Figures for Author Rebuttal

Interpolating Item and User Fairness in Multi-Sided Recommendations

Additional Experiments on MovieLens Data

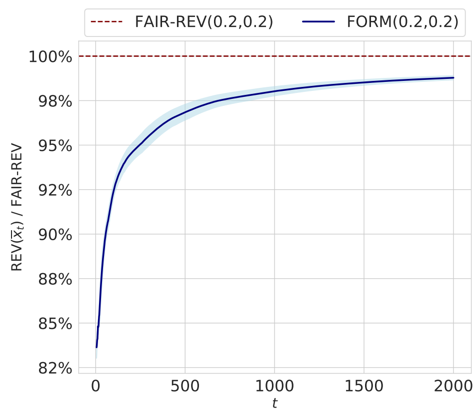

(a) Convergence of revenue.

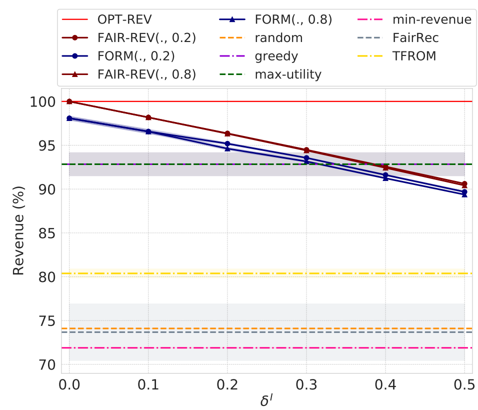

(b) Normalized revenue.

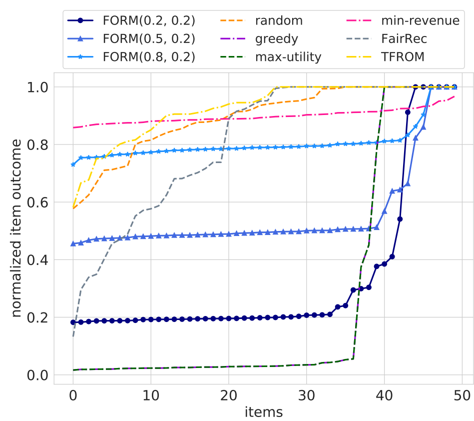

(c) Item outcomes.

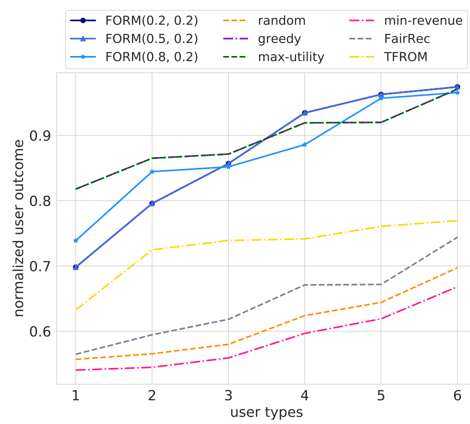

(d) User outcomes.

Figure 6: **Additional experiment results on MovieLens data.** Here, we act as a movie recommendation platform that shows a "trending action movie" to any arriving user. We considered the $N = 50$ action movies from MovieLens (ML-100K) data with the highest number of ratings and clustered users into 6 types based on their preferences. As the movies are not associated with revenues, we let $r_i = 1$ for all $i \in [N]$. Here, the platform's main objective is to maximize its expected marketshare. The item-fair solution adopts maxmin fairness w.r.t. each movie's marketshare, with user utilities captured by the MNL model. We consider a total of $200,000$ user arrivals, and solves Problem (FAIR-RELAX$(\hat{\boldsymbol{\theta}}_t, \eta_t)$) upon every 1000 arrivals.

The results are consistent with those in our Amazon review data case study (Section 4). Note that in a movie recommendation setting with homogeneous revenues, the interests of the platform and the users completely align. This explains why the curves of *greedy* and *max-utility* completely overlaps with each other in our figures. However, *greedy* still suffers from 7-8% loss in marketshare (Figure 6b)), which is precisely because inadequate exploration of user data makes it overlook potentially more popular items and stick with a sub-optimal item. Overall, our algorithm `FORM` adeptly balances the interests of both the platform and its stakeholders, while handling the tradeoff between learning and fair recommendation.