Reply to Reviewer #1

– *Practical motivation of our problem and comparison with [25]:* We would first like to remind that campaigns being initiated simultaneously by a social-network owner is a standard assumption in the literature on influence maximization with multiple campaigns [1,12,20,29]. We believe that this assumption is even more valid in the current era, with social-media companies being under public and legal scrutiny, requiring them to make more informed and transparent decisions. Still, we acknowledge that such assumption could reduce the generality of the problem; allowing initial seeds, selected at the prior rounds of the campaigning process, would constitute an interesting future work. Such an extension, however, would require to study the problem in an online and adaptive manner (along the axis of exploration vs. exploitation) while [25] simplifies the problem to a great extent and assumes that four different seed sets are all initiated in parallel. We should note that [25] maximizes a different objective, which we believe is somewhat artificial to reduce polarization, i.e., it also accounts for nodes *not being informed* by any of the campaigns, while co-exposure maximization is a natural objective, solely accounting for nodes that are informed by both sides of a controversial issue. Besides, among the three algorithms that [25] proposes, only one provides an approximation guarantee. Finally, their algorithms rely on choosing common seeds, hence, are not applicable when the seeds sets are required to be disjoint.

– *Conceptual differences between RRP set and RC set in [3]:* Both [3] and our algorithm TCEM extend the framework of IMM [38]. Therefore, there are high-level similarities. However, since [3] and TCEM solve different problems, the way these algorithms exploit reverse-reachability and their sample-size requirements differ greatly from each other. Notice that the domain where the random sets are sampled from is dictated by the objective function: [3] operates by sampling RC sets, i.e., random sets defined over user-item pairs, while TCEM samples RRP sets, i.e., random sets defined over user-user pairs. As these algorithms sample from different domains, the sample-complexity results and the sample of random sets obtained for one problem cannot be used to solve the other.

To further elaborate the differences of our work from [3], their problem assumes as input political leanings of nodes and items, quantified in the interval $[-1, 1]$, and aims to maximize the sum, over all nodes, of the range of the exposed political leanings including the node's own. Therefore, users who are exposed to only one item that has a different leaning from theirs, still contribute to the value of their objective. Translating this to our setting by assuming two items of leaning -1 and 1, implies that the objective of [3] can potentially achieve a relatively high value while the co-exposure being 0. Thus, the work in [3] does not guarantee that co-exposure is maximized. We thank the reviewer for their constructive feedback; we will clarify these points in the paper.

Reply to Reviewer #2

– *Weak approximation guarantee:* We acknowledge that obtaining a tighter guarantee is a highly interesting future work. We remind that our problem is not (bi-)submodular and has a submodularity ratio of 0, rendering the recent advances in monotone non-submodular function maximization non-applicable. Thus, we believe that our current result is theoretically interesting.

– *Round-robin greedy as a baseline:* We thank the reviewer for this suggestion and will consider it.

– *Bisubmodular or k-submodular functions:* our objective function is not bisubmodular. A simple counter-example is omitted due to space limitations. We will add this discussion in the paper.

Reply to Reviewer #3

– *Missing references and a pointer to such a reference:* The mentioned paper studies the problem of limiting misinformation, thus, is not related to the problem studied in our paper.

– *Real-world motivating scenarios for our work:* We would like to point out that competing campaigns being initiated by a social-network owner is a standard scenario in the viral-marketing literature [1,12,20,29] since it follows the real-world business model of social-network owners, such as, Facebook and Twitter, which provide social advertising service to marketers.

– *Disjoint seed sets:* Disjointness constraint is a widely adopted and credible design choice as an opinion leader would not be promoting two sides of a controversial issue, such as, gun control. We refer the reviewer to a WSDM2018 tutorial "Influence maximization in online social networks," and references therein.

– *Propagation model:* The independent-cascade (IC) model [28], is a widely-adopted information-propagation model in the literature, e.g., see the WSDM2018 tutorial. Please note also that campaigns operate on different IC model instances.

Reply to Reviewer #4

– *Scalability:* We remind that we are dealing with a more stringent estimation task than of the influence-maximization problem. This naturally translates to increased sample complexity, hence, less scalability compared to IMM [25]. In practise, our algorithm could scale up by a parallel implementation since RRP sets can be sampled independently.