We thank the reviewers for their helpful comments. We will address all suggested minor revisions, clarifications, corrections and typos. Due to space limitations, we comment here on major questions and suggestions.

Reviewer 1: We thank you for appreciating the balance between rigor and intuition.

• "why pure Nash equilibria, better response dynamics..." We employ the one-shot model as an approximation of the (much more) complex prevailing recommendation systems (RSs). Better response dynamics is the common and arguably the least demanding way to deal with any asynchronous dynamics in multiagent systems. Indeed, this is the standard practice in economic studies as well as in engineering contexts (e.g., communication and routing in networks). Of course, dealing with particular learning rules is of great interest but our emphasis when looking at BRD is on the set of *all* plausible behaviors. Furthermore, bloggers and vloggers focus on one or two niches and do not mix, which additionally supports our focus on pure fixed points versus mixed strategies (which are common in CE). Given this comment, we will address it better in the paper. Nevertheless, minimizing regret is a great tool to obtain CE and CCE, which also capture some form of centrality; we will propose it as future work, thanks.

• "can the pure Nash equilibrium be reached using best response dynamics?" Recall that a best response dynamic consists of a starting profile and a series of best-response improvements. Our algorithm can be seen as a best-response dynamic in which all players start from a *null profile*, assigning all players to a factitious topic with zero user mass, and then determines the order of best-responding. We will add that to the paper, thanks.

Reviewer 2: We thank you for appreciating our theoretical results.

• "when there are no ties in quality parameters, the game is essentially a stable marriage game." This is true (and is mentioned in prior work). We will write this fact explicitly in our paper, thanks.

• "The problem is not very well motivated." Our conceptual takeaway concerns any RS with strategic content providers. We claim that the system is not just about the short-sighted recommendation, but rather has a mediator role; hence, the system can (and ideally should) solve some of the market's inefficiencies. We agree that our theoretical results cannot be applied as-is (as we have many limiting assumptions). Still, we hope the ideas presented in this paper and our theoretical grounding will be useful to practitioners.

• "The marginal contribution over the previous work (e.g., AAAI'19 paper) is not that substantial." We respectfully disagree. We deal with the rate of convergence (Theorem 2) and efficient equilibrium computation (Theorem 3), none of which are even hinted in that AAAI'19 paper. Theorem 1, which we consider a minor contribution of this paper, is the only result that builds on techniques from prior work (yet significantly extends it). We stress that our techniques for proving these two theorems are genuine, and did not appear in prior work.

• "more discussion of the connection to congestion games and the deferred acceptance algorithm for stable marriage." If accepted, the extra page will allow for more elaboration on these works.

Reviewer 3: Thanks for seeing our attempt to justify the significance of the mediator role of RSs in a positive light.

• "BRD is not popular algorithm to find NE..." We agree with the reviewer. We analyze BRD as a means of decentralized computation, where the RS does nothing but matching demand with supply. In the presence of a centralized planner, the system can be much more efficient. This is one of the main claims of our paper.

• "social welfare only w.r.t. to consumers but not including content providers' utility" We entirely agree. Note that measuring the overall welfare of RSs with multiple stakeholders is a philosophical and technical challenge that goes way beyond this paper (for instance, see *Multisided Fairness for Recommendation* by Robin Burke, 2017). Our attempt is indeed to use the "Price of Correlation," which, if accepted, we intend to move to the main body based on this suggestion. We will also highlight that we see "comparing the social welfare under decentralized and centralized algorithms" as a promising direction for future work, thanks.

Reviewer 4: We thank you for appreciating our attempt to incorporate profit maximization into RS design.

• "social welfare "disappears" in the paper." As this reviewer mentions, the holy grail of multiple stakeholder RSs is to balance the utilities of producers and the welfare of consumers. However, coming up with the right mix is a longstanding challenge in RSs (see *Multisided Fairness for Recommendation* by Robin Burke, 2017) and is way beyond the scope of our paper. However, our results demonstrate that by steering the market to the 'right' equilibrium, future RSs could increase both utilities. As this reviewer mentions, our Assumption 1 sheds light on the connection between utility (producers) and welfare (consumers); we will better emphasize it in the paper, thanks.

• "the paper brings as motivation the study of long-term social welfare in these dynamic systems." We strongly agree. The main conceptual takeaway is that welfare-driven recommendations (instead of the current, short-sighted approach) can steer the market to better equilibrium points. However, this is just the tip of the iceberg; as this reviewer and the other reviewers mention, long-term welfare requires reaching the best equilibrium (the main technical open question of this paper), as well as other forms of system design. Due to this comment, we will extend the discussion on that, thanks.