



Individual Fairness for Social Media Influencers

Stefania Ionescu^(✉), Nicolò Pagan, and Anikó Hannák

Social Computing Group, University of Zürich, Zürich 8050, Switzerland
ionescu@ifi.uzh.ch, pagan@ifi.uzh.ch, hannak@ifi.uzh.ch

Abstract. Nowadays, many social media platforms are centered around content creators (CC). On these platforms, the tie formation process depends on two factors: (a) the exposure of users to CCs (decided by, e.g., a recommender system), and (b) the following decision-making process of users. Recent research studies underlined the importance of content quality by showing that under exploratory recommendation strategies, the network eventually converges to a state where the higher the quality of the CC, the higher their expected number of followers. In this paper, we extend prior work by (a) looking beyond averages to assess the fairness of the process and (b) investigating the importance of exploratory recommendations for achieving fair outcomes. Using an analytical approach, we show that non-exploratory recommendations converge fast but usually lead to unfair outcomes. Moreover, even with exploration, we are only guaranteed fair outcomes for the highest (and lowest) quality CCs.

Keywords: Social network formation · Individual fairness · Markov chains

1 Introduction

The past couple of decades brought a steep increase in the impact social media platforms have on our lives, e.g., by shaping the information we receive [1] and the opinions we form [2]. During this time, platforms previously designed to only connect real-life friends slowly encouraged users to follow strangers based on their content. Today, platforms such as YouTube, Twitter, Instagram, and TikTok are heavily centred around User Generated Content (UGC) and use recommender systems (RSs) to facilitate the exploration of content. In response to this change, some users specialize in creating even semi-professional content that can attract more and more followers to the point that they can make revenue based on their audience. Similarly to the labour market [3], it is then natural to expect that these online platforms guarantee fairness for the content creators (CCs) in a way that equally good CCs should be rewarded similarly in terms of visibility and audience, and ultimately of income. Since the network formation process is often heavily mediated by the RSs, it is thus appropriate to ask whether they produce fair outcomes for the CCs.

To first understand the structure and the properties of these social media platforms, [4] proposed a simple model in which (a) each CC has an intrinsic and objective quality, (b) in each round, users receive a recommendation for a CC (which could be drawn from the uniform distribution, or from a preferential attachment (PA) process), and (c) users follow the recommended CC if this CC has a higher quality than all the user's current followees. By simulating this model, Pagan et al. [4] showed that the expected number of followers (in-degree) of the CCs follows a Zipf's law [5]. In particular, the expected rankings of CCs given by their quality and by the number of followers are the same.

While the findings of [4] suggest UGC-centered platforms are fair for the CCs, some empirical evidence suggests differently. In cultural markets, for example, it is often difficult to anticipate success [6]. The experimental study of [7] revealed this is at least partially caused by the social influence. More precisely, as the amount of information users receive about the prior choices of others increases, the predictability of the items' popularity decreases. This unpredictability can be interpreted as a lack of fairness guarantees on some UGC-centered platforms.

We believe this apparent incongruity in previous work is mainly due to two model limitations of [4]. First, the analysis is restricted to two exploratory recommendation processes (i.e., processes which take risky actions in order to uncover better options and are thus limited in the amount of social influence they exert). However, in practice, not all RSs are as such. Recent literature in RSs argued why and how we should encourage diversity by finding the right balance between exploration and exploitation [8–11]. In practice, even pairs of popular items might not be jointly accessible to users, i.e., if a user is recommended and follows one of the two, they will not be recommended the other [12]. This puts real-world RSs in stark contrast with PA and uniform random (UR) recommendations where every user can be recommended any CC. Second, the authors only focus on the *expected* number of followers *at convergence* (i.e., after all users were eventually recommended the best CC). However, such ex-ante fairness does not imply ex-post fairness (i.e., even if in expectation CCs receive a number of followers proportional to their quality, many of the actual outcomes that could materialize are unfair) [13]. Moreover, as also noted by the authors [4], there might be long times to convergence. This means that even if a fair outcome would eventually be reached, this might not happen within reasonable time.

This paper aims to address these two challenges by bridging between the network science, RS, and algorithmic fairness communities. More precisely, the current work (a) defines both ex-ante and ex-post fairness metrics for CCs, (b) extends the model with extreme PA (a non-exploratory RS, inspired by the popularity RS[14], which only recommends the most followed CCs), (c) uses Markov Chains to theoretically study the network formation process and its fairness under extreme PA, and (d) compares the network formation processes and their fairness under extreme PA with the ones under PA and UR recommendations (either by referring to prior work, or by novel analytical and numerical results).

2 Related Work

Networks. After the pioneering work on the random graph model [15], the complex networks community started developing simple yet insightful mechanisms that explain the formation of social networks (e.g., the small-world network model [16], and the preferential attachment model -PA- [17]). In PA, newborns form connections to existing nodes with a probability proportional to the degree. This rich-get-richer phenomenon successfully reproduces the idea that popular users experience higher visibility, which in turns brings them more popularity. On the other hand, it gives little emphasis on the socio-economic microscopic foundations that explain why individuals make certain connections. Focusing on this alternative approach, one line of research in sociology (Stochastic Actor Oriented Models [18]) and one in economics (strategic network formation models [19]) take an utilitarian perspective: agents build their connections to maximize some benefit e.g., their network centrality. The quality-based model of Pagan et al. [4] combines these approaches by using a UR or PA-based RSs and a utilitarian following decision-making function for users. To enhance our understanding on the coupling between RS and human network behavior, we add a non-exploratory RS and investigate the fairness of the resulting outcomes.

Fairness. Researchers are not only concerned with the average performances of processes, but also with the equity of these processes in impacting individuals. This is reflected by the extensive work on developing fairness measures as well as a methodology to choose the most suitable one depending on the application domain [20–22]. From these various fairness metrics, we focus on *individual fairness* which assesses the degree similarly qualified individuals receive similarly quality outcomes (see [23] for an overview of its importance and apparent incompatibility with other fairness metrics). An important phenomenon is the *timing effect*, according to which it is not enough to specify a welfare function but also when this should be measured (ex-ante or ex-post) [13]. Building on these, we define and investigate both the ex-ante and ex-post individual fairness for CCs.

Recommender Systems. Recently, the RS community argued for the importance of looking beyond accuracy [8] in RS-evaluation. Design-wise, there is an ongoing struggle to develop diverse RSs [9–11] which, moreover, ensure that any two items could be recommended jointly to users [12]. This perhaps explains why the popularity-based algorithms implemented within the RS community [14, 24] differ from PA [4, 17] by not allowing for a complete exploration of alternatives. Our motivation to analyse extreme PA is grounded in this prior work, although our goal is to understand the network-formation process and its fairness. We also depart from the experimental study of [7] as we (a) formally define fairness metrics, (b) use theoretical tools to understand *why* social influence decreases fairness, and (c) also look at network-specific metrics (the time to convergence).

3 Model

As mentioned in the previous two sections, our work is based on the quality-based model proposed by [4] and briefly presented below. As any model, this is based on some simplifying assumptions, for which we refer the user back to the original paper. Note that our notation occasionally differs from [4] in order to allow for future extensions. The next two subsections contain our advancements of the model: (a) a RS-based formulation of the network formation process (plus the definition of the new RS, extreme PA), and (b) the fairness definitions.

We assume users can be partitioned into $n \geq 2$ content creators (CCs) and m (non-CC) users, with $m \gg n$. We refer to CCs as CC_1, CC_2, \dots, CC_n and assume they are ordered strictly by an existing objective quality. That is, CC_1 is the absolute best CC, CC_2 the next best CC, and so on. Every user can follow any content creator thus leading to a bipartite, unweighted, and directed network.¹ We represent this network by its adjacency matrix $A \in \{0, 1\}^{m \times n}$, such that (s.t.) $a_{u,i}$ is 1 if user u follows CC_i , and 0 otherwise.

The network formation is a sequential dynamic process where (a) the network is initially empty, and (b) at each timestep, each user is recommended a CC which they can follow or not. Formally, let A^t capture the state of the network at time t . By assumption (a), $a_{u,i}^0 = 0$ for each user $u \in \bar{m}$ and CC $i \in \bar{n}$.² The next subsection presents how recommendations are made and what are their effects.

3.1 Recommendation Function

We view recommendations as functions which map users to items, i.e., the recommendations at time t are given by $R^t : \bar{m} \rightarrow \bar{n}$.³ The recommender system is the algorithm which produces such a recommender function based on the current state of the network. In this paper, we look at three such algorithms which we list below from the most to the least exploratory one:

- *Uniform Random (UR)*. Each user can be recommended any CC with an equal probability, i.e., $R_{UR}^t(u) \sim \mathcal{U}(\bar{n})$.
- *Preferential Attachment (PA)*. Each content creator can be recommended with a chance proportional to their current number of followers. That is, if

¹ Although, in principle, on influencer-centered platforms, e.g., YouTube, any user could create content, reports show that only a small number of users do so [25]. This also means that in practice CCs can also follow other CCs. However, since they are relatively few compared to non-CCs, we can opt for simplicity and ignore such edges. Our results can be generalised though when allowing for such edges and thus not restricting the network to a bipartite structure. A more detailed discussion on this issue can be found in [4].

² Throughout the paper we use \bar{k} to denote the set of non-zero natural numbers that are at most equal to k , i.e. $\bar{k} := \{1, 2, \dots, k\}$.

³ Note that under full generality each user receives a *list* of recommendations. However, allowing for a single recommendation per user makes our model comparable with [4].

$a_{.,i} := \sum_{u \in \bar{m}} a_{u,i}$ is the number of followers of CC_i , then:

$$\mathbb{P}(R_{\text{PA}}^t(u) = i) = \frac{a_{.,i}^t + 1}{\sum_{j \in \bar{n}} (a_{.,j}^t + 1)}.$$

- *ExtremePA*. An extreme version of preferential attachment which recommends only (one of) the most popular items:

$$\mathbb{P}(R_{\text{ext}}^t(u) = i) = \begin{cases} 0 & , \text{ if } a_{.,i}^t < \max_j a_{.,j}^t, \\ 1/|\{i, \text{ s.t. } i \in \arg \max_j a_{.,j}^t\}| & , \text{ otherwise.} \end{cases}$$

After each user u receives their associated recommendation $R^t(u)$ they decide to follow $R^t(u)$ iff $R^t(u)$ is better (perhaps vacuously) than any followee of u :

$$a_{u,i}^{t+1} = \begin{cases} 1 & , \text{ if } R^t(u) = i \text{ and } i < j \ \forall j \text{ s.t. } a_{u,j}^t = 1 \\ a_{u,i}^t & , \text{ otherwise.} \end{cases}$$

3.2 Metrics of Interest

Since the network at time $t+1$ only depends on the network at time t , we can view this network formation process as a Markov Chain (MC). More details will follow in Sect. 4. With this interpretation, we first want to investigate (a) whether this is an absorbing MC (i.e., whether there exist some networks - absorbing states - which will eventually be reached and which will not change no matter how many timesteps proceed), and, if so, (b) the expected number of steps until absorption (i.e., until reaching such a state).

However, as mentioned in the introduction, we are also interested in *the fairness* of the overall process, both ex-ante and ex-post. Thus we define:

- *Ex-post individual fairness for CCs*. We say that a network A is (*individually*) *fair* if the number of followers of CCs have the same ranking as the CCs, i.e., if $a_{.,1} \geq a_{.,2} \geq \dots \geq a_{.,n}$. We also say that A is (*individually*) *CC_i-fair* if CC_i is one of the top i ranked CCs, i.e. if $|\{j : a_{.,j} > a_{.,i}\}| < i$. In particular, this means that a network is fair iff it is *CC_i-fair* for all $i \in \bar{n}$.⁴
- *Ex-ante individual fairness for CCs*. When this is an absorbing MC, we can also look at the ex-ante fairness of the network formation process. More precisely, we say that a process is *ex-ante (individually) fair* if the expected number of followers of CCs at absorption is decreasing, i.e., if $\mathbb{E}[a_{.,1}^\infty] \geq \mathbb{E}[a_{.,2}^\infty] \geq \dots \geq \mathbb{E}[a_{.,n}^\infty]$. Similarly, we say that a process is *ex-ante (individually) CC_i-fair* if $|\{j : \mathbb{E}[a_{.,i}^\infty] < \mathbb{E}[a_{.,j}^\infty]\}| < i$.

⁴ Note that this is a weak version of the fairness definition. Alternatively, we can say a network A is fair if $a_{.,1} > a_{.,2} > \dots > a_{.,n}$. However, as we will see later in the results section, as the number of users goes to infinity the chance of achieving equality goes to zero. So, in the limit, the two definitions are equivalent.

4 Results

In our results section we investigate the metrics of interest introduced in Sect. 3.2. As such, Sect. 4.1 shows how this process can be viewed as a MC and what are the absorbing states. Section 4.2 looks at the expected time to absorption. Finally, in Sect. 4.3, we build on the prior two and investigate fairness. Each of these subsections starts with (a) a summary paragraph which gives an overview of the results, and (b) a take-away paragraph which discusses the relevance of these results. The remainder of each subsection is used for the results themselves. In the interest of space, in a few of the later proofs we omit details and only provide the proof's outline and intuition.

4.1 An Absorbing Markov Chain

Summary. This subsection starts by proving that our process is an absorbing MC (see Theorem 1). During the proof we also define the transition matrix. Figure 1 illustrates this process for the small case of two CCs and two (non-CC) users. The remaining results characterize the absorbing states. First, Theorem 2 shows that under exploratory recommendation processes (i.e., PA or UR) the absorbing states are the ones where each user follows the best CC. Second, Theorem 3 shows that, under ExtremePA, a state reachable from the $\mathbf{0}$ state is absorbing if (a) there is a unique CC with the maximum number number of followers and (b) no new user would like to follow this CC. For the latter, we use Lemma 1, which states that any state reachable from $\mathbf{0}$ has, for each CC_i with the maximum number of followers, a user who does not follow any CC better than CC_i .

Take-away. The results of this subsection build a representation for the process which facilitates its understanding. Importantly, by describing the transition matrix we see the differences between the three RSs: (a) ExtremePA has more absorbing states, (b) out of these, the states where CC_1 is the most followed CC (i.e., the CC_1 -fair ones) are the absorbing states under the exploratory RSs, (c) PA situates itself between ExtremePA and UR in terms of exploration. By the latter we mean that, when compared to UR, PA could remain for longer times in transient states (which are fairness-wise similar to the states which are absorbing only under ExtremePA).

Theorem 1. $(A^t)_{t \geq 0}$ is an absorbing MC (for all three RSs).

Proof. First, $(A^t)_{t \geq 0}$ is a MC, as A^{t+1} depends only on the network configuration at the previous step. More precisely, it is a MC where: (a) the state space is $\{0, 1\}^{m \times n}$, (b) the initial distribution is λ where $\lambda_A = 1$ iff A is the zero matrix, and (c) the transition matrix is given by $p_{B,C} := \mathbb{P}(A^{t+1} = C | A^t = B)$, where:

- from the zero matrix we can only transit to a matrix where each user follows exactly one CC, i.e.: $p_{\mathbf{0},C} = \begin{cases} 1/n^m, & \text{if } c_{u..} = 1 \forall u \in \overline{m}; \\ 0 & , \text{ otherwise} \end{cases}$;




Fig. 1. An example of the MC representation for $n = m = 2$. The coloured node is the starting state. We use (a) full edges for transitions that are the same under the three RSs, and (b) dotted edges for transitions that differ (labeled with the respective probabilities for (1) ExtremePA, (2) PA, (3) UR). Dots replace probabilities when the starting state is not reachable from $\mathbf{0}$ under ExtremePA.

- from any other matrix we can only transit to a new matrix where each user either (a) follows the same CCs as before, or (b) follows exactly one more CC which is better than the best CC they followed so far; the probabilities of such transitions depend on the recommendation process.

Moreover, from the shape of the transition matrix it follows that this process is an absorbing MC. To see this, note that $p_{B,C}$ is non-zero iff $C = B$ or there is some user u who follows one more CC, i.e. $c_{u,\cdot} > b_{u,\cdot}$. So, any state B is either absorbing, or can transit to a state of a strictly higher sum of elements. Since the sum of elements of B is bounded by $m \cdot n$, such a sequence of transitions must be finite. Thus, in this latter case, an absorbing state will be eventually reached through a sequence of transitions. \square

Theorem 2. Under PA and UR recommendations, a state B is absorbing iff all users follow the best CC, i.e. iff $b_{u,1} = 1$ for all $u \in \bar{m}$.

Proof. (\Rightarrow) Assume there exist some user $u \in \bar{m}$ s.t. $b_{u,1} = 0$. Under PA and UR recommendations there is always a non-zero chance u is recommended CC_1 . If this happens, then u follows CC_1 . So we can transit to a state $C \neq B$ where $c_{u,1} = 1$ with a non-zero probability. Hence, B is not absorbing.

(\Leftarrow) The reverse is straightforward. If $b_{u,1} = 1$, then no recommendation will change the followers of u . This holds for all users u , so B is absorbing. \square

Lemma 1. For any state B reachable from $\mathbf{0}$ and any $i \in \bar{n}$ s.t. $b_{\cdot,i} = \max_j b_{\cdot,j}$ there exists some user $u \in \bar{m}$ s.t. $b_{u,j} = 0$ for all $j < i$.

Proof. We use an inductive argument. The claim is obviously true for any state directly reachable from $\mathbf{0}$ (i.e., any network achievable in one timestep), as, in such states, each user follows exactly one CC. Next, we assume the claim is true for some state B . From this state, we either remain in B , or we transit to a state C where a subset of these CCs have a higher maximum number of followers. As each CC_i in this subset increased their number of followers they needed to be recommended to a user u that preferred them to any CC they previously followed. As such, CC_i is the best CC followed by u in C , i.e., $i = \min\{j : c_{u,j} = 1\}$. \square

Theorem 3. *Under Extreme PA, all absorbing states B reachable from $\mathbf{0}$ are such that (a) there exists some $i \in \bar{n}$ s.t. $b_{\cdot,i} > b_{\cdot,j}$ for all $j \in \bar{n} - \{i\}$, and (b) for all $u \in \bar{m}$ there exists some $j \leq i$ s.t. $b_{u,j} = 1$. Moreover, every state which satisfies (a) and (b) is absorbing.*

Proof. For the first part, let B be an absorbing state reachable from $\mathbf{0}$. If (a) is not the case then there are at least two CCs, say CC_i and CC_j with $i < j$, that have the maximum number of followers. By Lemma 1, there exists some user u who does not follow CC_i yet. Therefore, if everybody is recommended CC_i , only the number of followers of CC_i will increase by at least one (as u will follow them). There is a non-zero probability of this happening, and, in such a case, a new state is reached. Therefore, B is not an absorbing state if (a) does not hold. Next, if (a) holds but (b) does not then in the next round (when the unique most followed CC_i is recommended to everybody) new users will follow CC_i . So, again, if (a) holds but (b) does not then B cannot be absorbing.

The second part is straightforward. If (a) holds then only the unique CC with the highest number of followers can be recommended in the next round. Since (b) holds, no new user will follow this CC. So, there is no chance of transiting to a new state, i.e. the current state is absorbing. \square

4.2 Expected Time to Absorption

Summary. Theorem 4 shows that as the number of users goes to infinity, the expected time to absorption under ExtremePA is less than 2. Two preliminary results are needed to prove this statement: (a) the chance of having ties in the number of followers after the first round of recommendations (Lemma 2) and (b) the chance that in the second round no one new would follow the most followed CC_i (if $i \neq n$, see Lemma 3) both go to 0 as the number of users goes to infinity.⁵ These results are used for the annotations in Fig. 2, which summarizes the process and provides the intuition for the proof of the theorem.

Take-away. This puts ExtremePA in sharp contrast with the exploratory RSs, as prior work indicates that, under the UR and PA recommendation scenarios, the convergence time increases logarithmically in the number of CCs and linearly (or sub-linearly) in the number of users (see Figs. 2 and 7b of [4]). Importantly, it indicates that, while for ExtremePA it could be sufficient to analyse the fairness in the absorbing states alone, this might not be the case for PA and UR: as these RSs could lead, in practice, to long convergence times the fairness in the transient states should also be evaluated. This is particularly true for PA, which, as discussed earlier, could remain in unfair states for longer than UR.

Lemma 2. *For any $i \neq j \in \bar{n}$, $\mathbb{P}(a_{\cdot,i}^1 = a_{\cdot,j}^1) \rightarrow 0$ as $m \rightarrow \infty$.*

⁵ Note that the two preliminary results are independent on the RS.

Proof. First, let Y_u be the following random variables depending on the recommendations in the first round:

$$Y_u = \begin{cases} 0 & , \text{ if user } u \text{ is recommended some } CC_k \text{ with } k \neq i, j; \\ 1 & , \text{ if user } u \text{ is recommended } CC_i; \\ -1 & , \text{ if user } u \text{ is recommended } CC_j. \end{cases}$$

Since the first recommendations are uniform random, $(Y_u)_{u \in \bar{m}}$ are i.i.d. (with $\mathbb{P}(Y_u = 0) = (n - 2)/n$ and $\mathbb{P}(Y_u = c) = 1/n$ when $c = \pm 1$). Hence, $\mathbb{E}[Y_u] = 0$ and $\text{Var}(Y_u) = \mathbb{E}[Y_u^2] = 2/n$. By the central limit theorem it follows that,

$$\frac{\sum_{u \in \bar{m}} Y_u - m \cdot 0}{2/n \cdot \sqrt{m}} \rightarrow^{\mathcal{D}} \mathcal{N}(0, 1).$$

In particular, this implies that $\lim_{m \rightarrow \infty} \mathbb{P}(a_{\cdot, i}^1 = a_{\cdot, j}^1) = 0$. □

Lemma 3. For any $i \in \overline{n-1}$, $\mathbb{P}((a_{\cdot, i}^1 > a_{\cdot, j}^1 \forall j \in \bar{n}) \wedge (\forall u \in \bar{m}, \exists j \in \bar{i} \text{ s.t. } a_{u, j}^1 = 1)) \rightarrow 0$ as $m \rightarrow \infty$.

Proof. This follows as the latter part of the conjunction corresponds to the scenario of all users being recommended one of the top i CCs in the first round. So, $\mathbb{P}((a_{\cdot, i}^1 > a_{\cdot, j}^1 \forall j \in \bar{n}) \wedge (\forall u \in \bar{m}, \exists j \in \bar{i} a_{u, j}^1 = 1)) \leq \mathbb{P}(\forall u \in \bar{m}, \exists j \in \bar{i} a_{u, j}^1 = 1) = (\frac{i}{n})^m \rightarrow 0$ as $m \rightarrow \infty$. □




Fig. 2. A summary of the MC for ExtremePA as $m \rightarrow \infty$. Rectangles represent sets of states. Transition probabilities are annotated in the limit. Dotted lines correspond to transitions who are negligible as $m \rightarrow \infty$.

Theorem 4. Under ExtremePA, the expected time to absorption goes to $2 - 1/n$ as m goes to infinity.

Proof. Based on Theorems 1 and 3, we can group states in subsets, as shown in Fig. 2. Let μ_B be the expected time from the state B to absorption. Then $\mu_{B: B \in S^*} = 0$ (as all states in S^* are absorbing) and $\mu_{B: B \in S} = 1$ (as all states in S lead in one timestep to a state in S^*). Therefore,

$$\mu_0 = 1 \cdot p_{0, S_n^*} + 1 \cdot p_{0, S - S_n^*} + 2 \cdot p_{0, S} + \sum_{B \in E} (1 + \mu_B) \cdot p_{0, B}$$

We can use Lemmas 2 and 3 to find the transition probabilities between the sets of states as $m \rightarrow \infty$: (a) $p_{\mathbf{0}, S_n^*} \rightarrow \frac{1}{n}$, (b) $p_{\mathbf{0}, S^* - S_n^*} \rightarrow 0$, (c) $p_{\mathbf{0}, S} \rightarrow \frac{n-1}{n}$, and (d) $p_{\mathbf{0}, E} \rightarrow 0$. Since, in addition, $\sum_{B \in E} p_{\mathbf{0}, B} \cdot (1 + \mu_B) \leq p_{\mathbf{0}, E} \cdot (1 + \mu_{B^*})$ (where $B^* = \arg \max_{B \in E} \mu_B$) and $\mu_{B^*} \leq c$ (c constant),⁶ the result follows. \square

4.3 Fairness for Content Creators

Summary. Building on previous results, we show that the exploratory RSs are both ex-ante and ex-post CC_1 -fair (Corollary 1), while ExtremePA is only ex-ante fair (Theorem 5). In fact, Corollary 2 shows the probability of achieving a CC_1 -fair absorbing state under ExtremePA goes to $1/n$ as the number of users goes to infinity, while Corollary 3 shows that the probability of achieving a fair outcome for all CCs goes to $1/n!$. Finally, Fig. 3 depicts the probability of achieving a CC_i fair outcome for each CC under any of the three RSs.

Take-away. This analysis carries several important messages. First, it shows that exploration is key in achieving CC_1 -fair outcomes. Second, although ExtremePA is ex-ante fair it is rarely ex-post fair, thus underlying the importance of looking beyond the number of followers in expectation. This finding is congruent with the experimental results of [7],⁷ thus it confirms that real-life RSs share important similarities with ExtremePA. Third, our numerical analysis reveals that although exploration always leads to a CC_1 -fair outcome, the fairness for the other CCs is not guaranteed. However, it suggests that exploratory RSs distribute fairness more homogeneously compared to ExtremePA.

Corollary 1. *Under UR and PA recommendations, $(A^t)_t$ is both ex-post and ex-ante CC_1 -fair.*

Proof. This is an immediate consequence of Theorem 2. Since B is absorbing iff $b_{.,1} = m$, all absorbing states are CC_1 fair (i.e., we have ex-post fairness) and $\mathbb{E}[b_{.,1}] = m \geq \mathbb{E}[b_{.,i}]$ (i.e., ex-ante CC_1 -fairness). \square

Corollary 2. *Under ExtremePA, the probability the outcome is ex-post CC_1 -fair goes to $1/n$ as $m \rightarrow \infty$.*

Proof. The final outcome (a) is always CC_1 -fair if CC_1 is the unique most followed CC after the first round, and (b) is CC_1 -fair only if CC_1 is one of the CCs with a maximum number of followers after the first round. So, using the notation from Fig. 2: $p_{\mathbf{0}, S_1 \cup S_1^*} \leq \mathbb{P}(CC_1 - \text{fair}) \leq p_{\mathbf{0}, S_1 \cup S_1^*} + p_{\mathbf{0}, E}$. The claim follows, since, as shown previously, $p_{\mathbf{0}, S_1 \cup S_1^*} \rightarrow 1/n$ and $p_{\mathbf{0}, E} \rightarrow 0$. \square

Corollary 3. *Under ExtremePA, the probability the outcome is ex-post fair goes to $1/n!$ as $m \rightarrow \infty$.*

⁶ This can be easily shown for $c = 4n/(n - 1)$ by using Lemma 1 to prove that $p_{B^*, S \cup S^*} \geq 1/2n$ (although better bounds can be obtained).

⁷ “Although, on average, quality is positively related to success, songs of any given quality can experience a wide range of outcomes” (from [7], p. 855).

Proof. By Lemma 2, when $m \rightarrow \infty$ the probability of achieving a tie in the number of followers after the first round goes to zero. After ignoring ties, only outcomes with $a_{\cdot,1}^1 > \dots > a_{\cdot,n}^1$ lead to a fair outcome. By symmetry, all strict orderings of $(a_{\cdot,i}^1)_i$ have an equal probability. Since there are $n!$ such orderings, $\mathbb{P}(a_{\cdot,1}^1 > \dots > a_{\cdot,n}^1) \rightarrow 1/n!$ as $m \rightarrow \infty$. The conclusion follows. \square




Fig. 3. Sampling probability of ex-post CC_i -fairness, under the three RSs: ExtremePA (blue), PA (purple), and UR (orange). For each RS, we run 10,000 simulations until convergence with $n = 100$ and $m = 10,000$. The dashed line denotes a reference value of 0.5. Fairness under ExtremePA increases from best- to lowest-quality CCs. PA and UR achieve higher fairness for the best CCs.

The behavior depicted in Fig. 3 is consistent with the theoretical results on ex-post fairness. In particular, while UR and PA guarantee CC_1 -fairness at convergence, ExtremePA only achieves it with a low probability (about 1%). More precisely, the probability of ex-post CC_i -fairness grows linearly from $1/n$ to 1 with the quality index i . In contrast, UR and PA generally achieve more (less) fair outcomes for the top (lowest) half of CCs. Between the two, the most exploratory recommendation strategy (UR) is more fair for most high-quality CCs, while PA is fair only for a small number of top (and bottom) quality CCs.

Theorem 5. *Under ExtremePA, the final outcome is ex-ante fair.*

Proof. The proof is based on the following observation: when CC_i becomes the most followed CC, they will be eventually followed by all CCs who did not follow a better-quality CC before. For example, if CC_2 is the most followed after the first round, all users (except those who were recommended CC_1) will follow CC_2 after round 2. Differently, if CC_1 was the most followed, then everybody will end up following CC_1 next, while if CC_3 was the most followed, then everybody except those who originally followed CC_1 or CC_2 will follow CC_3 . This intuitively leads to CC_2 having more followers in expectation than CC_3 and fewer than CC_1 .

For simplicity, we will only formalize this intuition for $n = 2$:

$$\mathbb{E}[a_{.,1}^{\infty}] = \sum_{k=0}^m \mathbb{P}(a_{.,1}^1 = k) \cdot \mathbb{E}[a_{.,1}^{\infty} | a_{.,1}^1 = k] = \sum_{k=0}^{\lfloor \frac{m-1}{2} \rfloor} \frac{\binom{k}{m}}{n^m} \cdot k + \sum_{k=\lfloor \frac{m-1}{2} \rfloor + 1}^m \frac{\binom{k}{m}}{n^m} \cdot m.$$

As the sum is larger than $\mathbb{E}[a_{.,2}^{\infty}] = \sum_{k=0}^m \frac{\binom{k}{m}}{n^m} \cdot k$, ExtremePA is ex-ante fair. \square

5 Conclusion

In this work we analyzed the role recommendations play in UGC-based social networks with respect to the individual fairness for CCs. To do so, we (a) extended prior models with a RS-framework, (b) introduced a non-exploratory RS, and (c) defined ex-ante and ex-post measures of fairness. Our results showed that the network formation process is an absorbing Markov Chain with different absorbing states, expected times, and fairness guarantees depending on the RS. In particular, the expected absorbing time under ExtremePA is bounded by 2, i.e., much faster than the ones under UR or PA (which are linearly or sub-linearly increasing in m). Furthermore, while all studied RSs guarantee ex-ante fairness, ex-post fairness is rarely attainable: without exploration (i.e., under ExtremePA), ex-post fairness is achieved with probability $1/n!$, and even with exploration (i.e., under UR or PA), ex-post fairness is guaranteed only for the best (and lowest) CCs. In essence, exploration in RSs trades faster absorption times for higher probabilities of achieving fair outcomes for the best CCs.

A key feature of our model is its simplicity. On the one hand, some level of model simplicity is both inevitable for a feasible theoretical analysis and useful for its clarity. On the other hand, the assumptions about both users and the RS are not always fulfilled by real-world systems. For example, here we reduced both the interests of users and the performance of CCs to one ordinal dimension: quality. Subsequent iterations of this model, which include multiple and different types of attributes, would automatically go beyond a uniform preference of users. Therefore, they could investigate the effect of personalized RSs (e.g., collaborative filtering), look beyond individual fairness (e.g., group fairness), and capture the evolution of users (e.g., changes in the taste of users). In short, UGC-centered platforms are complex sociotechnical systems, and our work is just one piece of this elaborate puzzle. Although there are still many unanswered questions, we believe this mathematical analysis carries one cornerstone message: if we want a fair community, one which rewards its creators based on the quality of their work, then we must encourage exploration.

Acknowledgements. The authors gratefully acknowledge financial support from the University of Zürich and the SNSF under the NCCR Automation grant (Grant Number 180545). We also thank the anonymous reviewers for their valuable feedback.

References

1. Bakshy, E., Rosenn, I., Marlow, C., Adamic, L.: The role of social networks in information diffusion. In: Proceedings of the 21st international conference on World Wide Web, pp. 519–528 (2012)
2. Hall, W., Tinati, Ramine, Jennings, W.: From Brexit to Trump: social media’s role in democracy. *Computer* **51**(1), 18–27 (2018)
3. Stacy Adams, J.: Towards an understanding of inequity. *J Abnormal Social Psychol.* **67**(5), 422 (1963)
4. Pagan, N., Mei, W., Li, C., Dörfler, F.: A meritocratic network formation model for the rise of social media influencers. *Nature Commun.* **12**(1), 1–12 (2021)
5. Zipf, G.K.: Human behavior and the principle of least effort: an introduction to human ecology. Ravenio Books (2016)
6. Vany, A.D.: Hollywood economics: how extreme uncertainty shapes the film industry. Routledge (2003)
7. Salganik, M.J., Dodds, P.S., Watts, D.J.: Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**(5762), 854–856 (2006)
8. McNee, S.M., Riedl, J., Konstan, J.A.: Being accurate is not enough: how accuracy metrics have hurt recommender systems. In: *CHI’06 Extended Abstracts on Human Factors in Computing Systems*, pp. 1097–1101 (2006)
9. Kunaver, M., Požrl, T.: Diversity in recommender systems—a survey. *Knowl. Syst.* **123**, 154–162 (2017)
10. Helberger, N., Karppinen, Kari, D’Acunto, L.: Exposure diversity as a design principle for recommender systems. *Inf. Commun. Soc.* **21**(2), 191–207 (2018)
11. Gravino, P., Monechi, B., Loreto, V.: Towards novelty-driven recommender systems. *Comptes Rendus Physique* **20**(4), 371–379 (2019)
12. Guo, W., Krauth, K., Jordan, M., Garg, N.: The stereotyping problem in collaboratively filtered recommender systems. In: *Equity and Access in Algorithms, Mechanisms, and Optimization*, pp. 1–10. ACM (2021)
13. Myerson, R.B.: Utilitarianism, egalitarianism, and the timing effect in social choice problems. *Econometrica J. Econ. Soc.* 883–897 (1981)
14. Chaney, A.G.B., Stewart, B.M., Engelhardt, B.E.: How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In: *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–232 (2018)
15. Erdős, P., Rényi, Alfréd: On random graphs, i. *Publicationes Mathematicae (Debrecen)* **6**, 290–297 (1959)
16. Watts, D.J., Strogatz, S.H.: Collective dynamics of “small-world” networks. *Nature* **393**(6684), 440 (1998)
17. Barabási, A., Albert, R.: Emergence of scaling in random networks. *Science* **286**(5439), 509–512 (1999)
18. Snijders, T.A.B.: Stochastic actor-oriented models for network change. *J. Math. Sociol.* **21**(1–2), 149–172 (1996)
19. Jackson, M.O.: *Social and Economic Networks*. Princeton University Press (2010)
20. Verma, S., Rubin, J.: Fairness definitions explained. In: *2018 IEEE/ACM International Workshop on Software Fairness (Fairware)*, pp. 1–7. IEEE (2018)
21. Garg, P., Villasenor, J., Foggo, V.: Fairness metrics: a comparative analysis. In: *2020 IEEE International Conference on Big Data (Big Data)*, pp. 3662–3666. IEEE (2020)
22. Mitchell, S., Potash, E., Barocas, S., D’Amour, A., Lum, K.: Algorithmic fairness: choices, assumptions, and definitions. *Annual Rev. Stat. Appl.* **8**, 141–163 (2021)

23. Binns, R.: On the apparent conflict between individual and group fairness. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, pp. 514–524 (2020)
24. Lucherini, E., Sun, M., Winecoff, A., Narayanan, A.: T-recs: a simulation tool to study the societal impact of recommender systems (2021). [arXiv:2107.08959](https://arxiv.org/abs/2107.08959)
25. Nielsen, J.: Participation inequality: lurkers vs. contributors in internet communities. *Jakob Nielsen's Alertbox* **107**, 108 (2006)