. Hence, IPR and PPR lead to fundamentally different implications instead of merely experimental performance.

As requested by the reviewer, we evaluated the performance of PPR with parameter 0.99 and summarized the results in the figures/table below. The experimental condition are the same as described in the manuscript. The performance gap between PPR 0.99 and IPR is smaller than the gap between PPR 0.95 and IPR. However, the gap is still significant (more than one standard deviation) for real world networks and even significantly more so for SBMs.

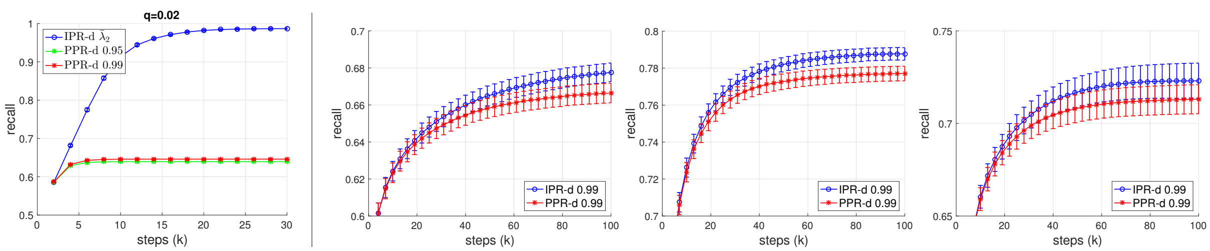

Figure 1: (Left): Recalls (mean $\pm$ std) for different PRs over SBMs with parameters $(500, 0.05, 500, 0.05, q)$, $q = 0.02$; (Right): Recalls (mean $\pm$ std) of different PRs over the Citeseer, Cora and PubMed networks (from left to right).

***Why not use spectral clustering? (Rev2)** Spectral clustering and seed-expansion approaches have fundamentally different objectives. Spectral clustering is used to find *all* communities in networks and is therefore a global clustering algorithm that does not scale for large networks. Seed-expansion approaches are local, and their complexity is dictated by the size of the community we are interested in [25]. Furthermore, in its classic form, spectral clustering partitions network nodes and may not be used to obtain overlapping communities; seed-expansion approaches can easily detect overlapping communities, but to keep our exposition focused we only analyzed the non-overlapping setting also pursued in [20, 24]. In conclusion, spectral clustering is not expected to scale and perform well on networks with overlapping communities such as the Amazon and DBLP networks. All forms of community detection, including classical spectral clustering and PR methods, require knowledge of some input parameters. In the former case, one needs to know the number of communities, while for PRs one needs to select the parameters α (for PPR), h (for HPR) or θ (for IPR). For IPRs, the choice of θ (and not the parameter q that is only used in the SBM) also allows for adapting IPRs to different networks and different tasks and is in general easy to estimate. Even when the parameter is estimated imprecisely, it does not influence the bulk performance gain of IPR compared to PPR, as the crucial point is that the parameter is used to control the increase (rather than decrease) in the weights.

Step size k	5	10	15	20	5	10	15	20
	Amazon (std: ± 0.12)				DBLP (std: ± 0.09)			
IPR0.99	46.63	48.03	**48.43**	**48.53**	27.58	28.78	**29.18**	**29.27**
IPR0.90	46.67	48.08	**48.45**	**48.53**	27.64	**29.14**	**29.26**	**29.32**
PPR0.95	46.57	47.92	48.30	48.43	27.46	28.49	28.90	29.06
PPR0.99	46.59	47.94	48.34	48.45	27.51	28.58	29.00	29.14

***Conditions used to establish the theoretical results. (Rev2)** Please note that our results already improved the condition $d_{\max}/d_{\min} = \Theta(1)$ in Avrachenkov et al. [23,24] to $\log n * d_{\max} = o(d_{\min}^2)$. This improvement essentially allows for much larger heterogeneity of degrees. A significant contribution of our work is the first known characterization of the variance of LPs under assumptions weaker than any other previously reported ones.

***Future directions and improvement. (Rev3)** In the Supplement, we listed four future research directions regarding how to further improve the GPR framework for seed-expansion community detection, especially in settings for which IPR may not be optimal. It would be of interest to characterize the correlation between LPs of different steps, as the correlation between k-step and $k+1$-step LPs increases with k. Correlations may help in identifying the optimal number of steps of LPs to accumulate. Moreover, our current theoretical analysis requires communities to be non-overlapping (also used in [20, 24]). Overlapping as well as *sparse* community GPR methods are other interesting new directions.

***Minor issues.** In the revision we will clarify the inequality on line 134 (**Rev1**), enlarge the fonts in the figures (**Rev1**), provide a more detailed algorithm in the Supplement (**Rev2**). Regarding aggregating results of different communities (**Rev2**), if we understand correctly, this is what we did in the experiments.