

Asymptotic Optimality of Myopic Optimization in Trial-Offer Markets with Social Influence

Andrés Abeliuk

Data61 & U. of Melbourne
aabeliuk@student.unimelb.edu.au

Felipe Maldonado

Data61 & ANU
felipe.maldonado@anu.edu.au

Gerardo Berbeglia

Melbourne Business School
g.berbeglia@mbs.edu

Pascal Van Hentenryck

University of Michigan
pvanhent@umich.edu

Abstract

We study dynamic trial-offer markets, in which participants first try a product and later decide whether to purchase it or not. In these markets, social influence and position biases have a greater effect on the decisions taken in the sampling stage than those in the buying stage. We consider a myopic policy that maximizes the market efficiency for each incoming participant, taking into account the inherent quality of products, position biases, and social influence. We prove that this myopic policy is optimal and predictable asymptotically.

1 Introduction

Social influence is ubiquitous in online cultural markets. From book recommendations to song popularities, social influence has become a critical aspect of the customer experience. Social influence may appear through different signals such as the number of past purchases; consumer ratings; and/or consumer recommendations, depending on the market and/or platform. Social influence often comes reinforced by position bias (e.g., [Lerman and Hogg, 2014]), as consumer preferences are also affected, sometimes significantly, by the visibility of the choices. In digital markets, the impact of visibility on consumer behavior has been widely observed [Kempe and Mahdian, 2008; Aggarwal *et al.*, 2008].

Yet, despite its ubiquity, there is still considerable debate about the benefits of social influence and its effects on the market. Salganik *et al.* [2006] argued that social influence makes markets more unpredictable. They created an artificial music market, a trial-offer market called the MUSICLAB and showed that social influence may create significant unpredictability and inefficiencies. Follow-up experiments have been made confirming these initial findings [Salganik and Watts, 2009; van de Rijt *et al.*, 2014].

The MUSICLAB experiments, however, relied on an implicit but critical design choice: The songs were displayed to participants in decreasing order of popularity, reinforcing the social signal with position bias. Thus, creating a herding effect [Muchnik *et al.*, 2013; Hogg and Lerman, 2015], leading to a “rich get richer” phenomenon. Unfortunately, popular-

ity is not always a good proxy for quality and may lead to self-fulfilling prophecies [Salganik and Watts, 2008]. Different search engine rankings have been studied experimentally (see e.g., [Ghose *et al.*, 2012; 2014]).

Kleinberg (2008) articulated the need to develop an expressive computational model to understand the long-term effect of social influence. A first step in answering this open issue appeared in Krumme *et al.* [2012]: The authors proposed a discrete choice model for modeling the MUSICLAB. This model was then used in Abeliuk *et al.* [2015] to analyze a new policy for displaying the products, referred to as *performance ranking*. This ranking is a myopic policy that maximizes the efficiency of the market for each incoming participant, taking into account the inherent quality of products, position bias, and social influence. Their computational results show that the performance ranking significantly decreases the unpredictability of the market. A different policy, referred to as *quality ranking*, was proposed in Van Hentenryck *et al.* [2016]. The quality ranking only takes into account the inherent quality of products and yet, the authors showed that this policy is optimal and predictable asymptotically.

In this paper, we reconsider the performance ranking from Abeliuk *et al.* [2015] and study its asymptotic convergence. We show that this myopic policy is also optimal and predictable asymptotically, in addition to being optimal at each step. In contrast to the quality ranking, which is a static policy, the performance ranking is dynamic. Thus, from a technical standpoint, our analysis is the first to provide theoretical guarantees over a dynamic policy in cultural markets. Moreover, computational results show that the rate of convergence for the performance ranking is considerably faster than the quality ranking. Additionally, we correct and simplify a proof from Abeliuk *et al.* [2015] showing that the market always benefits from position bias and social influence in expectation under this policy. These results indicate that the performance ranking is an attractive policy for trial-offer markets: It avoids the pathological effects of the popularity ranking and it optimizes market efficiency both locally and asymptotically.

2 Trial-Offer Markets

This section introduces trial-offer markets in which participants can try a product before deciding to buy it. Such mod-

els are now pervasive in online cultural markets (e.g., books and songs). The market is composed of n products and each product $i \in \{1, \dots, n\}$ is characterized by two values:

1. Its *appeal* a_i which represents the inherent preference of trying product i ;
2. Its *quality* q_i which represents the conditional probability of purchasing product i given that it was tried.

Each market participant is presented with a product list π : she then tries a product s in π and decides whether to purchase s with a certain probability. The product list is a permutation of $\{1, \dots, n\}$ and each position p in the list is characterized by its *visibility* $v_p > 0$ which is the inherent probability of trying a product in position p . Since the list π is a bijection from positions to products, its inverse is well-defined and is called a ranking. Rankings are denoted by the letter σ , π_i denotes the product in position i of the list π , and σ_i identifies the position of product i in the ranking σ . Hence v_{σ_i} denotes the visibility of the position of product i .

Our primary objective is to maximize the market efficiency, i.e., the expected number of purchases. Note also that the higher this objective is, the lower the probability that consumers try a product but then decide not to purchase it. Hence, maximizing the expected efficiency of the market also minimizes unproductive trials.

Dynamic Market Since we are interested in the long-term effects of social influence, we consider a multi-period, dynamic market where consumers arrive sequentially, one per time period. Upon arrival, a consumer is able to observe the aggregate purchase decisions of her predecessors. Denote by $d^t = (d_1^t, \dots, d_n^t)$ the total number of consumers who purchased product i until the beginning of period t . The probability that the consumer arriving at period t will try product i if items are displayed using position assignment σ is given by

$$P_i(\sigma, d^t) = \frac{v_{\sigma_i}(a_i + d_i^t)}{\sum_{j=1}^n v_{\sigma_j}(a_j + d_j^t)}.$$

Observe that consumer choice preferences for trying the products are essentially modeled as a discrete choice model based on a multinomial logit [Luce, 1965] in which product utilities are affected by their position. The market uses the number of purchases d^t of product i at time t as the social signal. However, other social signals such as the market share, used in online site such as iTunes, are equivalent as we show next. Let $\phi^t = (\phi_1^t, \dots, \phi_n^t)$ denote the market shares at time t in terms of the total number of purchases d^t , i.e.,

$$\phi_i^t = \frac{d_i^t}{\sum_{j=1}^n d_j^t}.$$

The probability of trying product i can be rewritten as a function of ϕ^t , yielding,

$$P_i(\sigma, \phi^t) = \frac{v_{\sigma_i} \phi_i^t}{\sum_{j=1}^n v_{\sigma_j} \phi_j^t},$$

where, for simplicity, the vector d^0 is initialized with the products appeals, i.e., $d_i^0 = a_i$. Both notations are convenient

to stress different results; we use market shares when analyzing the asymptotic behavior of the market and the number of purchases for analyzing the static behavior of the market. Since q_i is the conditional probability of purchasing product i given that it was tried and hence, the expected number of purchases at time t is given by $\sum_{i=1}^n P_i(\sigma, d^t) q_i$.

Our goal is to study how the market shares evolve over time when social influence is present. Observe that the probability of trying a product depends on its position in the list, its appeal, and its number of purchases at time t . As the market evolves over time, the number of purchases dominates the appeal of the product and the trying probability of a product becomes its market share. Note also that in a dynamic market when no social signals are displayed, the purchase history plays no role and hence, the market behaves as a static market. Following Salganik and Watts [2008], we refer to this setup as the independent condition.

3 Rankings Policies

This section presents the ranking policies studied in this paper. Without loss of generality, we assume that the qualities and visibilities are non-increasing, i.e., $q_1 \geq q_2 \geq \dots \geq q_n$ and $v_1 \geq v_2 \geq \dots \geq v_n$. We also assume that the qualities and visibilities are known. In practical situations, the product qualities are obviously unknown but Abeliuk *et al.* (2015) have shown that they can be recovered accurately and quickly, either before or during the market execution.

The *performance ranking* was proposed by Abeliuk *et al.* (2015) to show the benefits of social influence in cultural markets. It maximizes the expected number of purchases at each iteration, exploiting all the available information globally, i.e., the appeal, the visibility, the purchases, and the quality of the products. More precisely, the performance ranking at step t produces a ranking σ_t^* defined as

$$\sigma_t^* = \arg\text{-max}_{\sigma \in S_n} \sum_{i=1}^n P_i(\sigma, d^t) \cdot q_i. \quad (1)$$

The performance ranking can be computed in strongly polynomial time and the resulting policy scales to large markets [Abeliuk *et al.*, 2015].

The *quality ranking* which simply orders the products by quality, assigning the product of highest quality to the most visible position and so on. With the above assumptions, the quality ranking σ satisfies $\sigma_i = i$ ($1 \leq i \leq n$). The quality ranking was shown to be asymptotically optimal and entirely predictable in Van Hentenryck *et al.* [2016].

These results contrast with the *popularity ranking* used in Salganik *et al.* [2006] to show the unpredictability caused by social influence in cultural markets. At iteration t , the popularity ranking orders the products by the number of purchases d_i^t but these purchases do not necessarily reflect the inherent quality of the products, since they depend on how many times the products were tried, which in turn depends on the position and social signal of the product.

4 Theoretical Analysis

In this section, we present a number of theoretical results on the performance ranking. In particular, we show that

the performance ranking always benefits from position bias and social influence and is an optimal and predictable policy asymptotically. For simplicity, the results assume that $q_1 > q_2 > \dots > q_n$.

Our first results characterize some fundamental properties of the performance ranking.

Lemma 1. *Let π^* be the optimal list for the static problem (Equation (1)) given market share ϕ and λ^* be the expected number of purchases given π^* and ϕ . Then*

$$\pi^* = \arg\text{-max}_{\pi} \sum_{i=1}^n v_i \phi_{\pi_i} (q_{\pi_i} - \lambda^*).$$

Proof. First observe that

$$\lambda^* = \frac{\sum_{i=1}^n v_i \phi_{\pi_i^*} q_{\pi_i^*}}{\sum_{i=1}^n v_i \phi_{\pi_i^*}} \Leftrightarrow 0 = \sum_{i=1}^n v_i \phi_{\pi_i^*} (q_{\pi_i^*} - \lambda^*). \quad (2)$$

Now assume that there exists π' such that

$$\sum_{i=1}^n v_i \phi_{\pi'_i} (q_{\pi'_i} - \lambda^*) > \sum_{i=1}^n v_i \phi_{\pi_i^*} (q_{\pi_i^*} - \lambda^*) = 0.$$

By reordering the terms, it comes that $\frac{\sum_{i=1}^n v_i \phi_{\pi'_i} q_{\pi'_i}}{\sum_{i=1}^n v_i \phi_{\pi'_i}} > \lambda^*$, which contradicts the optimality of π^* . \square

Lemma 1 through the rearrangement inequality provides an important characterization of the optimal ranking at time t .

Corollary 1. *Let λ_t^* be the expected number of purchases at time t under the performance ranking. The performance ranking π_t^* satisfies*

$$\phi_{\pi_{1,t}^*} (q_{\pi_{1,t}^*} - \lambda_t^*) \geq \dots \geq \phi_{\pi_{n,t}^*} (q_{\pi_{n,t}^*} - \lambda_t^*). \quad (3)$$

This corollary indicates that a product with quality greater or equal to λ_t^* is ranked higher than a product with quality smaller than λ_t^* . This property is independent of the market shares at time t .

The optimal expected number of purchases (Equation (1)) can be written as a function of the market shares:

$$\lambda(\phi) = \max_{\pi \in S_n} \frac{\sum_{i=1}^n v_i \phi_{\pi_i} q_{\pi_i}}{\sum_{i=1}^n v_i \phi_{\pi_i}}.$$

The continuity of $\lambda(\phi)$ is necessary to apply stochastic approximation methods, which are key to the derivation of the asymptotic behavior of the performance ranking.

Lemma 2. *$\lambda(\phi)$ is continuous for all $\phi \in \Delta^{n-1}$.*

Proof. When π is fixed, it is easy to verify that $\lambda(\phi)$ is continuous. The only source of discontinuity could come when the list π changes as a function of ϕ . However, Corollary 1 states that the list only changes when the inequality (3) flips sign for some i, j given market share ϕ' . This happens when

$$\phi'_i (q_i - \lambda(\phi')) = \phi'_j (q_j - \lambda(\phi')). \quad (4)$$

Since $q_i \neq q_j$, this implies that $\phi'_i \neq \phi'_j$. Solving Equation (4) for $\lambda(\phi')$ yields

$$\lambda(\phi') = \frac{\phi'_i q_i - \phi'_j q_j}{\phi'_i - \phi'_j}, \quad \phi'_i \neq \phi'_j.$$

Thus, $\lim_{\phi \rightarrow \phi'+} \lambda(\phi) = \lim_{\phi \rightarrow \phi'-} \lambda(\phi) = \lambda(\phi')$. \square

The Benefits of Social Influence An important question in cultural markets is whether the revelation of past purchases to consumers improves market efficiency. The theorem below states that, under the performance ranking policy, the expected marginal number of purchases in the studied trial-offer model increases when past purchases are revealed. The key to the proof is Lemma 3 which uses the following notations:

- $V_1(\sigma | d)$ is the expected number of purchases at period t when position assignment σ is displayed conditional to $d_t = d$, i.e.,

$$V_1(\sigma | d) = \sum_{i \geq 1} P_i(\sigma, d) q_i.$$

- $V_2(\sigma | d)$ is the expected number of purchases at period $t + 1$ conditional to $d_t = d$ when ranking σ is used at periods t and $t + 1$ consecutively, i.e.,

$$V_2(\sigma | d) = \sum_{i \geq 1} (P_i(\sigma, d) q_i V_1(\sigma | d + e_i)) + (1 - \sum_{i \geq 1} P_i(\sigma, d) q_i) V_1(\sigma | d),$$

where e_i denotes the i th unit vector.

The right term captures the case where no product is purchased at period t , while the left term captures the cases where a product i is purchased, which increases its social influence for the next period.

Lemma 3. *Let σ^* be the optimal position assignment of the static problem at period t given any $d_t = d$. We have*

$$V_2(\sigma^* | d) \geq V_1(\sigma^* | d).$$

Proof. We use $a_{i,t} = a_i + d_{i,t}$ to denote the perceived appeal of product i under social influence at time t . When the period t is not relevant, we omit it and use a_i instead for simplicity. Without loss of generality, we can rename the products so that $\sigma_i^* = i$. Let λ^* denote the optimal expected number of purchases at period t given $d_t = d$, i.e.,

$$V_1(\sigma^* | d) = \frac{\sum_i v_i a_i q_i}{\sum_i v_i a_i} = \lambda^*.$$

We need to prove that $V_2(\sigma^* | d) \geq V_1(\sigma^* | d) = \lambda^*$, which amounts to showing that

$$\sum_j \left(\frac{v_j a_j q_j}{\sum_i v_i a_i} \frac{\sum_i v_i a_i q_i + v_j q_j}{\sum_i v_i a_i + v_j} \right) + \left(1 - \frac{\sum_j v_j a_j q_j}{\sum_i v_i a_i} \right) \lambda^* \geq \lambda^*,$$

which reduces to proving

$$\sum_j \left[\frac{v_j^2 a_j q_j}{\sum_i v_i a_i + v_j} (q_j - \lambda^*) \right] \geq 0. \quad (5)$$

By Corollary 1, the performance ranking at period t satisfies

$$a_1 (q_1 - \lambda^*) \geq a_2 (q_2 - \lambda^*) \geq \dots \geq a_n (q_n - \lambda^*).$$

Hence, the performance ranking allocates all products with a negative term after the products with a positive term. Define $P_1 = \{i \in N \mid (q_i - \lambda^*) \geq 0\}$ and $P_2 = \{i \in N \mid (q_i - \lambda^*) < 0\}$. It follows that $v_i \geq v_j$ for all $i \in P_1, j \in P_2$. Inequality (5) can be rewritten as follows:

$$\sum_{j \in P_1} \frac{v_j a_j q_j v_j (q_j - \lambda^*)}{\sum_i v_i a_i + v_j} + \sum_{j \in P_2} \frac{v_j a_j q_j v_j (q_j - \lambda^*)}{\sum_i v_i a_i + v_j} \geq 0.$$

By definition of P_1 and P_2 , the terms in the left summation are positive and the terms in the right summation are negative. Additionally, for any $c > 0$,

$$\begin{aligned} \frac{v_i}{c + v_i} \geq \frac{v_j}{c + v_j} &\Leftrightarrow (c + v_j)v_i \geq (c + v_i)v_j \\ &\Leftrightarrow cv_i \geq cv_j \Leftrightarrow v_i \geq v_j \end{aligned}$$

Hence, since $v_1 \geq v_2 \geq \dots \geq v_n$,

$$\frac{v_1}{\sum_i v_i a_i + v_1} \geq \frac{v_2}{\sum_i v_i a_i + v_2} \geq \dots \geq \frac{v_n}{\sum_i v_i a_i + v_n}.$$

Let $\underline{v} = \min_{i \in P_1} v_i$ the smallest visibility assigned to an item in P_1 . Thus, $\underline{v} \geq \max_{i \in P_2} v_i$. Similarly, let $\underline{q} = \min_{i \in P_1} q_i$. By definition of P_1 and P_2 , $\underline{q} \geq \max_{i \in P_2} q_i$. These observations leads to the following inequality:

$$\begin{aligned} \sum_{j \in P_1} \frac{v_j a_j q_j v_j (q_j - \lambda^*)}{\sum_i v_i a_i + v_j} + \sum_{j \in P_2} \frac{v_j a_j q_j v_j (q_j - \lambda^*)}{\sum_i v_i a_i + v_j} &\geq \\ \frac{\underline{v}\underline{q}}{\sum_i v_i a_i + \underline{v}} \left(\sum_{j \in P_1} a_j v_j (q_j - \lambda^*) + \sum_{j \in P_2} a_j v_j (q_j - \lambda^*) \right) & \\ = \frac{\underline{v}\underline{q}}{\sum_i v_i a_i + \underline{v}} \sum_{j=1}^n a_j v_j (q_j - \lambda^*). & \end{aligned}$$

The result follows from Equation (2) in Lemma 1 which implies that

$$\sum_{j=1}^n (a_j v_j (q_j - \lambda^*)) = 0. \quad \square$$

Lemma 3 states that, if the optimal position assignment in period t is used at period $t+1$, the expected number of purchases at $t+1$ is at least as high as the expected number of purchases at period t . Clearly, re-optimizing at period $t+1$ by using the optimal position assignment for period $t+1$ can only increase the expected number of purchases. Together with Lemma 3, this observation leads to the following theorem.

Theorem 1. *The expected rate of purchases is non-decreasing over time for the performance ranking under social influence.*

Proof. For any t and $d_t = d$, let σ^* be the optimal position assignment in period t . By Lemma 3, we know that $V_2(\sigma^* \mid d) \geq V_1(\sigma^* \mid d)$. The performance ranking policy computes

$$\begin{aligned} \sum_{i \geq 1} (P_i(\sigma^*, d) \max_{\sigma \in S_n} V_1(\sigma \mid d + e_i)) \\ + (1 - \sum_{i \geq 1} P_i(\sigma^*, d)) V_1(\sigma^* \mid d) \geq V_2(\sigma^* \mid d) \end{aligned}$$

and the result follows. \square

This result contrasts with the popularity ranking, under which Theorem 1 does not hold [Abeliuk *et al.*, 2016].

As a direct consequence of Theorem 1, the sequence $\lambda_1^*, \dots, \lambda_t^*$ of expected numbers of purchases over time is a sub-martingale. By Doob's sub-martingale convergence theorem [Doob, 1953], the process $\{\lambda_t^*\}_{t>0}$ converges almost surely. The next section studies the convergence points.

Asymptotic Behavior of Performance Ranking We now prove the key result of the paper: The trial-offer market becomes a monopoly for the best product when the performance ranking is used at each step. As a consequence, the performance ranking is optimal asymptotically since the best product has the highest probability to be purchased. The result also indicates that trial-offer markets are predictable asymptotically when the performance ranking is used. The proof needs the following lemma that characterizes the probability that the next purchase is product i .

Lemma 4. *The probability p_i that the next purchase (after any number of steps) is product i given ranking σ is*

$$p_i = \frac{v_{\sigma_i} a_i q_i}{\sum_{j=1}^n v_{\sigma_j} a_j q_j}.$$

Proof. The probability that product i is purchased in step m while no product was purchased in earlier steps is

$$p_i^m = \left(1 - \frac{\sum_{j=1}^n v_{\sigma_j} a_j q_j}{\sum_{j=1}^n v_{\sigma_j} a_j} \right)^{m-1} \frac{v_{\sigma_i} a_i q_i}{\sum_{j=1}^n v_{\sigma_j} a_j}.$$

Hence the probability that the next purchased product is product i is given by $p_i = \sum_{m=0}^{\infty} p_i^m$. The result follows from $\sum_{m=0}^{\infty} (1-a)^m = \frac{1}{a}$. \square

Since, by Lemma 4, the steps in which no product is purchased can be ignored, we can use the following variables ($1 \leq i \leq n$) to specify the market:

$$X_{i,t} \doteq a_{i,t} \hat{q}_i(\sigma) \quad (6)$$

$$Z_{i,t} \doteq \frac{X_{i,t}}{\sum_{i=1}^n X_{i,k}} \quad (7)$$

where $\hat{q}_j(\sigma) = v_{\sigma_j} q_j$. The trial-offer market can thus be modeled as generalized Pólya scheme [Renlund, 2010], where $X_{i,t}$ represents the number of balls of type i at step t and $Z_{i,t}$ is the proportion of balls of type i at step t . Since $X_{i,t+1} = X_{i,t} + \hat{q}_i(\sigma)$ if product i is purchased, the generalized Pólya scheme add $\hat{q}_i(\sigma)$ balls each time product i is purchased. As a result, the Pólya scheme uses the stochastic replacement matrix

$$\mathbf{R}(\sigma) = \begin{pmatrix} \hat{q}_1(\sigma) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \hat{q}_n(\sigma) \end{pmatrix}. \quad (8)$$

Higuera et al (2003) showed that a generalized Pólya urn scheme with n colors can be modeled as a Robbins-Monro Algorithm as follows. The process $\{Z_i\}$ can be written as

$$Z_{i+1} = Z_i + \gamma_{i+1}(F(Z_i) + \epsilon_{i+1} + \beta_{i+1}), \quad (9)$$

where: (1) the process $\{\gamma_i\}$ is a decreasing sequence of positive random variables, $\sum_{i \geq 1} \gamma_i = \infty$ and $\sum_{i \geq 1} \gamma_i^2 < \infty$; (2) $\{\epsilon_i\}$ is a sequence of martingale differences with respect to $\{\mathcal{F}_i\}$, where \mathcal{F}_i is the natural filtration of the entire process; (3) $\{\beta_i\}$ is a negligible sequence such that $\beta_i \rightarrow 0$ almost surely; (4) $F(x) = xR(I - 1^T x)$ must be continuous, where R is the replacement matrix of the Pólya scheme. The ODE method [Ljung, 1977] relates the recurrence Equation (9) with the ordinary differential equation $\dot{x} = F(x)$. If this ODE has a globally asymptotically stable equilibrium point $u \in \mathbb{R}^n$, then the discrete process $\{Z_i\}$ converges almost surely to this point. The equilibrium points u are obtained by solving $F(\phi) = 0$ when $\sum_{i=1}^n \phi_i = 1$.

Lemma 5. *Let $F(\phi) = \phi R(\sigma)(I - 1^T \phi)$. The solutions to $F(\phi) = 0$ when $\sum_{i=1}^n \phi_i = 1$ are given by the set $\{e_i : 1 \leq i \leq n\}$, where e_i denotes the i th unit vector.*

Proof.

$$F(\phi) = \phi \left[R(\sigma) - \begin{pmatrix} \hat{q}_1(\sigma) \\ \vdots \\ \hat{q}_n(\sigma) \end{pmatrix} \phi \right] = [F_1(\phi), \dots, F_n(\phi)],$$

with

$$F_i(\phi) = \phi_i(v_{\sigma_i} q_i - \sum_{j=1}^n v_{\sigma_j} q_j \phi_j). \quad (10)$$

An equilibrium point must satisfy $F_i(\phi) = 0$ for all i . Thus, the points in $\{e_i : 1 \leq i \leq n\}$ are trivial solutions. We now show these are the only solutions. Assume that $S = \{i : \phi_i > 0\}$ is such that $|S| > 1$. Then, an equilibrium must satisfy

$$v_{\sigma_i} q_i = \sum_{j \in S} v_{\sigma_j} q_j \phi_j \quad \forall i \in S. \quad (11)$$

Define $l = \arg\text{-min}_{i \in S} q_i$ and $m = \arg\text{-max}_{i \in S} q_i$. Equation (11) when applied twice to l and m implies that $v_{\sigma_m} q_m = v_{\sigma_l} q_l$. Since $q_m > q_l$, $v_m < v_l$ must hold to satisfy the equilibrium condition (11). This means that product l is assigned in a better position than product m . Thus, by Corollary 1, the performance ranking must satisfy that,

$$\phi_l(q_l - \lambda(\phi)) \geq \phi_m(q_m - \lambda(\phi)). \quad (12)$$

However, the expected number of purchases $\lambda(\phi)$ is bounded from below and above by q_l and q_m respectively, i.e.,

$$q_l \leq \lambda(\phi) \leq q_m.$$

Hence $q_l - \lambda(\phi) \leq 0$, $q_m - \lambda(\phi) \geq 0$ and $q_m > q_l$, resulting in a contradiction since it violates Equation (12). Hence, condition (11) can never met when $|S| > 1$. \square

We now proceed to check the stability of the equilibrium points. The behavior of the solutions depend on σ , which itself depends on ϕ by Corollary 1. We will exploit this fact.

Theorem 2. *[Monopoly] Consider a trial-offer market under the performance ranking. Then the market converges almost surely to a monopoly for the product with the best quality.*

Proof. In order to use the ODE Method, we study the asymptotic behavior of the solutions of $\dot{x} = F(x)$, or equivalently

$$\dot{x}_i = F_i(x) = x_i(v_{\sigma_i} q_i - \sum_{j=1}^n v_{\sigma_j} q_j x_j), \quad \forall i \in \{1, \dots, n\}.$$

If $x_i \neq 0$, we can rewrite the previous equation as follows:

$$\frac{\dot{x}_i}{x_i} - v_{\sigma_i} q_i = - \sum_{j=1}^n v_{\sigma_j} q_j x_j,$$

where the right-hand-side of the equation is constant for every product. Therefore

$$\begin{aligned} \frac{\dot{x}_{i,t}}{x_{i,t}} - v_{\sigma_i} q_i &= \frac{\dot{x}_{k,t}}{x_{k,t}} - v_{\sigma_k} q_k \\ \Leftrightarrow \frac{d}{dt} [\ln(x_{i,t}) - v_{\sigma_i} q_i t] &= \frac{d}{dt} [\ln(x_{k,t}) - v_{\sigma_k} q_k t] \Rightarrow \\ \int_0^t \frac{d}{ds} [\ln(x_{i,s}) - v_{\sigma_i} q_i s] ds &= \int_0^t \frac{d}{ds} [\ln(x_{k,s}) - v_{\sigma_k} q_k s] ds \\ \Leftrightarrow \ln(x_{i,t}) - v_{\sigma_i} q_i t - \ln(x_{i,0}) &= \ln(x_{k,t}) - v_{\sigma_k} q_k t - \ln(x_{k,0}) \\ \Leftrightarrow \ln\left(\frac{x_{i,t}}{x_{k,t}}\right) &= t[v_{\sigma_i} q_i - v_{\sigma_k} q_k] + \ln\left(\frac{x_{i,0}}{x_{k,0}}\right) \end{aligned} \quad (13)$$

From any initial condition where the appeals are non-zero, no product will ever reach a market share of exactly one or zero. Hence, we analyze the behavior of the performance ranking when arbitrarily close to the equilibrium points. Define ϕ^i to be a small perturbation from e_i , i.e.,

$$\{\phi^i : \phi_i^i = 1 - \epsilon, \phi_j^i > 0, 1 \leq j \leq n, \sum_{k \neq i} \phi_k^i = \epsilon\},$$

where ϵ is an arbitrarily small positive quantity. The expected number of purchases for any ϕ^i is $\lambda(\phi^i) \approx q_i$. For $i = 1$, any perturbation of the market shares will slightly decrease the expected number of purchases, i.e., $\lambda(\phi^1) = q_1 - \delta$, $\delta > 0$. Then, $q_1 - \lambda(\phi^1) > 0$ and, for any $k \geq 2$, $q_k - \lambda(\phi^1) \leq 0$. By the condition in Corollary 1 and the fact that all $\phi_k^1 > 0$, the best quality product q_1 is assigned in the top slot, i.e., $\sigma_1 = 1$. Hence, $v_{\sigma_1} q_1 - v_{\sigma_k} q_k > 0$ for any $k \neq 1$.

Consider the process given by Equation (13) with initial condition $x_0 = \phi^1$. The process begins in the simplex, i.e., $0 < x_{i,0} < 1, \forall i$, and hence $\ln(\frac{x_{i,0}}{x_{k,0}})$ is bounded. Since the ranking does not change, the behavior of the solutions is given by the asymptotic behavior of $t[v_{\sigma_i} q_i - v_{\sigma_k} q_k]$ which depends on the sign of $v_{\sigma_i} q_i - v_{\sigma_k} q_k$. Taking $i = 1$ and $k \in \{2, \dots, n\}$, we have that $v_{\sigma_1} q_1 - v_{\sigma_k} q_k > 0$, and hence $t[v_{\sigma_1} q_1 - v_{\sigma_k} q_k] \rightarrow +\infty$ as $t \rightarrow +\infty$. Hence, $\ln(\frac{x_{1,t}}{x_{k,t}}) \rightarrow +\infty$ for all $k > 1$, and consequently $x_k(t) \rightarrow 0$. As $\sum_{i=1}^n x_{i,t} = 1$, $x_{1,t} \rightarrow 1$ and the equilibrium e_1 is stable.

For $i > 1$, consider a perturbation where some product $j : q_j > q_i$ has a small increase in its market such that the expected number of purchases increases very slightly, i.e., $\lambda(\phi^i) = q_i + \delta$, $\delta > 0$. Thus, for small δ , $q_j - \lambda(\phi^i) > 0$ and

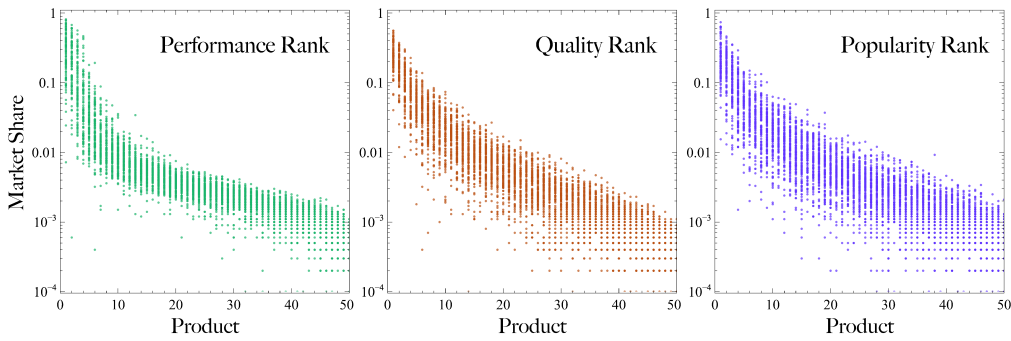


Figure 1: The Distribution of Market Shares. The products on the x-axis are ranked by decreasing quality from left to right. Each dot is the market share of a product in one of the 100 experiments. Note that the y-axis is in log scale.

$q_i - \lambda(\phi^i) = -\delta < 0$, which implies that $\phi_j^i(q_j - \lambda(\phi^i)) > \phi_i^i(q_i - \lambda(\phi^i))$. Therefore, by Corollary 1, product j is assigned in a better position than product i , i.e., $v_{\sigma_j} \geq v_{\sigma_i}$ and hence, $v_{\sigma_j}q_j - v_{\sigma_i}q_i > 0$. Consider the process given by Equation (13) with initial conditions $x_0 = \phi^i$, $i > 1$. At $t = 0$, as described above, it holds that $v_{\sigma_j}q_j - v_{\sigma_i}q_i > 0$ and consequently, at the next period of time $t = 1$, $\ln(\frac{x_{j,1}}{x_{i,1}}) > \ln(\frac{x_{j,0}}{x_{i,0}})$. This implies that $x_{j,1} > x_{j,0}$ or $x_{i,1} < x_{i,0}$, which, in either case, indicates that the new state is farther away from the initial state. Hence, e_i is an unstable equilibrium. \square

Theorem 2 states that, given any initial condition with non-zero appeals, the market eventually reaches to a monopoly for the product of highest quality, entailing the following.

Corollary 2. *The performance ranking is asymptotically optimal in trial-offer markets.*

5 Computational Experiments

We conclude by presenting computational results that illustrate and complement the theoretical results about the asymptotic behavior of the performance rank. We use an agent-based simulation and each simulation consists of N participants. For each participant at time t , the simulator randomly selects a product i according to the probabilities $p_i(\sigma, d)$, where σ is the ranking policy under evaluation and d is the social influence signal. Then, the simulator randomly determines, with probability q_i , whether selected product i is purchased; In the case of a purchase, the social influence signal for product i increases, i.e., $d_{i,t+1} = d_{i,t} + 1$. Otherwise, $d_{i,t+1} = d_{i,t}$. For every participant, a new list σ is computed using one of the ranking policies described above. We consider 50 songs (as in the MUSICLAB) and 40,000 participants. For every experiment, the product qualities q_i , product appeals a_i and the position visibilities v_i are independently drawn from a normal distribution with mean 0.5 and deviation 0.2. The values are then normalized to be between 0 and 1 and the appeals a_i are scaled by a factor of 50.

Figure 1 depicts computational results on the distribution of market shares under various ranking policies (in log scale). Each dot is the market share of a product in one of the 100 experiments. Figure 1 shows that the best product almost always receives the most purchases in the performance rank-

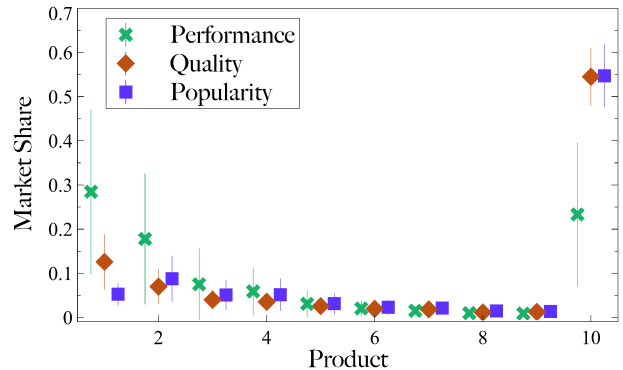


Figure 2: The average distribution of the top-10 market shares when the 10th best product starts with 500 purchases. The products are ranked by decreasing quality from left to right. Each dot is the average market share of a product over the 100 experiments. Error bars are the standard deviation.

ing. Quality ranking also performs well although the variance in its market shares is larger. The popularity ranking, while it shows the same overall correlation between quality and market share, exhibits many outliers. In terms of market efficiency, the performance ranking achieves 10% more purchases than the popularity ranking and 8% more than the quality ranking overall. For a single simulation, the performance ranking can achieve up to 23% more purchases than the other rankings.

Figure 2 illustrates the instability of sub-optimal monopolies and the benefits of the performance ranking to escape them. Each dot is the average market share product in one of the 100 experiments, where the 10th best product starts with 500 purchases, i.e., we start close to a monopoly for the 10th product. The figure shows that the performance ranking decreases the market share of the 10th product and significantly boosts the best product. This is not the case of the quality ranking at the completion of the experiments.

In conclusion, the theoretical and computational results indicate that the performance ranking has attractive properties for dynamic trial-offer markets. It is optimal and predictable asymptotically and it optimizes market efficiency at each time point. Computational results also show that it recovers from poor initial conditions much faster than the quality ranking.

Acknowledgement

Data61 (formerly known as NICTA) is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

- [Abeliuk *et al.*, 2015] A. Abeliuk, G. Berbeglia, M. Cebrian, and P. Van Hentenryck. The benefits of social influence in optimized cultural markets. *PLOS ONE*, 10(4), 2015.
- [Abeliuk *et al.*, 2016] A. Abeliuk, G. Berbeglia, M. Cebrian, and P. Van Hentenryck. Assortment optimization under a multinomial logit model with position bias and social influence. *4OR*, 14(1):57–75, 2016.
- [Aggarwal *et al.*, 2008] Gagan Aggarwal, Jon Feldman, S Muthukrishnan, and Martin Pál. Sponsored search auctions with markovian users. In *Internet and Network Economics*, pages 621–628. Springer, 2008.
- [Doob, 1953] Joseph L Doob. *Stochastic processes*, volume 101. New York Wiley, 1953.
- [Ghose *et al.*, 2012] Anindya Ghose, Panagiotis G Ipeirotis, and Beibei Li. Designing ranking systems for hotels on travel search engines by mining user-generated and crowd-sourced content. *Marketing Science*, 31(3):493–520, 2012.
- [Ghose *et al.*, 2014] Anindya Ghose, Panagiotis G Ipeirotis, and Beibei Li. Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science*, 60(7):1632–1654, 2014.
- [Higueras *et al.*, 2003] I. Higueras, J. Moler, F. Plo, and M. San Miguel. Urn models and differential algebraic equations. *Journal of applied probability*, pages 401–412, 2003.
- [Hogg and Lerman, 2015] Tad Hogg and Kristina Lerman. Disentangling the effects of social signals. *Human Computation Journal*, 2(2):189–208, 2015.
- [Kempe and Mahdian, 2008] David Kempe and Mohammad Mahdian. A cascade model for externalities in sponsored search. In *Internet and Network Economics*, pages 585–596. Springer, 2008.
- [Kleinberg, 2008] Jon Kleinberg. The convergence of social and technological networks. *Communications of the ACM*, 51(11):66–72, 2008.
- [Krumme *et al.*, 2012] Coco Krumme, Manuel Cebrian, Galen Pickard, and Sandy Pentland. Quantifying social influence in an online cultural market. *PloS one*, 7(5):e33785, 2012.
- [Lerman and Hogg, 2014] Kristina Lerman and Tad Hogg. Leveraging position bias to improve peer recommendation. *PLOS ONE*, 9(6):1–8, 2014.
- [Ljung, 1977] Lennart Ljung. Analysis of recursive stochastic algorithms. *Automatic Control, IEEE Transactions on*, 22(4):551–575, 1977.
- [Luce, 1965] D. Luce. *Individual Choice Behavior*. John Wiley and Sons, 1965.
- [Muchnik *et al.*, 2013] Lev Muchnik, Sinan Aral, and Sean J. Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146):647–651, 2013.
- [Renlund, 2010] Henrik Renlund. Generalized Polya Urns Via Stochastic Approximation. *ArXiv e-prints 1002.3716*, February 2010.
- [Salganik and Watts, 2008] Matthew J Salganik and Duncan J Watts. Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly*, 71(4):338–355, 2008.
- [Salganik and Watts, 2009] Matthew J Salganik and Duncan J Watts. Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in Cognitive Science*, 1(3):439–468, 2009.
- [Salganik *et al.*, 2006] Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854–856, 2006.
- [van de Rijt *et al.*, 2014] Arnout van de Rijt, Soong Moon Kang, Michael Restivo, and Akshay Patil. Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences*, 111(19):6934–6939, 2014.
- [Van Hentenryck *et al.*, 2016] Pascal Van Hentenryck, Andrés Abeliuk, Franco Berbeglia, Felipe Maldonado, and Gerardo Berbeglia. Aligning popularity and quality in online cultural markets. In *Tenth International AAAI Conference on Web and Social Media (ICWSM-16)*, 2016.