



Reconciling the Quality vs Popularity Dichotomy in Online Cultural Markets

ROSSANO GAETA, MICHELE GARETTO, and GIANCARLO RUFFO, Università di Torino,
Dipartimento di Informatica, Italy
ALESSANDRO FLAMMINI, Indiana University Network Science Institute, USA

We propose a simple model of an idealized online cultural market in which N items, endowed with a hidden quality metric, are recommended to users by a ranking algorithm possibly biased by the current items' popularity. Our goal is to better understand the underlying mechanisms of the well-known fact that popularity bias can prevent higher-quality items from becoming more popular than lower-quality items, producing an undesirable misalignment between quality and popularity rankings. We do so under the assumption that users, having limited time/attention, are able to discriminate the best-quality only within a random subset of the items. We discover the existence of a harmful regime in which improper use of popularity can seriously compromise the emergence of quality, and a benign regime in which wise use of popularity, coupled with a small discrimination effort on behalf of users, guarantees the perfect alignment of quality and popularity ranking. Our findings clarify the effects of algorithmic popularity bias on quality outcomes, and may inform the design of more principled mechanisms for techno-social cultural markets.

CCS Concepts: • **Information systems** → **Retrieval models and ranking**; **Probabilistic retrieval models**; *Users and interactive retrieval*; *Presentation of retrieval results*;

Additional Key Words and Phrases: Online cultural markets, ranking algorithms, popularity bias, retrieval diversity

ACM Reference format:

Rossano Gaeta, Michele Garetto, Giancarlo Ruffo, and Alessandro Flammini. 2023. Reconciling the Quality vs Popularity Dichotomy in Online Cultural Markets. *ACM Trans. Inf. Syst.* 41, 1, Article 25 (January 2023), 34 pages.
<https://doi.org/10.1145/3530790>

1 INTRODUCTION

It is a common experience to receive recommendations from an expert system. It happens when we buy a book or music from Amazon, when we scroll the lists that Google returns in response to our queries, when we are shown potential friends on a social media site such as Twitter or Facebook, or when we search a generic news article on a news aggregation service. Indeed, the availability of cultural and informational products, and the number of people we could potentially

A. Flammini acknowledges support from DARPA award HR001121C0169.

Authors' addresses: R. Gaeta, M. Garetto and G. Ruffo, Università di Torino, Dipartimento di Informatica, C.so Svizzera 185, Turin, Italy, 10149; emails: {rossano.gaeta, michele.garetto, giancarlo.ruffo}@unito.it; A. Flammini, Indiana University Network Science Institute, 919 East Tenth st., Bloomington, Indiana, USA, 47408; email: aflammin@indiana.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

1046-8188/2023/01-ART25 \$15.00

<https://doi.org/10.1145/3530790>

connect to via social media is so vast that without recommender systems [42] we would have little chance to find what we like or need. Such systems, broadly speaking, try to infer how pleased a user would be with a particular item, given her history of items consumption and taking into consideration the choices made from other, possibly similar, users, and return a list of personalized top-N items [17]. While there is a strong incentive to create systems that produce personalized recommendations and therefore pay special attention to choices made by users similar to the target of the recommendation, the popularity of an item (measured e.g., as the number of copies sold, views, downloads, likes) is a signal that is often leveraged to produce recommendations. In [26], for example, authors compare several recommendations strategies from different perspectives, including accuracy, catalog coverage and their bias to recommend popular items, and show that recent algorithmic techniques end up recommending mostly top sellers.

The notion that popularity is an indicator of quality is predicated on the notion of wisdom of the crowds [49], the fact that an assessment by many independent - even non expert - individuals could be more precise/valuable/correct than that of few experts. Another justification for using popularity as a proxy for quality is that quality is a concept that is hard to define and measure, as it is intimately tied to highly subjective notions such as beauty, novelty, and virality.

There is, however, ample empirical evidence showing that social influence (resulting for example by popularity-based ranking algorithms) can bias the success of different items in ways that do not reflect their intrinsic quality [29]. Intuitively, *popularity bias* can reinforce initial random fluctuations and crystallize a ranking in popularity that is misaligned with that based on quality, severely undermining the wisdom of crowd and producing giant distortions in the relative success of products.

The notion that popularity begets popularity is ubiquitous. It is generally known as Matthew effect, a term introduced by R. K. Merton [34] to describe the disparity of recognition attributed to known and relatively unknown scholars for producing work of comparable quality. The issue of the possible distortions introduced by popularity has been studied even more directly by Salganick and collaborators. In their music lab experiment [44] they divided participants in non-communicating groups, and asked them to select songs from a menu common across groups. When exposed to other members choices, groups produced very different popularity rankings due to the “market” enhancing the idiosyncratic initial choices of few via the popularity driven dynamics. This supports the evidence that success is very hard to predict and engineer [52].

The problem becomes harsher when popularity bias compounds with the limited attention that users have to make quality discriminations. Evaluating products comes with a cognitive cost for the consumer that, in turn, influences the amount of *attention* he/she puts in this effort. Previous research has shown that the interplay between popularity bias that could be introduced by recommender systems and consumer attention strongly impacts the quality of items sale ranking. Although traditional evaluation metrics of novel collaborative filtering recommender systems [23] take into account such interplay, the problem is still considered hard to be solved. Qiu et al. [40] studied a system of social recommendation where agents’ choices depend both on items quality and popularity among neighbors. They found (i) a very non-linear relationship between popularity and quality of items, (ii) a poor alignment between the quality and the popularity ranking, and (iii) an inverse relationship between said alignment and diversity of the market. Ciampaglia et al. [15] studied a market where a combination of choice by quality and popularity is performed and showed the existence of a regime where top-quality items are pushed at the top of the sale ranking.

The purpose of this paper is to better understand, from an analytical perspective, the interplay between popularity and quality as determinants of users choices in an idealized online cultural

market, and the conditions that realize a desirable alignment between items ranking in popularity and quality.

For simplicity, we will not consider personalized recommendations, i.e., recommendation lists tailored to the specific user, or enhanced, say, by collaborative filtering techniques. There are, indeed, scenarios in which such personalization is not really needed, for example because the user has already issued a query for a specific category of objects matching her/his interests. For example, consider a user searching for books on the “Python programming language” on an e-commerce platform: the platform might have a few tens of books in its catalog about the Python language, and must decide a way to present them to the user, in a list more or less biased by popularity, knowing that the average user will inspect only a subset of the items, in the case of a long list. This scenario is different from the case of users who just log in the platform without having specific ideas about what to buy/watch, thus receiving personalized recommendations based on her/his navigation history, or history of similar users.

Moreover, for the sake of generality and of analytical tractability our approach assumes a generic and idealized online cultural market, and it is not meant to propose any specific technique to be incorporated into a real recommendation system. Nevertheless, we suggest potential applications of our findings. In particular, our results can help define the regimes in which a more or less aggressive dependence on popularity in the recommendation algorithm can have the undesired effect to promote items that are not those that qualify as top quality according to the average user perception.¹

Our contributions can be summarized as follows.

- We propose a simple model of an idealized online cultural market, in which N items have a hidden, intrinsic quality metric. Users pay attention to just a random subset of the items, but are able to select the best-quality item within this subset. Items are recommended to users by a ranking algorithm possibly biased by the current items’ popularity.
- We analyze the asymptotic system behavior as the number of user interactions grows large, with the goal of understanding whether the ranking associated to quality eventually emerges, aligning quality with popularity.
- We discover that the system can have multiple equilibria depending on the ranking algorithm (parameterized by a power-law exponent $\alpha \geq 0$) and the quality discrimination power of the users (parameterized by the number K of inspected items). Notably, different equilibria produce misalignments which affect first top-quality items, before involving also lower-quality ones.
- We characterize the minimum user discrimination power K_{\min} that guarantees that the desired alignment between quality and popularity is always achieved, in the long run.
- We extend the model to multiple user classes, to account for the fact that different users can perceive a different quality in the same item, and we conduct a preliminary investigation within this context.

Our main findings can be summarized as follows.

- In the special case in which items inspected by users are selected uniformly at random ($\alpha = 0$), a minimum discriminating power $K_{\min} = 2$ allows alignment of quality and popularity, at the expense of a low average quality of items selected by users.

¹In Example 4.1, at the end of Section 4, we provide a numerical example of application of our findings in a concrete scenario.

- In the case of ranking algorithms biased by popularity ($\alpha > 0$), there is a harmful regime of mild popularity bias ($0 < \alpha \leq 1$) in which quality struggles to emerge, requiring a disproportionately large K to guarantee the desired alignment.
- With stronger popularity bias ($\alpha > 1$), K_{\min} is bounded and typically small. Increasing α , K_{\min} approaches again the value of 2, at the expense of increasing unfairness among items (the top-quality item monopolizes the market).
- An increasing fraction of naïve users, who deterministically select the most popular item, makes things worse and worse, up to the point that alignment cannot be restored by the other users, no matter their quality discrimination power.

The paper is organized as follows: Section 2 describes our model of online cultural markets and our main assumptions. In Section 3, we analyze the case where items inspected by users are selected uniformly at random, whereas Section 4 deals with ranking algorithms biased by popularity. Section 5 deals with the problem of parameter optimization to achieve a desired level of fairness among the items. In Section 6, we extend the model to the case of multiple user classes. We discuss related work in Section 7. At last in Section 8, we discuss our contributions and findings, point out the limitations of our approach, and outline directions of future research.

2 SYSTEM MODEL AND MAIN ASSUMPTIONS

We consider a fixed set \mathcal{N} of $N = |\mathcal{N}|$ items that the recommender system/search-engine/e-commerce platform/social network can offer to users issuing a certain query/interested in a given object category/exploring a given topic. We assume that each item i , $i = 1, 2, \dots, N$, has an intrinsic quality $q_i \in \mathbb{R}$, which is unknown to the system (and to the users). The notion of quality—although the term is frequently (informally) used—is subtle and hard to define. For the purpose of our analysis, quality is simply an (unobserved) measure that underlies a ranking. The ranking is the one upon which users would converge upon had they the chance to choose items independently, i.e., in the absence of social influence in general and of algorithmic popularity bias in particular. We assume that both the recommender system and the users are eventually interested in promoting/discovering items having higher values of this hidden metric. The absolute values of q_i are not important, provided that they produce the same ranking. Hence, without loss of generality, we assume that items are indexed in increasing order of their quality, i.e., $q_1 < q_2 < \dots < q_N$. Here, item N has the highest-quality and item 1 the lowest. The item ranking is assumed to be fixed over the time period on which we study popularity dynamics.

Items also have an integer-valued ‘popularity weight’ $w_i \in \mathbb{N}^+$, which represents, e.g., the number of times they have been purchased/selected by the users, the number of views/likes/comments received, etc. For each item, this number increases over time, when the item ‘wins’ the competition with the other items.

Let n be the number of competition rounds performed so far within the system among the N items, and $w_i[n]$ the weight accumulated by item i after the n -th competition. For simplicity, we will assume that weight w_i just counts the number of competitions won by item i , i.e., it is increased by one each time item i wins. At time step n , a user receives a *randomized* list of recommended items drawn from the available set of N objects. The randomization introduced by the system is biased by the current popularity weights $w_i[n]$ of the objects (see next). Due to limited time/attention, the user inspects only K objects, and selects the one with the best-quality among the K . By so doing, we account for the fact that users do not have enough time/skills to discover the best-quality item among all alternatives, but they are at least able to do so on a restricted set. In the following, we

will refer to the subset of K items inspected by the user as *pre-selection*. If the user arriving at time n selects objects j , popularity of j increases by one²: $w_j[n] = w_j[n-1] + 1$.

Popularity weights start from arbitrary initial values $w_i[0] > 0$. For example, we can assume that at the beginning all items are equally popular with weight 1, $w_i[0] = 1$, $i = 1 \dots N$. Most of the results in this paper do not depend on the set of initial weights, and we will point out explicitly when they do.

We denote by $b_i[n]$ the probability that item i wins the n -th competition. Such probability depends crucially both on the way in which the system presents available items to the user (all of them or just a subset of them), for example through a vertical scrollable list, and on the user behavior, especially her/his patience to explore the alternatives and her/his ability to identify higher-quality items, as expressed by discrimination parameter K . Probabilities $b_i[n]$ allow us to define the concept of *average quality index* of the online cultural market after n -th competition as $\bar{q}[n] = \sum_{i=1}^N i b_i[n]$.

Our main interest is to investigate what happens, in the long run (as $n \rightarrow \infty$), to popularity weights $w_i[n]$ and to average quality index $\bar{q}[n]$ as we vary the quality-discrimination parameter K . In particular, will asymptotic weights $w_i[n]$ be ordered according to the intrinsic quality of items? Which is the minimum quality-discrimination power (i.e., the minimum K) requested to the users so that the best-quality item will eventually emerge? How does the average quality index depend on the quality-discrimination power?

To answer these questions, we need to specify also the other fundamental ingredient of the model, that is the law by which the restricted, random set of K items is chosen from set \mathcal{N} , which must reflect both the way items are internally promoted by the system and visually presented to the user, and the additional randomness introduced by the user interaction with the online platform.

Similarly to previous work, we adopt the following popularity-biased *ranking model*. The ranking model was introduced in the context of an abstract model of network growth [19] and used in a context similar to the present one in [15]. It aims at describing how items are selected by a ranked list, and it assumes that the entity performing the selection appreciates the differences in ranking, although not necessarily the quantity that underlies the ranking. Let $r_i[n] \in \{1 \dots N\}$ be the rank of item i at the beginning of round n , in terms of popularity weight (i.e., $r_i[n]$ is the number of items - including i - in set \mathcal{N} whose weight is at least as large as $w_i[n]$). Considering ranks, instead of absolute values $\{w_i[n]\}_i$, is a standard technique to avoid disproportionate bias towards items which have accumulated too much weight w.r.t. to others (for example because they have stayed in the system for a much longer time). Then, we assume that K items are successively drawn from set \mathcal{N} (or the set of remaining items after previous extractions) with probability proportional to $r_i[n]^{-\alpha}$, where $\alpha \geq 0$ is an exponent reflecting the bias of the system/user towards popularity. The extreme values are $\alpha = 0$, corresponding to a system in which the K items inspected by the user are chosen uniformly at random, irrespective of their popularity, and $\alpha \rightarrow \infty$, corresponding to a system in which users deterministically focus their attention only on the current top- K items in terms of popularity. The ranking model naturally assumes a monotonically decreasing probability to select an item as a function of its popularity rank. While the power-law form of such dependency is somewhat arbitrary, it is general enough to allow (via the parameter α) to gauge how biased towards the most popular items is the preselection mechanism operated by the putative recommender system.

²We assume for simplicity that competitions occur sequentially one after the other. In real systems many competitions can take place concurrently as several users interact at the same time with the online platform, but we neglect such micro-scale effects.

Table 1. Notation

| symbol | description |
|----------------------------------------------------|----------------------------------------------------------|
| N | number of items |
| n | number of competition rounds |
| $w_i[n]$ | popularity weight of item i after round n |
| $\tilde{w}_i[n]$ | normalized popularity weight of item i after round n |
| $r_i[n]$ | popularity rank of item i at beginning of round n |
| $b_i[n]$ | winning probability of item i at round n |
| $b_i = \lim_{n \rightarrow \infty} b_i[n]$ | asymptotic winning probability |
| $\bar{q}[n]$ | average quality index at round n |
| $\bar{q} = \lim_{n \rightarrow \infty} \bar{q}[n]$ | asymptotic average quality index |
| $\tilde{b}_i \sim (N - i + 1)^{-\beta}$ | desired winning probability of item i |
| $\tilde{q} = \sum_{i=1}^N i \tilde{b}_i$ | desired average quality index |
| α | pre-selection power-law exponent |
| β | fairness power-law exponent |
| K | quality discrimination power of users |
| p_k | distribution of K (if used) |
| f_m | fraction of naïve users |
| C | number of user classes |
| f_c | probability the a user belongs to class c |

A simple, concrete example of mechanism described by the above model could be the following. The system, in response to the user query, generates a random permutation of the N items, according to the above power law of the rank, and presents them to the user in a scrollable list; the user, because of its limited budget of attention, explores only the first K of the list, and after careful inspection is able to select the one with the best quality. Note that our model is not limited to this simple mechanism. More in general, we can describe, through a single parameter α , the generic bias towards popularity of the restricted, random set of items explored by the user as a result of its interaction with the platform (i.e., the items inspected by the user might not necessarily be the first K of the list proposed by the system).

Since the systems we aim to describe are highly heterogeneous in terms of user behavior, we will consider several extensions to the base setting. In particular:

- we consider a case in which the parameter K is not the same for all users to model user diversity in terms of attention/discrimination power. This is achieved by introducing a probability distribution $\{p_k\}_k$ over the quality-discrimination parameter K , which becomes an i.i.d. random variable across different users.
- we also consider the case in which a fraction f_m of users are too ingenuous/impatient and simply select the item which is currently the most popular. N (some platforms indeed mark such item with a special flag, such as “best seller”). We will refer to such users as naïve users. We finally consider the case in which users differ in term of preferences: quality is highly subjective and can be perceived differently by users. We assign users to one of a finite number C of classes, according to fixed (generally heterogeneous) probabilities $\{f_c\}_{c=1}^C$. Users belonging to the same class equally rank the quality of the N items.

Table 1 summarizes the paper notation.

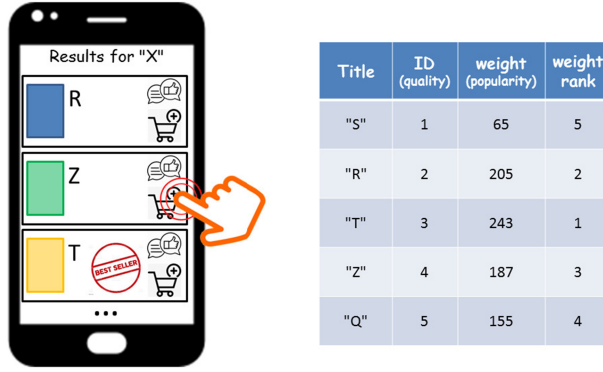


Fig. 1. Illustration of competition among $N = 5$ items, with $K = 3$ items viewed (and inspected) by a certain user: instance of user view (left) and item features (right table). The stylized hand indicates the click over item Z which will be selected by a non-naïve user, being the highest-quality item among the inspected 3 items.

2.1 Toy Example

For clarity, we present a numerical example of a competition round in a system operating under the base setting described above, see Figure 1. Suppose that a user searches for books about some topic “X” on an e-commerce site, and that five books match the query. Books 1, 2, 3, 4, 5, ordered according to increasing intrinsic quality, have title “S”, “R”, “T”, “Z”, “Q”, respectively, and their current popularity rank is 5, 2, 1, 3, 4, respectively, computed for example on the current number of purchased copies, which provides weights $w_i[n]$ at the beginning of the n -th competition (see right table on Figure 1). Suppose the user has time to inspect the reviews of the first three books of the list shown to her/him (left view in Figure 1), discovering the one with the best quality among titles “R”, “Z”, “T”, which happens to be “Z”. This particular (sub)list, in the case of a rank-based power-law pre-selection with exponent $\alpha = 1$, is generated by the recommender system with probability:

$$\frac{\frac{1}{2}}{1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{5}} \cdot \frac{\frac{1}{3}}{1 + \frac{1}{3} + \frac{1}{4} + \frac{1}{5}} \cdot \frac{1}{1 + \frac{1}{4} + \frac{1}{5}} \approx 0.028.$$

The formula above is the product of the three terms that correspond to the probabilities to select R first, Z second and T third. Each of this probabilities is proportional to the inverse of the current rank of the item ($\alpha = 1$ in the example). This accounts for the numerators of the three terms. The denominators are the needed normalization to reflect the fact that once an item has been selected, it cannot be selected again. So, for example, once R (rank=2) has been selected, only four items are still available, those of rank 1, 3, 4, 5, which explains the denominator of the second term.

Note that, in the case $f_m = 0$, the user will end up buying book “Z”, which has the highest quality among the three inspected books. Otherwise, with probability f_m the user is a naïve one, who would instead decide to buy the best-seller “T”.

3 UNIFORM PRE-SELECTION

As a warm-up, we start considering the simple case $\alpha = 0$, where any combination of $K > 1$ (distinct) objects has the same probability to be pre-selected. In this case, the winning probability b_i

of item i does not depend on n and reads:

$$b_i = \begin{cases} 0 & i < K \\ \frac{\binom{i-1}{K-1}}{\binom{N}{K}} & i \geq K \end{cases} \quad (1)$$

since it is equal to the fraction of combinations in which item i is present, and the other $K - 1$ elements of the pre-selection are chosen from the set of $i - 1$ lower-quality items.

The above winning probabilities, for $i \geq K$, are positive and strictly increasing with index i . Moreover, they do not depend on current weights $w_i[n]$. Therefore, $\mathbf{w}[n] = (w_1[n], w_2[n], \dots, w_N[n]) \sim \mathbf{w}[0] + \text{Multi}(b_1, b_2, \dots, b_N; n)$, i.e., $\mathbf{w}[n]$ follows a multinomial probability distribution with parameters $(b_1, b_2, \dots, b_N; n)$, shifted by initial weights $\mathbf{w}[0]$. Let E_n be the event that, after n competitions, weights associated to items with positive winning probability are correctly ordered according to the quality of items, i.e.,

$$E_n = \{w_K[n] < w_{K+1}[n] < \dots < w_N[n]\}.$$

Exploiting classic concentration results we can prove the following:

PROPOSITION 3.1 (ASYMPTOTIC WEIGHTS WITH INCREASING WINNING PROBABILITIES). *When non-null winning probabilities $\{b_i\}_i$ are strictly increasing with index i , as the number of competitions n grows to infinity, the associated weights are almost surely correctly ordered, i.e.,*

$$\mathbb{P} \left[\lim_{n \rightarrow \infty} E_n \right] = 1.$$

PROOF. See Appendix A. □

The above result suggests that a minimum discrimination effort on behalf of the users (i.e., $K = 2$) is sufficient, in the case $\alpha = 0$, to eliminate the effects of popularity bias and guarantee, in the long run, that high-quality items will eventually emerge as popular, with item popularity perfectly aligned with item quality. Note that in the case of $K = 1$, instead, quality never comes into play in determining selections, hence any permutation of the items is equally likely to provide the final ranking of the items in terms of popularity.

The above results can be easily extended to the case of heterogeneous users characterized by a distribution $\{p_k\}_k$ of the discrimination parameter K . Indeed, denoting by $b_i(k)$ the dependency of (1) on K , we simply have the weighted average:

$$b_i = \sum_{k \geq 1} p_k b_i(k).$$

Note that sequence $\{b_i\}_i$ is still strictly increasing for all distributions $\{p_k\}_k$ except the one in which $p_1 = 1$ (i.e., fixed $K = 1$), hence Proposition 3.1 still applies except for this degenerate case.

At last, considering for simplicity the case of constant K , we examine the impact of f_m , the fraction of naïve users who always select the most popular item. Denoting by \hat{b}_i the winning probability of item i with $f_m > 0$, we have

$$\hat{b}_i = f_m \cdot I_{i=\arg \max_j w_j} + (1 - f_m) b_i$$

where $I_{i=\arg \max_j w_j}$ is the indicator function of the event that item i is currently the most popular, while b_i is the winning probability for non-naïve users. In this case the most popular item is not guaranteed to be item N . Indeed an item $i^* < N$ can become the most popular, provided that its winning probability is higher than that of all other items, and most crucially of that of item N , i.e., provided that:

$$f_m + (1 - f_m) b_{i^*} > (1 - f_m) b_N \Rightarrow f_m > \frac{b_N - b_{i^*}}{1 + b_N - b_{i^*}}$$

which means that, if f_m is large enough, item i^* can stably occupy the most popular position. As we increase f_m , starting from 0, we see that $N - 1$ is the first item that can replace N on top of the list. Since configurations in which item N is not the most popular are undesirable, we can compute the maximum value of f_m such that only item N can be the most popular. Indeed, since $b_N - b_{N-1} = \frac{K(K-1)}{N(N-1)}$, we obtain the condition:

$$f_m < \frac{K(K-1)}{N(N-1) + K(K-1)}.$$

We observe that, when $K \ll N$, a small fraction of naïve users is enough to (potentially) disrupt the optimal configuration in which item N is the most popular.

Having analyzed the simple case of uniform pre-selection, in next section we move on to investigate what happens in the more challenging case of popularity bias in the pre-selection of items, i.e., when current popularity weights $w_i[n]$ are used by the system to build the list of items recommended to the user.

4 POPULARITY-BASED PRE-SELECTION

When $\alpha > 0$, more popular items are more likely to be proposed to the user, and if they win the competition (which depends on their intrinsic quality) they become even more popular. This endows the system with a self-reinforcing property reminiscent of the rich-get-richer phenomenon. What happens in the long run? Will item popularity reflect the intrinsic item quality, or can the process lead to (undesirable) configurations in which the most popular item is not the best?

Our system can be modeled as a special case of Pólya urn [30] with N colors, where $w_i[n]$ is the number of balls of color i after the n -th round. At competition round n , K balls are selected from the urn, and a new ball of the winning color is added to the urn. In contrast to classic Pólya urn models, the analysis here is complicated by the fact that the selection of K balls is a complex function of the current rank of the N colors in terms of their ball count. Indeed, recall that our system, mapped onto a Pólya urn, works as follows: we start with an empty set \mathcal{S} of pre-selected balls; at each of K iterations, a ball of color i is chosen with probability proportional to $r_i[n]^{-\alpha}$, where $r_i[n]$ is the rank of color i at the beginning of the round, and added to set \mathcal{S} . If j is the highest-quality color in set \mathcal{S} , a ball of color j is added to the urn at the end of the competition: $w_j[n] = w_j[n-1] + 1$, whereas $w_i[n] = w_i[n-1]$ for all $i \neq j$.

Moreover, we may or may not allow the repetition of colors (items) in the set of K balls. This distinction leads to two different pre-selection schemes, that we will call *with-item-repetition* and *without-item-repetition*. As the name suggests, *with-item-repetition* means that the K balls are independently selected with a probability that depends only on the ranking of the colors at the beginning of the round. Specifically, each of the K selected balls belongs to color i with probability

$$p_i[n] = \frac{r_i[n]^{-\alpha}}{\sum_j r_j[n]^{-\alpha}}. \quad (2)$$

Note that, by so doing, we could obtain more than one ball of a given color in our pre-selection.

In *without-item-repetition* pre-selection, we enforce that K balls of *distinct* colors are selected by re-normalizing the probabilities to choose a ball of a given color among the remaining colors. Let $p_{i,k}[n]$ be the probability to select the k -th ($k = 1, \dots, K$) ball among those of color i , and C_k the set of colors extracted in the first k choices, starting with $C_0 = \emptyset$. We have:

$$p_{i,k}[n] = \begin{cases} 0 & i \in C_{k-1} \\ \frac{r_i[n]^{-\alpha}}{\sum_{j:j \notin C_{k-1}} r_j[n]^{-\alpha}} & i \notin C_{k-1}. \end{cases} \quad (3)$$

The *without-item-repetition* pre-selection scheme more realistically describes a system that offers K distinct alternatives to the user (see example in Section 2.1). Unfortunately, this scheme is more difficult to analyze than *with-item-repetition*, as it requires to consider all dispositions of N elements over K positions. Conversely, *with-item-repetition* is easy to analyze, but it's less realistic, since items can appear multiple times in the pre-selection. After extensive numerical experiments, we discovered that almost all properties of the system operating under *without-item-repetition* are the same as those of the system operating under *with-item-repetition*. Moreover, as we will see, analytical results for the *with-item-repetition* scheme are very close to those numerically obtained for *without-item-repetition*, especially when $K \ll N$. For these reasons, in the following we will mainly focus on the *with-item-repetition* scheme.

We start with a general result valid for a generic pre-selection scheme based on the popularity rank of the items. Let \tilde{w}_i be the asymptotic fraction of balls of color i in the urn as the number of competitions tends to infinite:

$$\tilde{w}_i = \lim_{n \rightarrow \infty} \frac{w_i[n]}{\sum_j w_j[n]} = \lim_{n \rightarrow \infty} \frac{w_i[n]}{n + n_0}$$

where n_0 is the initial number of balls in the system.

Let $B : [0, 1]^N \rightarrow [0, 1]^N$ be the function that, given a vector $\tilde{\mathbf{w}}$ of normalized weights, provides the corresponding vector \mathbf{b} of winning probabilities. We will restrict ourselves to functions $B(\cdot)$ that depends on $\tilde{\mathbf{w}}$ only through the rank of the items.

ASSUMPTION 4.1 (RANK-BASED WINNING PROBABILITIES). *Function $B : [0, 1]^N \rightarrow [0, 1]^N$ is an injective function of the rank of weights $\tilde{\mathbf{w}}$.*

Definition 4.1 (System Stable Point). A system stable point is a stochastic vector $\mathbf{e} \in [0, 1]^N$ (a vector with non-negative entries that add up to one), which is a fixed point of function $B(\cdot)$, i.e., $B(\mathbf{e}) = \mathbf{e}$, and, in addition, $e_i \neq e_j$ for $i \neq j$.

PROPOSITION 4.1 (ASYMPTOTIC SYSTEM BEHAVIOR). *Under Assumption 4.1, as $n \rightarrow \infty$, normalized weights $\tilde{\mathbf{w}}[n]$ converge to a system stable point.*

PROOF. See Appendix B. □

Note that the system can have multiple stable points. Each stable point is associated to a particular permutation of the N items. For this reason, we will also say that a permutation is stable if it is associated to a system stable point.

We will call desirable equilibrium the stable point associated to the natural permutation $1, 2, \dots, N$ in which item popularity is perfectly aligned with item quality. We will call instead spurious equilibrium a stable point associated to a permutation different from the natural one.

The number and identity of stable points clearly depends on function $B(\cdot)$. In the *with-item-repetition* scheme, we have a simple expression for winning probabilities $\{b_i\}_i$. Indeed, denoting by $s_i = \sum_{j \leq i} p_j$ the cumulants of item selection probabilities (2) (we have dropped the dependency on n for simplicity), we can write, for $i > 1$:

$$b_i = s_i^K - s_{i-1}^K \quad (4)$$

(for $i = 1$, we have the special case $b_1 = p_1^K$).

Equation (4) can be understood as follows: for item i to be the highest-quality item selected among K independent choices with repetition we need that no higher-quality object has been selected for K times (term s_i^K); this event includes however also the event that color i has never been selected among the K choices (equivalently, objects of index up to $i - 1$ have been chosen K times), hence the probability s_{i-1}^K of this event has to be subtracted from previous event.




Fig. 2. Permutation graph for the system $N = 4$, $K = 2$, $\alpha = 1$, *with-item-repetition*.

In the *without-item-repetition* scheme, instead, the evaluation of $\{b_i\}_i$ requires, unfortunately, to enumerate all possible ordered sequences σ_i of K objects, containing item i and other $K - 1$ lower-index items (without repetition):

$$b_i = \sum_{\sigma_i} \prod_{k=1, \dots, K} p_{j,k} \quad (5)$$

$j = (\sigma_i)_k$

where $p_{j,k}$ are given by (3). In this case the computation of $\{b_i\}_i$ is feasible only for small values of N, K .

To find out all of the stable points of a system, we have followed different approaches. For small systems (say $N < 20$) one can simply test all $N!$ permutations, and check whether each of them is stable (according to Definition 4.1). To better illustrate the system dynamics, one can build, for a given choice of N, K, α the *permutation graph*, containing $N!$ nodes, one for each permutation, and add a direct edge from node i to node j if function $B(\cdot)$ maps permutation i into permutation j . Figure 2 shows the permutation graph for the system $N = 4$, $K = 2$, $\alpha = 1$, in the case *with-item-repetition*. In this case we have $4! = 24$ nodes, and 3 fixed points of function $B(\cdot)$, associated to permutations $\{1, 2, 3, 4\}$, $\{1, 2, 4, 3\}$ and $\{1, 4, 3, 2\}$. However, only $\{1, 2, 3, 4\}$ and $\{1, 2, 4, 3\}$ are stable points, because with $\{1, 4, 3, 2\}$ we have $b_2 = b_3$. Detailed simulations of the corresponding Pólya urn confirm that the actual system can only converge to the above two stable points.

The permutation graph can also be used to define the *attractiveness* of a given fixed point of function $B(\cdot)$, i.e., of a permutation f such that $B(f) = f$. The attractiveness $a(f)$ of fixed point f is defined as the size of the (weakly) connected component containing f , divided by the number $N!$ of all permutations. Note that $\sum_{f: B(f)=f} a(f) = 1$. For example, in the scenario of Figure 2 we have $a(\{1, 2, 3, 4\}) = \frac{1}{2}$, $a(\{1, 2, 4, 3\}) = \frac{1}{3}$, and $a(\{1, 4, 3, 2\}) = \frac{1}{6}$.

For large systems, we cannot examine all possible permutations. To overcome this problem, one approach is to perform a randomized search of stable points by starting from an arbitrary stochastic vector \tilde{w} , i.e., some random positive values $\{\tilde{w}_i\}_i$ summing up to 1, and iteratively apply function $B(\cdot)$ until we hit a fixed point. This procedure is described in detail in Algorithm 1. In the permutation graph, this means to start from a random node, and follow the direct path up to the node with the self-loop. One problem of this approach is that we need to perform a number of searches large enough to fall at least once in each connected component of the permutation graph, and therefore we are not guaranteed to discover all of the stable points.

ALGORITHM 1: Randomized computation of equilibria**Require:** $N \in \mathbb{N}^+$, $K \in \{1, \dots, N\}$, $\alpha \in \mathbb{R}^+$

```

1: Choose random positive values  $\{\tilde{w}_i\}_{i=1}^N : \sum_{i=1}^N \tilde{w}_i = 1$  ▷ Initialization of  $w_i$ 's
2:  $fixedpoint \leftarrow false$ 
3: while ( $fixedpoint = false$ ) do
4:    $r_i \leftarrow$  index of  $i$  in sorted  $\tilde{w}$ ,  $r_i \in \{0, \dots, N-1\}$ 
5:    $p_i \leftarrow (N - r_i)^{-\alpha}, \forall i$  ▷ popularity of  $i$ 
6:   normalize  $p_i$ 's such that  $\sum_{i=1}^N p_i = 1$ 
7:   compute  $b_i = \mathbb{P}[\text{object } i \text{ wins}], \forall i$  ▷ using either (4) or (5)
8:   if  $b = \tilde{w}$  then
9:      $fixedpoint \leftarrow true$ 
10:  else
11:     $\tilde{w} \leftarrow b$ 
12:  end if
13: end while
14: if  $b_i \neq b_j, \forall i \neq j$  then
15:    $b$  is a stable point
16: end if

```

As we will see, in most cases the only stable permutations are those that, starting from the natural permutation, change the order of the top M items, with $M \ll N$. Therefore, in practice one can find all stable points by sequentially exploring all permutations generated in lexicographic order, and stopping the search when no more stable points are found comprising swaps of the top $M + 1$ items. The random search performed by Algorithm 1 is then to be considered a last resort, when exact approaches become unfeasible (N or M too large). We emphasize that the above computational issue pertains only the analytical prediction of system stable points, while the actual algorithm implemented by the recommendation system, which produces just randomized lists biased by popularity (see toy example in Section 2.1), can be applied to arbitrarily large N without scalability problems, but being oblivious of where it is going to converge.

As another interesting example, consider the *without-item-repetition* scheme with $\alpha = 1$. Figure 3 shows, on a logarithmic vertical scale, the number of stable points as a function of N , for $K = 1, 2, 3, 4$. Note that for $K = 1$ all $N!$ permutations are feasible. As we increase K , the number of stable points is brought down drastically, though it still increases exponentially for large N . For example, in the case of $N = 10$ and $K = 2$, there are four stable points, corresponding to permutations:

$\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}, \{1, 2, 3, 4, 5, 6, 7, 8, 10, 9\}, \{1, 2, 3, 4, 5, 6, 7, 9, 8, 10\}, \{1, 2, 3, 4, 5, 6, 7, 10, 9, 8\}.$

Notice that, in addition to the (desired) natural permutation, another three spurious permutations can emerge in which some of the highest-quality items are permuted (the top two or top three items).

Interestingly, for given N , if we take K large enough we obtain a single stable point. We will see that, when the stable point is unique, it corresponds to the natural permutation, i.e., the desired equilibrium. For example, $K = 3$ is enough to obtain as stable point only the natural permutation for all values of N up to 19. As another case, $K = 4$ is required to achieve only the desired equilibrium with $N = 50$ items. In the next subsection we will precisely address the problem of determining when a given system (N, K, α) can only achieve the desired equilibrium, while no spurious permutations can emerge from the competition among items.




Fig. 3. Number of stable points for *without-item-repetition*, $\alpha = 1$.

4.1 Computation of K_{\min}

We define K_{\min} as the minimum value of K that admits the natural permutation $1, \dots, N-1, N$ as the only stable permutation. By virtue of the following lemma, to compute K_{\min} it is sufficient, in most cases, to determine whether the alternate permutation $1, \dots, N, N-1$, in which the two highest-quality objects are permuted, is stable.

LEMMA 4.1 (CRITICAL PERMUTATION). *Under popularity-biased pre-selection, the minimum value K_{\min} is determined by the stability of the critical permutation $1, \dots, N, N-1$.*

PROOF. See Appendix C. □

We first consider the case *with-item-repetition*, for which K_{\min} can be characterized exactly. Suppose that the current permutation, induced by increasing weights, is the critical permutation $1, \dots, N, N-1$. Then, we have $p_N = 2^{-\alpha}/G$, $p_{N-1} = 1/G$, where $G = \sum_{i=1}^N i^{-\alpha}$ is the normalization constant of the power-law popularity bias. We also introduce the cumulants $s_N = 1$, $s_{N-1} = 1 - p_N$, $s_{N-2} = 1 - p_N - p_{N-1}$, with which we can write the winning probabilities of the top two objects as:

$$\begin{aligned} b_N &= s_N^K - s_{N-1}^K \\ b_{N-1} &= s_{N-1}^K - s_{N-2}^K. \end{aligned}$$

The critical permutation is not stable when $b_{N-1} < b_N$, i.e., when

$$2s_{N-1}^K < s_N^K + s_{N-2}^K. \quad (6)$$

Substituting the cumulants, the above inequality becomes:

$$2(1 - G^{-1}2^{-\alpha})^K \leq 1 + (1 - G^{-1} - G^{-1}2^{-\alpha})^K \quad (7)$$

which provides, as function of G and α , the condition that K must satisfy so that the critical permutation is not stable. Therefore, by Lemma 4.1 we have:

$$K_{\min} \triangleq \min_K \left\{ K : 2(1 - G^{-1}2^{-\alpha})^K \leq 1 + (1 - G^{-1} - G^{-1}2^{-\alpha})^K \right\}. \quad (8)$$




Fig. 4. Minimum value of K such that the natural permutation is the only stable point, for different values of N , *with-item-repetition*.

The following proposition provides a summary of interesting properties of K_{\min} as function of parameters N, α . In particular, since the catalog size can in some cases be very large, it is interesting to see what happens when $N \rightarrow \infty$, though the item set is to be considered finite in any feasible scenario.

PROPOSITION 4.2 (PROPERTIES OF K_{\min}). *For any given $N \geq 1, \alpha \geq 0$, there exists an integer value $K_{\min} \geq 1$ such that for any $K \geq K_{\min}$ the only stable permutation is the natural one. For fixed α , K_{\min} is a non-decreasing function of N . As a special case, $K_{\min} = 2$ for $\alpha = 0$. When $0 < \alpha \leq 1$, $K_{\min} \rightarrow \infty$ as $N \rightarrow \infty$. For fixed $\alpha > 1$, $K_{\min} \rightarrow K_{\min}^{\infty}$ as N increases, where K_{\min}^{∞} is a constant which depends on α .*

PROOF. See Appendix D. □

In the following, we will see that properties of K_{\min} listed in Proposition 4.2 hold also under *without-item-repetition*. The following result instead is specific to the case *with-item-repetition*:

COROLLARY 4.1. *For given $N > 1$, in the case *with-item-repetition*, $K_{\min} \rightarrow \infty$ as $\alpha \rightarrow \infty$.*

PROOF. We can simply upper bound the right-hand side of (7), for any α and K , by $2 - G^{-1}$, obtaining the weaker inequality:

$$2 \left(1 - \frac{1}{G2^{\alpha}} \right)^K \leq 1 + \left(1 - \frac{1}{G} - \frac{1}{G2^{\alpha}} \right)^K \leq 2 - \frac{1}{G}$$

which is satisfied by

$$K \geq \frac{\log(1 - (2G)^{-1})}{\log(1 - (G2^{\alpha})^{-1})} \quad (9)$$

The assert follows by noticing that the right-hand side of (9) tends to infinity as $\alpha \rightarrow \infty$. □

Figure 4 shows K_{\min} as function of α , for different values of N , obtained by solving numerically (8). Results in Figure 4 confirm the properties in Proposition 4.2. In particular, K_{\min} becomes unbounded, as N increases, for $0 < \alpha \leq 1$. As predicted by Corollary 4.1, for large α the value K_{\min} (slowly) increases with α .

Remark 4.1. When $K < K_{\min}$, the system is not guaranteed to produce, in the long run, the desired alignment between quality and popularity, and can end up operating under one of many suboptimal configurations containing severe misalignments in the top of the list. The frequencies with which we can observe such undesirable operating points depend crucially on the distribution $\mathbf{w}[0]$ of initial weights. Extensive numerical simulations reveal that undesirable configurations are not at all ‘rare events’, and show up with frequency comparable to that of the natural configuration, even when we start with all weights initially equal to 1. If we start with random initial weights, e.g., an i.i.d weight for each item, we increase the likelihood to ‘crystallize’ initial suboptimal configurations, and correspondingly reduce the probability to achieve the desired alignment.

Remark 4.2. Proposition 4.2 provides some fundamental insights into the behavior of popularity-biased cultural markets. Specifically, it suggests there exists a harmful regime of mild popularity bias ($0 < \alpha \leq 1$) in which desired alignment of popularity and quality requires large quality discrimination power on behalf of users. Since in this regime K_{\min} diverges as N increases, systems with large catalog size are essentially unpredictable in this case, and can possibly produce severe misalignments among high-quality items. Surprisingly, the special case $\alpha = 0$ (uniform preselection) behaves as a singularity, requiring the minimum discrimination power $K_{\min} = 2$ for any N .

4.1.1 Extensions to Heterogeneous Discrimination Power. Previous analysis of K_{\min} can be extended to account for the presence of heterogeneous users with different behavior. One easy extension is to consider a fraction f_m of users deterministically choosing the current most popular item, while the remaining fraction $1 - f_m$ of users employ quality discrimination power K . In this case, under the crucial permutation we have:

$$\begin{aligned} b_N &= (1 - f_m) (s_N^K - s_{N-1}^K) \\ b_{N-1} &= f_m + (1 - f_m) (s_{N-1}^K - s_{N-2}^K) \end{aligned}$$

and condition $b_{N-1} < b_N$ becomes:

$$f_m + (1 - f_m) \left[2 \left(1 - \frac{1}{G^{2\alpha}} \right)^K - 1 - \left(1 - \frac{1}{G} - \frac{1}{G^{2\alpha}} \right)^K \right] < 0.$$

We observe that the left-hand side is positive for $K = 1$, and becomes negative for sufficiently large K provided that $f_m < 1/2$ (the term in square brackets tends to -1 as $K \rightarrow \infty$). Therefore, $K_{\min} \rightarrow \infty$ as $f_m \rightarrow 1/2$, and no K_{\min} exists for $f_m \geq 1/2$. With fixed $f_m < 1/2$, all properties stated in Proposition 4.2 are still valid, as one can verify along the same lines of the proof reported in Appendix D. Moreover, K_{\min} is an increasing function of f_m , which can be proven by applying the implicit function theorem, and noticing that $\frac{\partial F}{\partial f_m} = 1$ at the points in which $F(x, f_m) = 0$ (see Appendix D).

As another extension we can consider the case of users with i.i.d. values of K , distributed according to the discrete law $\{p_k\}_k$. Asymptotic stability results in Proposition 4.1 still apply upon substituting the vector of winning probabilities \mathbf{b} , which was implicitly depending on fixed parameter K , with its average with respect to K , i.e., $\mathbb{E}_K[\mathbf{b}(K)]$.

Moreover, one can check whether a given distribution of K can only lead to the natural permutation. In particular, the critical permutation $1, \dots, N, N-1$ is not stable provided that

$$\sum_k p_k \left[2 \left(1 - \frac{1}{G^{2\alpha}} \right)^k - 1 - \left(1 - \frac{1}{G} - \frac{1}{G^{2\alpha}} \right)^k \right] < 0 \quad (10)$$

and one can numerically check whether (10) indeed holds under a given distribution $\{p_k\}_k$.

4.1.2 Extension to Without-item-repetition. The case *without-item-repetition* can in principle be handled in the same way as *with-item-repetition*, but unfortunately the exact computation of winning probabilities b_i requires now the enumeration of all dispositions of K objects from N objects, whose number becomes intractable for large values of N, K . Therefore, we propose the following approximate computation of K_{\min} for the case *without-item-repetition*, which turns out to be very accurate when compared to exact (computationally feasible) results. The approximation is based on the following idea: we assume that the top two objects are extracted without repetition, whereas all other objects can be extracted with repetition. By so doing, we capture the fact that the two most important objects for our purposes are correctly extracted without repetition, while all of the others, whose precise identity is not important to determine winning probabilities b_N, b_{N-1} , are approximately extracted with repetition to allow scalable computation. Let $p^* = 1 - p_N - p_{N-1}$ be the probability to choose a ball within set $\{1, \dots, N-2\}$.

Given an extraction of K balls, we have that neither object N nor object $N-1$ wins with probability $(p^*)^K$. Then, if we determine b_{N-1} , we can easily obtain the complementary $b_N = 1 - (p^*)^K - b_{N-1}$. Therefore, $b_N \geq b_{N-1}$ is equivalent to $b_{N-1} \leq \frac{1-(p^*)^K}{2}$. To compute the winning probability of object $N-1$, we consider that object $N-1$ wins when: m objects different from $N, N-1$ are initially extracted, $0 \leq m \leq K-1$ (and put back in the urn); object $N-1$ is extracted (and not put back in the urn); $N-1-m$ objects different from N are extracted with *renormalized* probability $p^*/(1-p_{N-1})$.

It follows:

$$b_{N-1} \sim \sum_{m=0}^{K-1} (p^*)^m p_{N-1} \left(\frac{p^*}{1-p_{N-1}} \right)^{K-m-1} = (p^*)^{K-1} \frac{1 - (1-p_{N-1})^K}{(1-p_{N-1})^{K-1}}$$

and we obtain the approximate value of K_{\min} :

$$K_{\min} = \min_K \left\{ K : 2 \left(\frac{1 - G^{-1} - G^{-1}2^{-\alpha}}{1 - G^{-1}} \right)^{K-1} \leq \frac{1 - (1 - G^{-1} - G^{-1}2^{-\alpha})^K}{1 - (1 - G^{-1})^K} \right\}. \quad (11)$$

Under the above approximation of K_{\min} for the case *without-item-repetition*, all properties listed in Proposition 4.2 still holds, as one can check along the same lines of the proof reported in Appendix D. The fundamental difference with respect to *with-item-repetition* is the behavior of K_{\min} as α grows large:

COROLLARY 4.2. *For given $N > 1$, in the case without-item-repetition, under approximation (11), $K_{\min} \rightarrow 2$ as $\alpha \rightarrow \infty$.*

PROOF. The right hand side of (11) is larger than one for any α . The left-hand side of (11) is equivalent to $2(1 - \frac{1}{(G-1)2^\alpha})^{K-1}$. Since $(G-1)2^\alpha = 1 + \sum_{i=3}^{\infty} (\frac{2}{i})^\alpha$ tends to 1 as $\alpha \rightarrow \infty$, the left hand side of (11) tends to zero as $\alpha \rightarrow \infty$, hence $K = 2$ is enough to make the inequality valid for large α . \square

Figure 5 shows K_{\min} as a function of α , for different values of N , obtained either by the approximate inequality (11) (thick curves), or by exact numerical computation of all stable points, which is feasible only for small values of N, K (we were not able to obtain exact results beyond $N = 30$). A number of interesting observations can be drawn from Figure 5: (i) the approximate formula (11), which can scale to arbitrarily large values of N, K , gets very close to exact values, except for little discrepancies around the points at which K_{\min} varies by 1; (ii) properties in Proposition 4.2 are indeed verified also for *without-item-repetition*, and in particular K_{\min} still diverges for $0 < \alpha \leq 1$;




Fig. 5. Minimum value of K such that the natural permutation is the only stable point, for different values of N , *without-item-repetition*. Thin curves refer to approximation (11), whereas thick curves show exact results (when computation is feasible).

(iii) for fixed N and α , K_{\min} is smaller under *without-item-repetition* with respect to *with-item-repetition*, as one can check by comparison with Figure 4; and (iv) in contrast to *with-item-repetition*, $K_{\min} = 2$ for large α , as predicted by Corollary 4.2.

Remark 4.3. Results in Figure 5 confirms the existence of a harmful regime of mild popularity bias ($0 < \alpha < 1$) also in the more realistic case *without-item-repetition*. To give the reader an intuitive explanation of this phenomenon, consider the case $K = 2$. Denote with H the highest-quality object, M the second highest-quality object, and L a generic object of lower quality. The critical permutation in which M becomes more popular than H can be a system stable point for N sufficiently large and, say, α around 0.5. Indeed, suppose to start already in this configuration. Then H will struggle to regain popularity versus M , because pre-selections of type $\{M, L\}$ (note that there are many different choices of low-quality item L), which reinforce the popularity of M vs H , become jointly more likely than those of type $\{M, H\}$ or $\{H, L\}$, which would reinforce the popularity of H .

Note however that, in the extreme case of $\alpha = 0$, any pre-selection of the N objects is equally likely to be generated, and those in which H wins are more numerous than those in which M wins (indeed, H wins also in the direct match $\{M, H\}$). Hence, for α sufficiently small we expect that the only stable point will be the natural permutation.

For large α , the stability of the critical permutation depends on the fact that item repetition is allowed or not. First of all, note that for large α the probability to pre-select an item is concentrated on the most popular items. In *without-item-repetition*, for large α the direct match $\{M, H\}$ will appear more and more frequently, while pre-selections containing also low quality objects will become negligible, so the critical permutation cannot be stable, and K_{\min} tends to 2. Indeed, consider an extremely large α , and suppose to start from the critical permutation: after extracting M , M can no longer be extracted, and we have to choose as second object H , which wins over M . In *with-item-repetition*, the opposite behavior occurs for very large α : now almost all extractions of the first two objects will both be M , reinforcing M itself. Actually for any K there will be an α large enough that with high probability we extract M consecutively K times, making the critical

permutation stable. Hence, *with-item-repetition* $K_{\min} \rightarrow \infty$ as $\alpha \rightarrow \infty$ (Corollary 4.1). Of course, in a real system that randomizes the list of items proposed to the user objects are never repeated, so *without-item-repetition*, though more difficult to analyze exactly, is more meaningful.

Example 4.1. As an example of possible application of our results, consider a Video-On-Demand platform offering 100 titles in response to queries for a given movie genre. Suppose that the system shows to users a randomized list of the above titles, using as popularity metric the current number of views. If we want the best-quality movie (assuming that this notion exists for movies) to always emerge from the competition with other movies of the same genre (which occurs when the natural permutation is the only stable point), the popularity-based amount of randomness introduced in the generation of the recommendation list (the value of α in our ranking model) must be carefully controlled. For example, from Figure 5 we see that values of α around 0.5 are a very bad choice, unless the quality discrimination power of the users is large ($K \geq 9$).

5 ANALYSIS OF THE AVERAGE QUALITY INDEX

The performance of an online cultural market depends on a complex combination of factors including the way items are internally ranked, the way in which they are visually presented to the user in response to a given query, and the user behavior in exploring and selecting among alternatives. In our simple model, the system has been described by tuple (N, K, α) . Among this set, α can be considered as a parameter tunable by the online platform, and one can naturally ask whether an optimal value α^* can be computed so as to maximize a given performance metric.

In the previous section, we have considered as primary objective the guaranteed emergence of the natural alignment between popularity and quality. As a secondary objective, we consider here another important optimization criterion, which is the average quality of items ultimately chosen by users. Indeed, it might not be desirable to just maximize such average quality, i.e., to make the top-quality item monopolize the market. In general, a system might prefer to achieve a desired level of ‘fairness’ among the different alternatives, so that also mid-quality items have non-negligible chances to be chosen.

Note that in this section we will analyze the average quality index assuming that the system has a single stable point (the natural alignment). In Section 6 we will extend the definition of average quality index to the case in which there are multiple stable points, when we will introduce the extension of the model to multiple classes of users.

5.1 Obtaining a Desired Average Quality

Assuming that $K \geq K_{\min}$, so that the system is guaranteed to achieve the natural alignment, the average quality index \bar{q} takes value in the range $[1, N]$. Since the extreme case $\bar{q} = N$, corresponding to a system in which only the top-quality item is chosen, is in general undesirable, we assume that our goal is to achieve a given, intermediate value $1 \leq \bar{q} \leq N$.

To understand whether a given \bar{q} is indeed feasible, we need to compute how metric \bar{q} associated to the natural permutation depends on system parameters (N, K, α) , separately considering the cases *with-item-repetition* and *without-item-repetition*.

PROPOSITION 5.1 (AVERAGE QUALITY INDEX UNDER WITH-ITEM-REPETITION). *For given N and fixed K , the average quality \bar{q} of the natural permutation is an increasing function of α . Possible values of \bar{q} lie in the range $[\bar{q}_{\min}, N)$, where \bar{q}_{\min} is the value attained with $\alpha = 0$:*

$$\bar{q}_{\min} = N - \frac{\sum_{i=1}^{N-1} i^K}{N^K}$$

while $\bar{q} \rightarrow N$ as $\alpha \rightarrow \infty$.

For given N and α , the average quality \bar{q} of the natural permutation is an increasing function of K .

PROOF. See Appendix E. \square

Under *without-item-repetition*, a formal proof of the monotonicity of \bar{q} with respect to α and K is more difficult, since winning probabilities b_i 's lack a simple closed form expression. However, we have numerically verified that \bar{q} indeed increases with α and K , and results analogous to those stated in Proposition 5.1 holds under *without-item-repetition*. However, the lower extreme of possible values of \bar{q} is different, since here, for $\alpha = 0$, we get:

$$\bar{q}_{\min} = \sum_{i=K}^N i \cdot \frac{\binom{i-1}{K-1}}{\binom{N}{K}} = \frac{\sum_{i=K}^N K \binom{i}{K}}{\binom{N}{K}} = K \frac{\binom{N+1}{K+1}}{\binom{N}{K}} = K \frac{N+1}{K+1}. \quad (12)$$

It is interesting to see how \bar{q}_{\min} in (12) depends on K : with $K = 1$ we obtain $\frac{N+1}{2}$, which is the baseline performance of a system in which quality does not come into play, and any item is equally likely to win the competition.³ For $K = N$, the top-quality item always win, and we obtain, as expected, the maximum possible value N .

As immediate consequence of Proposition 5.1, for given N and K , any desired average quality index $\tilde{q} \in [\bar{q}_{\min}(N, K), N)$ can be obtained by a unique choice of $\alpha^* \in \mathbb{R}^+ \cup 0$, where we have specified for clarity the dependency of \bar{q}_{\min} on both N and K .

5.2 Approaching a Desired Winning Probability Distribution

Instead of the average quality index, one could try to obtain a specific winning probability \tilde{b}_i for each item, assuming that $\{\tilde{b}_i\}_i$ form an increasing sequence with i , corresponding to a system operating under the natural permutation.

As an example, we consider a family of desired winning probabilities $\{\tilde{b}_i\}_i$ specified by a rank-based power-law of exponent $\beta > 0$:

$$\tilde{b}_i \triangleq \frac{(N-i+1)^{-\beta}}{\sum_{j=1}^N j^{-\beta}} \quad [\text{desired winning probability}]. \quad (13)$$

Here, β is a pre-defined exponent reflecting the desired level of fairness: when $\beta = 0$ we have the extreme case in which all items have the same winning probability, irrespective of their intrinsic quality; for $\beta \rightarrow \infty$ we approach the other extreme case in which only item N wins. Correspondingly, we have the desired average quality index $\tilde{q} \triangleq \sum_{i=1}^N i \cdot \tilde{b}_i$.

Since winning probabilities \tilde{b}_i cannot be perfectly achieved, we consider as optimal α^* the value of α that minimizes a suitable distance metric between probability distributions $\{b_i\}_i$ and $\{\tilde{b}_i\}_i$. We will also consider the difference:

$$\Delta q \triangleq |\tilde{q} - \bar{q}| \quad (14)$$

between the desired average quality index and the actual average quality index.

5.3 An Optimization Example

As an example, we consider an online cultural market containing $N = 20$ items where the user discrimination parameter is $K = 5$. Suppose that our goal is to approach the winning probabilities $\{\tilde{b}_i\}$ obtained by using $\beta = 2$ in (13), i.e., the second-best item has winning probability equal to 1/4 of the winning probability of the best item.

³Recall that with $K = 1$ any permutation is equally likely to emerge, but since $b_i = 1/N$ for any item, the specific permutation is not important.




Fig. 6. Distance between probability distributions $\{b_i\}$ and $\{\tilde{b}_i\}$ as a function of α .

From Figure 5 we observe that, in the case *without-item-repetition*, $K = 5$ is enough to guarantee emergence of the natural permutation for *any* α , so here we do not need to check whether we indeed have $K > K_{\min}$.⁴

Figure 6 shows the results of our numerical optimization. We plot the value of three distance metrics between discrete probability distributions:

$$\begin{aligned} & \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^N \left(\sqrt{b_i} - \sqrt{\tilde{b}_i} \right)^2} && [\text{Hellinger distance}] \\ & \sum_{i=1}^N b_i \log \left(\frac{b_i}{\tilde{b}_i} \right) && [\text{Relative entropy}] \\ & -\ln \left(\sum_{i=1}^N \sqrt{b_i \tilde{b}_i} \right) && [\text{Bhattacharyya distance}] \end{aligned}$$

as functions of α . Interestingly, all considered distances between probability distributions $\{b_i\}$ and $\{\tilde{b}_i\}$ reach their minimum value for α^* close to 0.58. At the same time, there exists a unique value for α ($\alpha^* = 0.40$) that nullifies Δq , as a consequence of Proposition 5.1. Note that both values of α^* computed above fall in the harmful regime $0 < \alpha \leq 1$: though we can safely operate in this regime with just 20 items, doing so with larger values of N could be unfeasible, since K_{\min} would become too large (see Figure 5).

6 HETEROGENEOUS PERCEPTION OF QUALITY

So far we have assumed the existence of an intrinsic quality metric of each item, such that all users would rank the N items in the same order (provided that they had the time to independently examine all of them). Since quality is highly subjective, and its assessment may vary among users, we now extend the model to the case of a (finite) number C of user classes: the user arriving at time n is assumed to belong to class $1 \leq c \leq C$ with probability f_c , independently of other users. Of course, we have the constraint $\sum_{c=1}^C f_c = 1$.

⁴In general, α^* has to be found in the restricted set of α values for which $K > K_{\min}$.

Users belonging to the same class are assumed to perceive items' quality in such a way that they would all rank the items in the same order (provided that they had the time to independently examine all of them). Therefore, user class c is characterized by a unique (fixed) vector \mathbf{v}_c , listing the items' id's in increasing order of their quality, as perceived by users belonging to that class. Note that in the base setting there is a single class ($C = 1$), with $\mathbf{v}_1 = [1, 2, \dots, N]$. Our model is quite general, since we do not impose any restriction on vectors \mathbf{v}_c , that could differ completely from one class to another. For example, users belonging to a class may perceive an item as high-quality, while users belonging to another user may perceive the same item as low-quality.

It turns out that the multi-class version of the system can be analyzed essentially in the same way as the single-class instance, by substituting the winning probability of each item by a weighted average (where weights correspond to probabilities $\{f_c\}$) of the winning probabilities computed for the same item according to the different classes. Therefore, we can introduce a generalized mapping function $B(\cdot)$ as done in Section 4 and easily extend to the multi-class case the Definition 4.1 of system stable point.

The numerical computation of system stable points can thus be performed in a similar way as for the single-class setting. Algorithm 2 is a simple and quite straightforward generalization of Algorithm 1 to the case of multiple user classes.

ALGORITHM 2: Randomized computation of equilibria with multiple users classes

Require: $N \in \mathbb{N}^+$, $K \in \{1, \dots, N\}$, $\alpha \in \mathbb{R}^+$

- 1: Choose random positive values $\{\tilde{w}_i\}_{i=1}^N : \sum_{i=1}^N \tilde{w}_i = 1$ ▷ Initialization of w_i 's
 - 2: $fixedpoint \leftarrow \text{false}$
 - 3: **while** ($fixedpoint = \text{false}$) **do**
 - 4: $r_i \leftarrow \text{index of } i \text{ in sorted } \tilde{\mathbf{w}}, r_i \in \{0, \dots, N-1\}$
 - 5: $p_i \leftarrow (N - r_i)^{-\alpha}, \forall i$ ▷ popularity of i
 - 6: normalize p_i 's such that $\sum_{i=1}^N p_i = 1$
 - 7: $\mathbf{b} \leftarrow \mathbf{0}$
 - 8: **for** $c=1$ **to** C **do**
 - 9: $p_{i,c} \leftarrow p_{\mathbf{v}_c[i]}, i \in \{1, \dots, N\}$ ▷ popularity of objects, ordered by quality, according to class c
 - 10: compute $b_{i,c} = \mathbb{P}[\text{object } i \text{ wins}], i \in \{1, \dots, N\}$ ▷ using either (4) or (5), and $p_{i,c}$ in place of p_i
 - 11: $b_{\mathbf{v}_c[i]} \leftarrow b_{\mathbf{v}_c[i]} + f_c b_{i,c}, i \in \{1, \dots, N\}$ ▷ contribution of class c to winning probability of object i
 - 12: **end for**
 - 13: **if** $\mathbf{b} = \tilde{\mathbf{w}}$ **then**
 - 14: $fixedpoint \leftarrow \text{true}$
 - 15: **else**
 - 16: $\tilde{\mathbf{w}} \leftarrow \mathbf{b}$
 - 17: **end if**
 - 18: **end while**
 - 19: **if** $b_i \neq b_j, \forall i \neq j$ **then**
 - 20: \mathbf{b} is a stable point
 - 21: **end if**
-

Unfortunately, instead, the analytical results on K_{\min} derived in Section 4.1 no longer apply to the multi-class case, since they are crucially based on the assumption that there exists a unique

Table 2. Items Ordered by Increasing Quality, According to 3 Users Classes

| | | | | | | | | | | |
|----------------------------|---|---|---|---|---|---|---|---|---|---|
| (class 1) \mathbf{v}_1 : | 5 | 7 | 4 | 9 | 8 | 1 | 2 | 6 | 3 | 0 |
| (class 2) \mathbf{v}_2 : | 3 | 1 | 2 | 5 | 9 | 7 | 4 | 8 | 0 | 6 |
| (class 3) \mathbf{v}_3 : | 7 | 9 | 0 | 4 | 2 | 3 | 8 | 1 | 5 | 6 |

quality ranking for the items. Given the wide variety of multi-class scenarios that can be considered (notice that, beyond the number C of classes, detailed characteristics of each class as specified by vector \mathbf{v}_c can affect the transient and asymptotic system behavior), we have limited ourselves to a preliminary investigation of just a few scenarios.

As an example, we have considered the case of $N = 10$ items, with at most three user classes whose characteristics are summarized in Table 2: for each class c , $c = 1, 2, 3$, the id's of the items, ordered by quality as perceived by users of class c (i.e., vector \mathbf{v}_c), are listed from the lowest-quality (left-most item) to the highest quality (right-most item). Note that vectors \mathbf{v}_c were chosen simply at random from the space of $10!$ permutations.

Figure 7 reports, on a log vertical scale, the number of system stable points as function of K , for four different values of $\alpha = 0, 0.5, 1, 2$, comparing three systems with $N = 10$, *with-item-repetition*: a system with $C = 1$ (only class 1, left plot); a system with $C = 2$ (classes 1 and 2, middle plot); a system with $C = 3$ (classes 1, 2, and 3, right plot). In all scenarios we assume that classes are equally represented (i.e., $\forall c, f_c = \frac{1}{C}$).

Note that the system with $C = 1$, equivalent to the single-class case, confirms results in Figure 4, according to which $K_{\min} = 2$ for $\alpha = 0$, $K_{\min} = 3$ for $\alpha = 0.5$, $K_{\min} = 4$ for $\alpha = 1$ or $\alpha = 2$. When $K < K_{\min}$, it is interesting to observe how the number of stable points (equal to $10!$ for $K = 1$) is drastically reduced by increasing K . A similar strong impact of K on the number of stable points is observed also in the other two cases with multiple classes, but the decay is slower, suggesting that systems with wildly different user classes are much more difficult to stabilize. Moreover, while the system with $C = 3$ can actually reach a unique configuration for K sufficiently large (except for $\alpha = 2$), the system with $C = 2$ always leads to multiple stable points, except for $\alpha = 0$, for which a unique configuration is stable provided that $K \geq 2$ (this special behavior for $\alpha = 0$ holds in general).

Interestingly, for $C = 2$, $\alpha = 0.5$, the number of stable points is even non-monotonic with K . This anomaly prompted us to investigate also the behavior of the average quality index, Figure 8. In the presence of multiple stable point, the overall average quality index \bar{Q} is computed as the weighted sum of average quality indexes of individual stable points, weighted by their attractiveness:

$$\bar{Q} = \sum_f a(f) \bar{q}(f)$$

Recall from Section 4 that attractiveness is defined as the number of permutations falling into a given fixed point by recursive applications of function $B(\cdot)$. We use it as a proxy of the likelihood that the system converges to a given equilibrium point, which is unfortunately hard to characterize since it depends on initial weights $\mathbf{w}[0]$.

Interestingly, \bar{Q} always increases with K , suggesting, in line with intuition, that stronger quality-discrimination of users leads to better average quality of selected items, also in the multi-class case. However, as we increase the number of classes, achievable values of \bar{Q} tend to decrease, as consequence of the fact that a certain stable point (unique or not) inevitably cannot make happy users belonging to many (very different) classes.

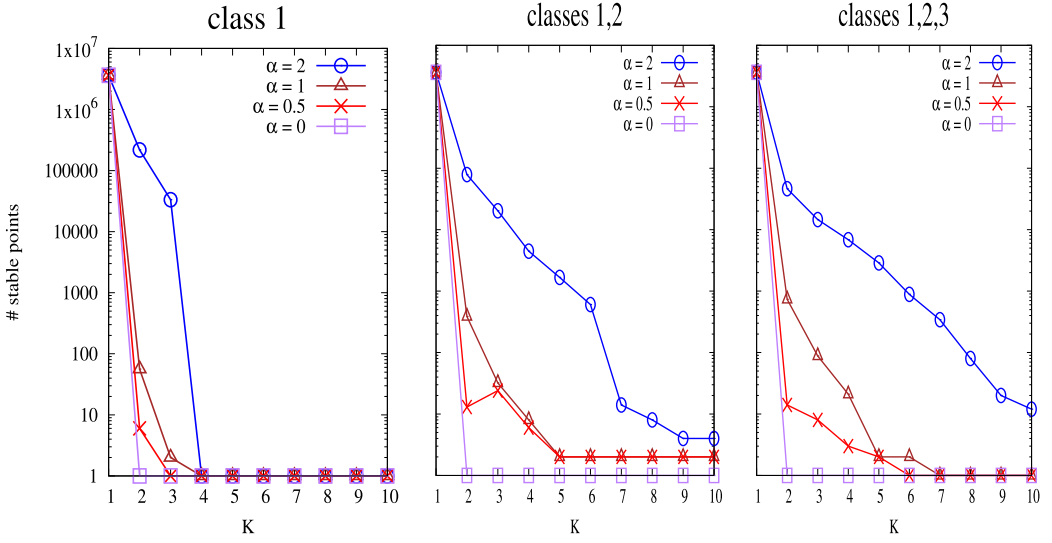


Fig. 7. Number of system stable points as a function of K , for different values of α , in the case of $N = 10$, *with-item-repetition*: with only class 1 (left plot); with classes 1,2 (middle plot); with classes 1,2,3 (right plot).




Fig. 8. Average overall quality index \bar{Q} as function of K , for different values of α , in the case of $N = 10$, *with-item-repetition*: with only class 1 (left plot); with classes 1,2 (middle plot); with classes 1,2,3 (right plot).

At last, we performed also simulations of the above multi-class scenario, to derive performance metrics that cannot be easily derived analytically, not even for the single-class case. One such metric is the distribution of the smallest time for the system to reach one of its stable points. We assume that the system starts from the initial condition in which popularity weights are all equal to 1, i.e., $\forall i, w_i[0] = 1$. We simulated 10^6 run for each scenario, as we vary the number C of classes and parameter α .




Fig. 9. Median of the minimum time to reach a system stable point as function of α , for different number of classes.




Fig. 10. Evolution of normalized weights $\tilde{w}_i[n]$ as function of time, in the case $C = 3$, $K = 5$, $\alpha = 0.5$.

Figure 9 shows the median of the distribution of the smallest time to reach any stable point, as function of parameter α , for $N = 10$, $K = 5$, *without-item-repetition*. We observe that, as a general trend, the stronger the popularity bias represented by α , the longer it takes for the system to reach one of its stable points. One exception is in the case $C = 3$ (note that the median diminishes passing from $\alpha = 0.5$ to $\alpha = 0.6$): this can be explained by the fact that, in the interval $\alpha \in [0.6, 0.8]$, the system has two stable points instead of 1 (intuitively, the higher the number of system stable points the lower the time to reach any of them).

Figure 10 shows, on a log-log scale, the evolution over time of normalized weights $\tilde{w}_i[n]$ in a particular simulation run obtained for one of the scenarios considered in Figure 9, i.e., $C = 3$, $\alpha = 0.5$. We observe that, after a turbulent initial phase characterized by large randomness affecting mid-popularity objects, trajectories converge to their asymptotic constant value, becoming almost flat after about 10^6 user interactions.

7 RELATED WORK

Our work studies the consequences of using popularity as proxy for quality. There has been work supporting this view. In [48], for example, authors consider two different social news aggregators, Reddit and Hacker News. They define quality as the number of votes an article would have received if shown, in a bias-free way, to an equal number of users. Using a Poisson regression method they find that popularity on Reddit and Hacker News is a relatively strong reflection of quality. It is therefore to be expected that, especially in condition of limited attention, one makes choices using popularity-based heuristics. There is a vast literature proving that popularity is indeed an important factor that influences choices: Cai et al., e.g., studied the choices from a restaurant menu and showed that the demand for the most popular dishes increased when this information was made explicitly available [12]. Salganick and collaborators conducted a randomized experiment based on an artificial market for music downloads and proved a marked difference in people choices when they are exposed to those of others [44]; such differences can ultimately result in popularity rankings (for the same set objects) that differ greatly from experiment to experiment and from the ranking that one would observe if choices were made in isolation and therefore independently. Such idiosyncrasies can be attributed to the noisy popularity bias induced by the somewhat random choices made by the first users to perform their choices, and they are the reason that make success hard to predict in cultural markets [52]. We can easily appreciate the powerful effect of noisy initial fluctuations in popularity ranking in modelling cultural market, e.g., in [54] items

equally appreciated by generic users can reach very different levels of popularity due to popularity biases. Other studies have observed empirically the somewhat pernicious effect of popularity bias [29], or more in general of social influence, in our choices [37]. A related, but distinct, issue is that when popularity informs the recommendation of a recommending systems it can over concentrate the attention of users on the most popular items at the expense of equally valid but poorly recognized others (see e.g., [1]). Indeed, even when the recommender system is well aligned with the preference of the majority, it can still happen that it produces “unfairly” poor recommendation for those whose preferences are misaligned with those of the majority of users [3]. Finally, it is probably interesting to notice that also in situations where a reasonably uncontroversial definition of quality can be assumed (where, e.g., it can be linked to measurable outcomes or performances) popularity bias can provoke misalignments between quality and popularity rankings [8]. In [18] the authors have shown that consumers are more sensible to revealed rather than stated preferences regarding apps downloads from Google Play. A similar result is reported for hotels choices in [51].

Many different, often overlapping, reasons lay behind our tendency to inform our choices by popularity. First, people’s behavior carries information and there are circumstances in which it is a perfectly rational and effective strategy to observe and copy, even though this may lead to information cascades, as demonstrated in the classical works [7, 10, 53]. Second, adopting popular choices can derive from social influence [6, 28, 39], which, in turn, can produce undesirable outcomes such as herding effects, that have been well documented in many areas of human behavior [41]. Third, the tendency of choosing the most popular items may reflect a cognitive “bias” or heuristics. In cognitive sciences, the notion of “recognition” heuristics has been widely studied. In the words of Goldstone and Gingerenzer [21]: “Consider the task of inferring which of two objects has a higher value on some criterion (e.g., which is faster, higher, stronger). The recognition heuristic for such tasks is simply stated: If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value”.

While it can be cost and time effective to use popularity as a proxy for quality and value, this does not come without a price. A side effect of popularity driven dynamics is that it implies a positive feedback loop by which items that are already popular tend to become even more popular. The net effect of such dynamics is that, even when it produces popularity rankings that are aligned with the quality of items, it concentrates the collective attention towards few items at the top. Indeed, since the seminal work of Simon [47], the rich-get-richer phenomenon has been quantitatively translated in the principle of linear growth and claimed to explain power law distribution of quantities that can be regarded as popularity [43]. Most recently, the principle has found new life in network science, where it is known as preferential attachment [9] and has been used to justify the widespread occurrence of power law degree distributions in network abstraction of real-world systems.

When this effect is brought to the extreme and all popularity is concentrated in only few items at the top of the ranking, then, naturally, diversity in the system disappears. This danger has been noted in several contexts, including that of recommender systems. In [24] authors deal with the problem of data scarcity affecting items in the tail of the popularity distribution, and develop techniques to improve their estimated quality and robustness to shilling attacks. In [2] authors propose a regularization-based framework to enhance the long-tail coverage of recommendation lists in a learning-to-rank algorithm, in order to achieve a desired trade-off between accuracy and coverage. [4] further explores the fairness problem in recommendation by considering how different user groups are affected by algorithmic popularity bias. The same issues of popularity concentration towards few top-ranking items has been considered in studies about search engines. These presumably use popularity as a signal of relevance. Also, they direct attention towards already

popular sites via the (popularity-based) ranking they produce [14]. Also, in this case there is an implicit risk of a positive feedback loop, although this tendency is dampened by the diversity of user interests [20].

In the recommendation systems community, it is well known that popularity bias and feedback loop in a long run operation have important consequences on the performance of the system itself. This has been studied - among the others - by Chaney et al. [13], Jiang et al. [27], D'Amour et al. [16], and Mansoury et al. [31].

On the empirical side, a systematic study of the interplay between quality and popularity is problematic for several reasons. The most prominent is that the notion of quality, especially where cultural markets are concerned, is to some extent subjective and therefore difficult to operationalize. A number of studies have nevertheless confirmed the finding of the music lab experiment [37, 45, 46, 50]. Motivated by [45], a number of theoretical papers have tried to avoid the empirical difficulties by postulating some mechanism of choice that takes popularity into account. Krumme et al. proposed an agent based model that closely reproduces the findings of the music lab experiment [28]. Van Hentenryck and collaborators considered a model of market in which customers can try products before committing to buying. They introduce a policy which successfully recovers quality ranking but asymptotically leads to a monopoly of the top-quality item [22].

Popularity bias is not easily eliminated, even when one is aware of it. In [32] authors investigate the **non-random missing data problem (NRMD)**, e.g., the fact that users are more likely to supply ratings for items that they do like, and less likely to supply ratings for items that they do not like. Incorrect assumptions about missing data have been found to lead to biased parameter estimation in collaborative ranking. In [33] authors focus on the impact of false-positives, i.e., suggestions that are disliked by users, discovering a surprising degree of disagreement with true positives, actually penalizing the most popular items. In [11] authors formally investigate how popularity bias is affected by random variables such as item relevance, item discovery by users, and the decision by users to interact with discovered items.

8 CONCLUSIONS AND FUTURE WORK

In this paper we have shown that a small effort on behalf of users to discriminate quality out of popularity can, in the long run, straighten out systems that, by themselves, could drift to undesirable configurations containing severe misalignments, especially among the top-quality items.

Readers familiar with load balancing strategies might recognize an analogy between our findings and the so called “power of two choices” paradigm [35], which has been applied also to other computer science problems, such as hashing and shared memory machines [36]. For example, in the classical balls-and-bins model, suppose that n balls are thrown into n bins. A logarithmic reduction in the maximum number of balls in any bin is obtained when, instead of throwing each ball uniformly at random, we select two bins uniformly at random, and put the ball in the least loaded one [5]. Note that in our case the goal is not to obtain a balanced allocation of balls among the bins, but an allocation satisfying a desired ranking (the one associated to intrinsic quality). Although our problem is different, the discovered phenomenon is similar: a minimum, local effort performed during the addition of an individual ball can, in the long run, rectify the distortions produced by a purely random choice, guaranteeing to achieve the desired configuration. In particular, in analogy to the “power of two choices”, we discovered that $K_{\min} = 2$ is enough to guarantee convergence to the desired configuration in any system run (Section 3), assuming that items are selected uniformly at random.

In contrast to classic results for the balls-and-bins model, we have considered also the case in which the candidate bins to receive a new ball are not selected uniformly at random, but according to their current load (ranking model of exponent α). Quite surprisingly, we have found that in this

case $K = 2$ is not always enough to guarantee convergence to the desired ranking. In particular, for $0 < \alpha < 1$, K_{\min} is even unbounded as N grows large, while for $\alpha > 1$ we recover the effect that a bounded, small $K \ll N$ is enough (Section 4), though $K_{\min} = 2$ only for large α . Unfortunately, the regime $\alpha > 1$ might not allow us to achieve the desired level of fairness among items (Section 5.1).

Our analysis is not without limitations, and several directions could be pursued to generalize it and make it closer to real systems. For example, we have assumed that users can always find the best-quality items among a subset of K : one could incorporate also a possibly imperfect outcome of the user evaluation process. Moreover, the K inspected items might not always be on the top of the list produced by the system (see example in Section 2.1), and one could try to combine the effect of users focusing their attention on random items of the list. Another important point is that modern platforms attempt to automatically estimate item quality in a way similar to what real users would do, e.g., by analyzing reviews left by other users (e.g., by sentiment analysis), and combine such estimates with popularity metrics in building their recommendation lists, de facto boosting the quality discrimination performed by users alone. Indeed, it would be interesting to develop models in which popularity, quality and randomness are jointly combined to determine the subset of items inspected by users.

At last, we acknowledge that real systems do not construct recommendation lists based on our idealized ranking model (power-law with exponent α): different laws (other than power-law) could be considered, as well as computationally simpler strategies to present to users randomized lists so as to favor diversity and allow serendipity. For example, some platforms (e.g., Amazon) seem to just perturb the list produced by their ranking algorithms by randomly moving up or down a few items (occasionally adding also some ‘intruders’), so that each non-repeated query produces a slightly different items view. It would be interesting to study whether such simpler ways to randomize lists could be mapped onto our ranking model with properly chosen α . Another interesting direction would be to consider dynamic systems [25] in which new items are progressively inserted into the catalog, and study strategies to let them quickly emerge over older (worse) ones.

Finally, our preliminary investigation in Section 6 suggests that the system behavior in the case of multiple user classes is much richer and more complex than the single class case. Future work should expand our understanding of the multi-class scenario, and, more in general, the impact of heterogeneous perception of quality.

APPENDIX

A PROOF OF PROPOSITION 3.1

PROOF. Let $r_{\min} \triangleq \min_i \frac{b_{i+1}}{b_i} > 1$ be the minimum ratio between two consecutive (positive) winning probabilities. For example, in the case of winning probabilities (1) we have, for $K \leq i < N$,

$$\frac{b_{i+1}}{b_i} = \frac{i}{i - K + 1}$$

hence $r_{\min} = \frac{N-1}{N-K} > 1$. Moreover, let $b_{\min} \triangleq \min_i \{b_i\}$.

Next, we recall the following classic bounds for the tails of the binomial distribution $\text{Bin}(n, p)$ (see e.g., Lemma 1.1 p. 16 in [38]). Let $\mu := np$, and

$$H(a) := 1 - a + a \log a, \quad a \in \mathbb{R}_+. \quad (15)$$

For any $0 < k < n$, we have:
if $k \geq \mu$, then

$$P(\text{Bin}(n, p) \geq k) \leq \exp \left(-\mu H \left(\frac{k}{\mu} \right) \right); \quad (16)$$

if $k \leq \mu$, then

$$P(\text{Bin}(n, p) \leq k) \leq \exp\left(-\mu H\left(\frac{k}{\mu}\right)\right). \quad (17)$$

We identify a suitable threshold t_i in between each pair of consecutive winning probabilities (b_i, b_{i+1}) . In particular, we set t_i as the geometric mean of b_i and b_{i+1} :

$$t_i = \sqrt{b_i b_{i+1}}, \quad K \leq i < N$$

and we also set $t_{K-1} = 0$, $t_N = 1$. We denote by $C_n^i = \{t_{i-1} < \frac{w_i[n]}{n} < t_i\}$ the event that, at round n , the empirical frequency of the number of times object i wins is comprised in the interval (t_{i-1}, t_i) . Observe that we can safely disregard initial weights $\mathbf{w}[0]$, which provide vanishing contribution to empirical frequencies $\frac{w_i[n]}{n}$.

Note that the joint occurrence of all events $C_n^i, \forall i$, is a sufficient but not necessary condition for the occurrence of E_n , which is the event that weights are correctly ordered. Therefore, denoting by \bar{E} the complement of a generic event E , we have:

$$E_n \supseteq \left(\cap_{i=K}^N C_n^i\right) \Rightarrow \bar{E}_n \subseteq \left(\cup_{i=K}^N \bar{C}_n^i\right)$$

By the union bound,

$$\mathbb{P}(\bar{E}_n) \leq \sum_{i=K}^N \mathbb{P}(\bar{C}_n^i).$$

Using the above deviation bounds of the binomial distribution, we have, for $K < i < N$:

$$\begin{aligned} \mathbb{P}(\bar{C}_n^i) &= \mathbb{P}(\text{Bin}(n, b_i) \geq nt_i) + \mathbb{P}(\text{Bin}(n, b_i) \leq nt_{i-1}) \leq \exp(-nb_i H(t_i/b_i)) + \exp(-nb_i H(t_{i-1}/b_i)) \\ &= \exp(-nb_i H(\sqrt{b_{i+1}/b_i})) + \exp(-nb_i H(\sqrt{b_{i-1}/b_i})) \leq \exp(-nb_i H(\sqrt{r_{\min}})) + \exp(-nb_i H(\sqrt{1/r_{\min}})). \end{aligned} \quad (18)$$

For the special case $i = K$ ($i = N$), we have to consider only the right (left) tail, obtaining:

$$\begin{aligned} \mathbb{P}(\bar{C}_n^K) &= \mathbb{P}(\text{Bin}(n, b_K) \geq nt_K) \leq \exp(-nb_K H(\sqrt{r_{\min}})) \\ \mathbb{P}(\bar{C}_n^N) &= \mathbb{P}(\text{Bin}(n, b_K) \leq nt_{N-1}) \leq \exp(-nb_N H(\sqrt{1/r_{\min}})). \end{aligned}$$

In conclusion, since $b_i \geq b_{\min}$, we have

$$\mathbb{P}(\bar{E}_n) \leq (N-1) \left[\exp(-nb_{\min} H(\sqrt{r_{\min}})) + \exp(-nb_{\min} H(\sqrt{1/r_{\min}})) \right]$$

meaning that the probability of event \bar{E}_n decreases exponentially to zero as n increases. Since $\sum_n \mathbb{P}(\bar{E}_n) < \infty$, the assert follows from the Borel-Cantelli lemma. \square

B PROOF OF PROPOSITION 4.1

SKETCH OF PROOF. We present a simplified proof based on the analysis of the trajectory of mean (normalized) weights $\tilde{\mathbf{w}} = \mathbb{E}[\tilde{\mathbf{w}}]$. As it usually occurs in Pólya urn models, the actual system exhibits wild random fluctuations at the beginning of the process, but the impact of stochasticity diminishes over time due to the accumulation of balls in the urns, which makes the randomness related to the addition of individual balls less and less important, so that the system behavior becomes more and more concentrated around mean values, which follow a deterministic trajectory on which we restrict our attention here.

Possible values of normalized weights $\tilde{\mathbf{w}}$ lie in the simplex $\mathcal{S} = \{\mathbf{w} \in \mathbb{R}^N : w_1 + \dots + w_N = 1, w_i \geq 0, i = 1, \dots, N\}$, which can be partitioned into $M = N!$ subregions $\{\mathcal{S}_i\}_1^M$ corresponding




Fig. 11. Example trajectories (left) and permutation graph (right) for the system $N = 3, K = 2, \alpha = 1$, *with-item-repetition*. The attraction basin of both stable points e_1 and e_2 is half of \mathcal{S} .

to the $N!$ possible permutations of indexes $\{1, \dots, N\}$ associated to sorted values of $\{\tilde{w}_i\}_i$. It is immediate to verify that subregions \mathcal{S}_i are convex. At the beginning of round n (i.e., after $n - 1$ balls have been added, starting from n_0 initial balls in the system), normalized weights $\tilde{w}[n - 1]$ lie in a subregion of \mathcal{S} associated to the vector of winning probabilities $b[n] = b(\tilde{w}[n - 1])$. Let $\Delta[n]$ be the random vector denoting which ball is added at round n , i.e., a vector of all zeros except for the index of the urn in which the ball is added, whose entry equals 1. We have

$$\tilde{w}[n] = \mathbb{E}[\tilde{w}[n]] = \mathbb{E} \left[\frac{(n_0 + n - 1)\tilde{w}[n - 1] + \Delta[n]}{n_0 + n} \right] = \frac{(n_0 + n - 1)\tilde{w}[n - 1] + b[n]}{n_0 + n}.$$

Therefore, from $n - 1$ to n , $\tilde{w}[n]$ moves within domain \mathcal{S} along the segment connecting $\tilde{w}[n - 1]$ to $b[n]$. Note that $b[n]$ can lie within a different subregion of \mathcal{S} with respect to vector $\tilde{w}[n - 1]$. In that case, $\tilde{w}[n]$ will eventually enter the subregion of $b[n]$, and since then start moving towards a possible different attraction point b . Instead, if $b[n]$ lies in the same subregion of \mathcal{S} as $\tilde{w}[n - 1]$, the trajectory will end up at point $b[n]$, which must be a fixed point of function $B(\cdot)$. Figure 11 illustrates examples of trajectories on the plane b_3, b_2 for the system $N = 3, K = 2, \alpha = 1$, *with-item-repetition*. In this case, simplex \mathcal{S} is partitioned into 6 subregions, and there are two stable points $e_1 = (\frac{4}{121}, \frac{21}{121}, \frac{96}{121})$ and $e_2 = (\frac{4}{121}, \frac{60}{121}, \frac{57}{121})$. The associated permutation graph is shown on the right of Figure 11.

Special fixed points of $B(\cdot)$ lying exactly on the hyperspaces in which two or more components are equal are not stable, due to the assumption that $B(\cdot)$ is an injective function of the rank of weights \tilde{w} . Indeed, small perturbations around such points, which naturally occurs in the actual urn model, would drive the system towards different attraction points.

To show that fixed points of $B(\cdot)$ are the only possible stable points of the system, it remains to prove that the system cannot have a limit cycle, or, in other words, the associated permutation graph is a **DAG (directed acyclic graph)**. This fact can be proven by contradiction: suppose that the system, as $n \rightarrow \infty$, visits in sequence subregions $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_C$, $C \geq 2$, characterized by winning probabilities b_i , $1 \leq i \leq C$, and staying an asymptotic fraction f_i of time in subregion i , $1 \leq i \leq C$. Then normalized weights will converge, as $n \rightarrow \infty$, to the unique point $b_\infty = \sum_{i=1}^C f_i b_i$, hence it cannot keep visiting different attraction points b_i . \square

C PROOF OF LEMMA 4.1

PROOF. We analyze the stability of a generic permutation in which the v -th highest quality item is the first one (in the natural sequence with increasing quality) that is not given the right popularity according to the natural permutation. In the natural permutation, such item (the v -th

highest quality) would get popularity $[v^\alpha G]^{-1}$, and suppose instead that it gets higher popularity $[(v - \Delta)^\alpha G]^{-1}$, with $\Delta \geq 1$, in the generic alternate permutation. For the alternate permutation to be stable, the winning probability b_v of this first disordered item must be higher than the winning probability of all items receiving lower popularity, and most crucially of the item that gets popularity value $[(v - \Delta + 1)^\alpha G]^{-1}$, whose quality cannot be smaller than that of the $(v - 1)$ -th highest quality item, since items with quality rank $v + 1, v + 2, \dots, N$ are given their natural popularity, by hypothesis. In this situation, the popularities assigned to items with quality rank $N, N - 1, \dots, v, v - 1$, are, respectively:

$$\frac{1}{N^\alpha G}, \frac{1}{(N - 1)^\alpha G}, \dots, \frac{1}{(v + 1)^\alpha G}, \frac{1}{(v - \Delta)^\alpha G}, \frac{1}{(v - \Delta + 1)^\alpha G}.$$

Let $s^*(v) = \sum_{i=v+1}^N (i^\alpha G)^{-1}$, and notice that this is the cumulative popularity of all items with quality rank higher than v , i.e., those items which receive by hypothesis their correct popularity according to the natural permutation. The winning probability of the first disordered item is then

$$b_v = \left[s^*(v) + \frac{1}{(v - \Delta)^\alpha G} \right]^K - [s^*(v)]^K$$

while the winning probability of the item with quality rank $v - 1$ is:

$$b_{v-1} = \left[s^*(v) + \frac{1}{(v - \Delta)^\alpha G} + \frac{1}{(v - \Delta + 1)^\alpha G} \right]^K - \left[s^*(v) + \frac{1}{(v - \Delta)^\alpha G} \right]^K.$$

The alternate permutation is stable provided that $b_v > b_{v-1}$, therefore we are left to equivalently analyze properties of function $F(\Delta, v, K)$:

$$F(\Delta, v, K) = 2 \left[s^*(v) + \frac{1}{(v - \Delta)^\alpha G} \right]^K - [s^*(v)]^K - \left[s^*(v) + \frac{1}{(v - \Delta)^\alpha G} + \frac{1}{(v - \Delta + 1)^\alpha G} \right]^K \quad (19)$$

where integers Δ, v, K can be chosen in the domain $\mathcal{D} : \Delta \in [1, N - 1], v \in [\Delta + 1, N], K \geq 1$, and we are especially interested in possible monotonicities of $F(\Delta, v, K)$ in the subdomain $\mathcal{D}^+ \subseteq \mathcal{D}$ in which $F(\Delta, v, K)$ is positive.

We observe that, in the special case $K = 1$:

$$F(\Delta, v, 1) = \frac{1}{G} \left[\frac{1}{(v - \Delta)^\alpha} - \frac{1}{(v - \Delta + 1)^\alpha} \right] > 0 \quad \forall (\Delta, v) \in \mathcal{D}$$

which also suggests that, for $K = 1$, any local swap of two items, starting from the natural permutation, leads to a stable alternate permutation. On the other hand, as $K \rightarrow \infty$, $F(\Delta, v, K)$ becomes eventually negative, $\forall (\Delta, v)$, meaning that, for sufficiently large K , any alternate permutation w.r.t. the natural one becomes unstable, and therefore there exists indeed K_{\min} such that for $K \geq K_{\min}$ the natural permutation is the only stable one. It remains to understand, for fixed K , which pair (Δ, v) provides the largest (positive) $F(\Delta, v, K)$ in \mathcal{D}^+ . It turns out that, for fixed Δ , $F(\Delta, v, K)$ is a decreasing function of v in subdomain \mathcal{D}^+ . This means that, for given displacement Δ of the first disordered item from its natural position, $F(\Delta, v, K)$ achieves its maximum for the minimum $v = \Delta + 1$. Moreover, $F(\Delta, \Delta + 1, K)$ is a decreasing function of Δ , and therefore it achieves its maximum for the minimum $\Delta = 1$. In conclusion, $F(\Delta, v, K)$ achieves its maximum for $\Delta = 1, v = \Delta + 1$, hence K_{\min} is determined by the *critical* permutation in which the top two highest quality objects are swapped. \square

D PROOF OF PROPOSITION 4.2

PROOF. For $N = 1$, we trivially have $K_{\min} = 1$. For $N \geq 2$, properties of K_{\min} can be proven by applying basic calculus to the real function of two real variables $F(x, G) : [1, \infty] \times [1 + 2^{-\alpha}, \infty] \rightarrow \mathbb{R}$:

$$F(x, G) = 2(1 - G^{-1}2^{-\alpha})^x - 1 - (1 - G^{-1} - G^{-1}2^{-\alpha})^x$$

obtained by rewriting inequality (7) relaxing integer K to the real value $x \geq 1$. We note that $F(x, G)$ is continuous over its domain in both x and G . Moreover, $F(1, G)$ is positive for any $\alpha > 0$, while $F(x, G) \rightarrow -1$ as $x \rightarrow \infty$. Computing $\frac{\partial F}{\partial x}$, one can easily see that there is at most one value of x at which $\frac{\partial F}{\partial x} = 0$, which is enough to conclude that, for $\alpha > 0$, there exists a unique value $x^* > 1$ at which $F(x^*, G) = 0$, $\frac{\partial F}{\partial x}|_{x=x^*} < 0$, before which $F(x, G)$ is positive, and after which $F(x, G)$ is negative. Since K must be an integer, we take the smallest integer larger than or equal to x^* , i.e., $K_{\min} = \lceil x^* \rceil$. In the special case $\alpha = 0$, $F(1, G) = 0$ and $F(x, G)$ is monotonically decreasing in x , and we can use $K_{\min} = 2$ to guarantee that $b_{N-1} < b_N$, for any N .

With respect to G , we observe that $F(x, G) \rightarrow 0^+$ as $G \rightarrow \infty$, and by computing $\frac{\partial F}{\partial G}$, we find that there is only one value of G at which $\frac{\partial F}{\partial G} = 0$. This is enough to conclude that $\frac{\partial F}{\partial G}|_{x=x^*} > 0$. The sign of the derivative of the implicit function $x = g(G)$ (implicitly defined by $F(x, G) = 0$) with respect to G can thus be determined by the implicit function theorem, which provides:

$$\frac{dg}{dG} = -\frac{\partial F / \partial G}{\partial F / \partial x}$$

at the points at which $F(x, G) = 0$. Since $\frac{\partial F}{\partial G} > 0$, while $\frac{\partial F}{\partial x} < 0$, we have that $\frac{dg}{dG}$ is positive, and since G monotonically increases with N we conclude that x^* is an increasing function of N , and thus $K_{\min} = \lceil x^* \rceil$ is a non-decreasing function of N .

Since, when $0 < \alpha \leq 1$, the Dirichlet series $G = \sum_{i=1}^{\infty} i^{-\alpha}$ diverges, it follows that, for $0 < \alpha \leq 1$, $K_{\min} \rightarrow \infty$ as $N \rightarrow \infty$. Instead, when $\alpha > 1$, as $N \rightarrow \infty$ G converges to Riemann zeta function $\zeta(\alpha)$. Therefore, $K_{\min}^{\infty} := \lceil g(\zeta(\alpha)) \rceil$ is an upper bound to K_{\min} for any N , when $\alpha > 1$. \square

E PROOF OF PROPOSITION 5.1

PROOF. We will need the following auxiliary lemma which can be proven by elementary calculus.

LEMMA E.1. *Given any two positive integers N and K , with $N > K$, the function $f(\alpha) : \mathbb{R}^+ \cup 0 \rightarrow (0, 1)$:*

$$f(\alpha) = \frac{\sum_{i=1}^K i^{-\alpha}}{\sum_{j=1}^N j^{-\alpha}}$$

is monotonically increasing with α .

PROOF.

$$\frac{df(\alpha)}{d\alpha} = \frac{1}{G^2} \left(\sum_{j=1}^N j^{-\alpha} \sum_{i=1}^K (-\log(i)) i^{-\alpha} - \sum_{j=1}^N (-\log(j)) j^{-\alpha} \sum_{i=1}^K i^{-\alpha} \right).$$

The term in parenthesis can be simplified to:

$$\sum_{i=1}^K i^{-\alpha} \sum_{j=K+1}^N \log(j) j^{-\alpha} - \sum_{i=1}^K \log(i) i^{-\alpha} \sum_{j=K+1}^N j^{-\alpha}$$

which is positive since, for any i , we take indexes $j > i$, and thus:

$$\sum_{i=1}^K i^{-\alpha} \sum_{j=K+1}^N \log(j) j^{-\alpha} > \sum_{i=1}^K i^{-\alpha} \sum_{j=K+1}^N \log(i) j^{-\alpha}.$$

□

Going back to Proposition 5.1, from the definition of \bar{q} :

$$\bar{q} = \sum_{i=1}^N i b_i = \sum_{i=1}^N i (s_i^K - s_{i-1}^K) = N - \sum_{i=1}^{N-1} s_i^K \quad (20)$$

where $s_i = \sum_{j=1}^i \frac{(N-j+1)^{-\alpha}}{G}$. Computing the derivative of the above expression with respect to α , we have

$$\frac{\partial \bar{q}}{\partial \alpha} = -K \sum_{i=1}^{N-1} s_i^{K-1} \frac{\partial s_i}{\partial \alpha}$$

which is positive provided that $\frac{\partial s_i}{\partial \alpha}$ is negative, for any i . Since

$$s_i = 1 - \frac{\sum_{j=1}^{N-i} j^{-\alpha}}{\sum_{j=1}^N j^{-\alpha}}$$

we can apply Lemma E.1 with $K = N - i$ to conclude that s_i is indeed a decreasing function of α , for any $i = 1 \dots N - 1$.

In the special case $\alpha = 0$, we have $s_i = i/N$, and we obtain the minimum possible value

$$\bar{q}_{\min} = N - \frac{\sum_{i=1}^{N-1} i^K}{N^K}.$$

As $\alpha \rightarrow \infty$, all cumulants s_i , $i = 1 \dots N - 1$ vanish, and \bar{q} approaches (from below) the limit N .

For given N and α , the fact that \bar{q} increases with K stems directly from (20), since each term s_i of the summation, being smaller than 1, decreases as we make K larger. □

REFERENCES

- [1] Himan Abdollahpour. 2019. Popularity bias in ranking and recommendation. In *Proc. of AAAI/ACM Conference on AI, Ethics and Society: Honolulu, USA*. <https://doi.org/10.1145/3306618.3314309>
- [2] Himan Abdollahpour, Robin Burke, and Bamshad Mobasher. 2017. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys'17)*. Association for Computing Machinery, New York, NY, USA, 42–46. <https://doi.org/10.1145/3109859.3109912>
- [3] Himan Abdollahpour, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The unfairness of popularity bias in recommendation. In *Proc. of RMSE Workshop at ACM Recsys*.
- [4] Himan Abdollahpour, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2020. The connection between popularity bias, calibration, and fairness in recommendation. In *Fourteenth ACM Conference on Recommender Systems (RecSys'20)*. Association for Computing Machinery, New York, NY, USA, 726–731. <https://doi.org/10.1145/3383313.3418487>
- [5] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. 1999. Balanced allocations. *SIAM J. Comput.* 29, 1 (1999), 180–200. <https://doi.org/10.1137/S0097539795288490>
- [6] Eytan Bakshy, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. Everyone's an influencer: Quantifying influence on Twitter. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining (WSDM'11)*. Association for Computing Machinery, New York, NY, USA, 65–74. <https://doi.org/10.1145/1935826.1935845>
- [7] Abhijit V. Banerjee. 1992. A simple model of herd behavior*. *The Quarterly Journal of Economics* 107, 3 (Aug. 1992), 797–817. <https://doi.org/10.2307/2118364>
- [8] Albert-László Barabási. 2018. *The Formula: The Universal Laws of Success*. Little, Brown and Company.
- [9] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (Oct. 1999), 509–512. <https://doi.org/10.1126/science.286.5439.509>

- [10] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100, 5 (1992), 992–1026. <https://www.jstor.org/stable/2138632>
- [11] Rocio Cañamares and Pablo Castells. 2018. Should I follow the crowd? A probabilistic analysis of the effectiveness of popularity in recommender systems. In *The 41st International ACM SIGIR Conference on Research And Development in Information Retrieval (SIGIR'18)*. Association for Computing Machinery, New York, NY, USA, 415–424. <https://doi.org/10.1145/3209978.3210014>
- [12] Hongbin Cai, Yuyu Chen, and Hanming Fang. 2009. Observational learning: Evidence from a randomized natural field experiment. *American Economic Review* 99, 3 (June 2009), 864–882. <https://doi.org/10.1257/aer.99.3.864>
- [13] Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. 2018. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys'18)*. Association for Computing Machinery, New York, NY, USA, 224–232. <https://doi.org/10.1145/3240323.3240370>
- [14] Junghoo Cho and Sourashis Roy. 2004. Impact of search engines on page popularity. In *Proceedings of the 13th International Conference on World Wide Web (WWW'04)*. Association for Computing Machinery, New York, NY, USA, 20–29. <https://doi.org/10.1145/988672.988676>
- [15] Giovanni Luca Ciampaglia, Azadeh Nematzadeh, Filippo Menczer, and Alessandro Flammini. 2018. How algorithmic popularity bias hinders or promotes quality. *Scientific Reports* 8, 1 (Oct. 2018), 15951. <https://doi.org/10.1038/s41598-018-34203-2>
- [16] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D. Sculley, and Yoni Halpern. 2020. Fairness is not static: Deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 525–534. <https://doi.org/10.1145/3351095.3372878>
- [17] Mukund Deshpande and George Karypis. 2004. Item-based top-n recommendation algorithms. *ACM Trans. Inf. Syst.* 22, 1 (Jan. 2004), 143–177. <https://doi.org/10.1145/963770.963776>
- [18] Per Engström and Eskil Forsell. 2018. Demand effects of consumers' stated and revealed preferences. *Journal of Economic Behavior & Organization* 150, C (2018), 43–61. <https://ideas.repec.org/a/eee/jeborg/v150y2018icp43-61.html>
- [19] S. Fortunato, A. Flammini, and F. Menczer. 2006. Scale-free network growth by ranking. *Physical Review Letters* 96, 21 (2006), 218701.
- [20] S. Fortunato, A. Flammini, F. Menczer, and A. Vespignani. 2006. Topical interests and the mitigation of search engine bias. *Proceedings of the National Academy of Sciences* 103, 34 (Aug. 2006), 12684–12689. <https://doi.org/10.1073/pnas.0605525103>
- [21] Gerd Gigerenzer, Peter M. Todd, and ABC Research Group. 2000. *Simple Heuristics That Make Us Smart* (illustrated edition ed.). Oxford University Press, Oxford.
- [22] Pascal Van Hentenryck, Andrés Abeliuk, Franco Berbeglia, Felipe Maldonado, and Gerardo Berbeglia. 2016. Aligning popularity and quality in online cultural markets. *Proceedings of the International AAAI Conference on Web and Social Media* 10, 1 (March 2016). <https://ojs.aaai.org/index.php/ICWSM/article/view/14720>
- [23] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. 2004. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.* 22, 1 (Jan. 2004), 5–53. <https://doi.org/10.1145/963770.963772>
- [24] Liang Hu, Longbing Cao, Jian Cao, Zhiping Gu, Guandong Xu, and Jie Wang. 2017. Improving the quality of recommendations for users and items in the tail of distribution. *ACM Trans. Inf. Syst.* 35, 3, Article 25 (June 2017), 37 pages. <https://doi.org/10.1145/3052769>
- [25] Gerardo Iñiguez, Carlos Pineda, Carlos Gershenson, and Albert-László Barabási. 2021. Universal dynamics of ranking. *arXiv preprint arXiv:2104.13439* (2021).
- [26] Dietmar Jannach, Lukas Lerche, Iman Kamehkhosh, and Michael Jugovac. 2015. What recommenders recommend: An analysis of recommendation biases and possible countermeasures. *User Modeling and User-Adapted Interaction* 25, 5 (Dec. 2015), 427–491. <https://doi.org/10.1007/s11257-015-9165-3>
- [27] Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. 2019. Degenerate feedback loops in recommender systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AI/ES'19)*. Association for Computing Machinery, New York, NY, USA, 383–390. <https://doi.org/10.1145/3306618.3314288>
- [28] Coco Krumme, Manuel Cebrian, Galen Pickard, and Sandy Pentland. 2012. Quantifying social influence in an online cultural market. *PLOS ONE* 7, 5 (5 2012), 1–6. <https://doi.org/10.1371/journal.pone.0033785>
- [29] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108, 22 (2011), 9020–9025. <https://doi.org/10.1073/pnas.1008636108> arXiv:<https://www.pnas.org/content/108/22/9020.full.pdf>
- [30] H. M. Mahmoud. 2008. *Pólya Urn Models*. Chapman & Hall/CRC.
- [31] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback loop and bias amplification in recommender systems. In *Proceedings of the 29th ACM International Conference on*

- Information & Knowledge Management (CIKM'20)*. Association for Computing Machinery, New York, NY, USA, 2145–2148. <https://doi.org/10.1145/3340531.3412152>
- [32] Benjamin M. Marlin and Richard S. Zemel. 2009. Collaborative prediction and ranking with non-random missing data. In *Proceedings of the Third ACM Conference on Recommender Systems (RecSys'09)*. Association for Computing Machinery, New York, NY, USA, 5–12. <https://doi.org/10.1145/1639714.1639717>
- [33] Elisa Mena-Maldonado, Rocío Cañamares, Pablo Castells, Yongli Ren, and Mark Sanderson. 2020. Agreement and disagreement between true and false-positive metrics in recommender systems evaluation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'20)*. Association for Computing Machinery, New York, NY, USA, 841–850. <https://doi.org/10.1145/3397271.3401096>
- [34] Robert K. Merton. 1968. The Matthew Effect in science: The reward and communication systems of science are considered. *Science* 159, 3810 (Jan. 1968), 56–63. <https://doi.org/10.1126/science.159.3810.56>
- [35] M. Mitzenmacher. 2001. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems* 12, 10 (2001), 1094–1104. <https://doi.org/10.1109/71.963420>
- [36] Michael Mitzenmacher, Andréa W. Richa, and Ramesh Sitaraman. 2000. The power of two random choices: A survey of techniques and results. In *Handbook of Randomized Computing*. Kluwer, 255–312.
- [37] Lev Muchnik, Sinan Aral, and Sean J. Taylor. 2013. Social influence bias: A randomized experiment. *Science (New York, N.Y.)* 341, 6146 (Aug. 2013), 647–651. <https://doi.org/10.1126/science.1240466>
- [38] M. Penrose. 2003. *Random Geometric Graphs*. Oxford University Press. <https://books.google.it/books?id=M38e7nPGSCsC>
- [39] Nicholas A. Christakis MD PhD and James H. Fowler PhD. 2009. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives* (first printing edition ed.). Little, Brown Spark, New York.
- [40] Xiaoyan Qiu, Diego Oliveira, Alireza Shirazi, Alessandro Flammini, and Filippo Menczer. 2017. Limited individual attention and online virality of low-quality information. *Nature Human Behaviour* 1 (2017), 0132. <https://doi.org/10.1038/s41562-017-0132>
- [41] Ramsey M. Raafat, Nick Chater, and Chris Frith. 2009. Herding in humans. *Trends in Cognitive Sciences* 13, 10 (2009), 420–428. <https://doi.org/10.1016/j.tics.2009.08.002>
- [42] Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor. 2010. *Recommender Systems Handbook*. (2010).
- [43] Daniel Rigney. 2010. *The Matthew Effect: How Advantage Begets Further Advantage*. Columbia University Press, New York. OCLC: ocn320798941.
- [44] Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762 (Feb. 2006), 854–856. <https://doi.org/10.1126/science.1121066>
- [45] Matthew J. Salganik and Duncan J. Watts. 2008. Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly* 71, 4 (Dec. 2008), 338–355. <https://doi.org/10.1177/019027250807100404>
- [46] Matthew J. Salganik and Duncan J. Watts. 2009. Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in Cognitive Science* 1, 3 (July 2009), 439–468. <https://doi.org/10.1111/j.1756-8765.2009.01030.x>
- [47] Herbert A. Simon. 1955. On a class of skew distribution functions. *Biometrika* 42, 3/4 (1955), 425–440. <https://doi.org/10.2307/2333389>
- [48] G. Stoddard. 2015. Popularity dynamics and intrinsic quality in Reddit and Hacker News. In *ICWSM*.
- [49] James Surowiecki. 2005. *The Wisdom of the Crowds*. Anchor Books.
- [50] Arnout van de Rijt, Soong Moon Kang, Michael Restivo, and Akshay Patil. 2014. Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences of the United States of America* 111, 19 (May 2014), 6934–6939. <https://doi.org/10.1073/pnas.1316836111>
- [51] Giampaolo Viglia, Roberto Furlan, and Antonio Ladrón-de Guevara. 2014. Please, talk about it! When hotel popularity boosts preferences. *International Journal of Hospitality Management* 42 (Sept. 2014), 155–164. <https://doi.org/10.1016/j.ijhm.2014.07.001>
- [52] Duncan J. Watts. 2011. *Everything Is Obvious: *Once You Know the Answer*. Currency.
- [53] Ivo Welch. 1992. Sequential sales, learning, and cascades. *The Journal of Finance* 47, 2 (1992), 695–732. <https://doi.org/10.1111/j.1540-6261.1992.tb04406.x>
- [54] Lilian Weng, Alessandro Flammini, Alessandro Vespignani, and Filippo Menczer. 2012. Competition among memes in a world with limited attention. *Scientific Reports* 2 (2012).

Received 23 May 2021; revised 22 March 2022; accepted 6 April 2022