Aligning Popularity and Quality in Online Cultural Markets

Pascal Van Hentenryck
University of Michigan
pvanhent@umich.edu

Andrés Abeliuk
Data61 & University of Melbourne
aabeliuk@student.unimelb.edu.au

Franco Berbeglia
Carnegie Mellon University
fberbegl@andrew.cmu.edu

Felipe Maldonado
Data61 & Australian National University
felipe.maldonado@anu.edu.au

Gerardo Berbeglia
Melbourne Business School
g.berbeglia@mbs.edu

Abstract

Social influence is ubiquitous in cultural markets and plays an important role in recommendations for books, songs, and news articles to name only a few. Yet social influence is often presented in a bad light, often because it supposedly increases market unpredictability. Here we study a model of trial-offer markets, in which participants try products and later decide whether to purchase. We consider a simple policy which recovers product quality and ranks the products by quality when presenting them to market participants. We show that, in this setting, market efficiency always benefits from social influence. Moreover, we prove that the market converges almost surely to a monopoly for the product of highest quality, making the market both predictable and asymptotically optimal. Computational experiments confirm that the quality ranking policy quickly identifies "blockbusters", outperforms other policies, and is highly predictable.

Introduction

Social influence is ubiquitous in cultural markets. From book recommendations in `Amazon`, to song popularities in `iTunes`, and article rankings in the online version of the `New York Times` or the `Reddit and Hacker` site, social influence has become a critical aspect of the customer experience. Social influence may appear through different social signals such as the number of past purchases; consumer ratings; and/or consumer recommendations, depending on the market and/or platform. However, not all social signals are equally important. Indeed, two recent studies (Engstrom and Forsell 2014; Viglia, Furlan, and Ladrón-de Guevara 2014) conducted to understand the relative importance of the different social signals on consumer behavior in the `Android` app platform and in hotel selection experiment have arrived to the same conclusion, namely that the *popularity* signal (i.e., the number of purchases) has a much stronger impact on consumer behavior than the average consumer rating signal. Social influence is also reinforced by position bias (e.g., (Lerman and Hogg 2014)), as consumer preferences are also affected considerably by the visibility of the choices. In digital markets, the impact of visibility on consumer behavior has been widely observed in internet advertisement where sophisticated mathematical models have been developed to determine the relative importance of the different ads positions, in online stores such as `Amazon` and `iTunes`, and in online travel agents such as `Expedia` and `Orbitz` among others.

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Yet, despite its ubiquity, there is still considerable debate about the benefits of social influence and its effects on the market. Indeed, Salganik, Dodds, and Watts (2006) argued that social influence makes markets more unpredictable, providing an explanation about why prediction in cultural markets is a wicked problem. To investigate this hypothesis experimentally, they created an artificial music market called the MUSICLAB. Participants in the MUSICLAB were presented a list of unknown songs from unknown bands, each song being described by its name and band. The participants were divided into two groups exposed to two different experimental conditions: the *independent* condition and the *social influence* condition. In the first group (independent condition), participants were provided with no additional information about the songs. Each participant would decide which song to listen to from a random list. After listening to a song, the participant had the opportunity to download it. In the second group (social influence condition), each participant was provided with an additional information: The number of times the song was downloaded by earlier participants. Moreover, these participants were presented with a list ordered by the number of downloads. Additionally, to investigate the impact of social influence, participants in the second group were distributed in eight "worlds" evolving completely independently. In particular, participants in one world had no visibility about the downloads and the rankings in the other worlds. The MUSICLAB is a *trial-offer market* that provides an experimental testbed for measuring the unpredictability of cultural markets. By observing the evolution of different worlds given the same initial conditions, the MUSICLAB provides unique insights on the impact of social influence and the resulting unpredictability. In particular, Salganik et al suggested that social influence contributes to unpredictability, inefficiencies, and inequality of success, with follow-up experiments confirming these initial findings (Salganik and Watts 2009; Muchnik, Aral, and Taylor 2013; 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 leading to a Matthew effect. Unfortunately, popularity, which is easily distorted by noise in the process, is not a good proxy for quality: It leads to market unpredictability and even self-fulfilling prophecies, in which a perceived but initially false popularity becomes real over time (Salganik and Watts 2008).

In this paper, we reconsider this choice and study a setting in which product qualities are first recovered (using sampling and/or reinforcement learning) and then used to display products in decreasing order of quality: This policy reinforces the appeal of quality products with position bias. We investigate the quality ranking both computationally, using the generative model of the MUSICLAB proposed in (Krumme et al. 2012), and theoretically by modeling the trial-offer market as a discrete choice model based on a multinomial logit (Luce 1965) with social influence.

Our work is a step toward the understanding and development of expressive computational models for long-term effect of social influence (including unpredictability), an open question raised by Kleinberg (2008). Our main contributions can be summarized as follows:

1. The theoretical results show that the trial-offer market is optimal asymptotically when the quality ranking is used and converges almost surely. They also show that the market always benefits from position bias and social influence in expectation when the quality ranking is used.
2. The computational results show that the quality ranking under social influence significantly improves market efficiency, decreases unpredictability, and identifies "blockbusters". It provides significant improvements over the popularity ranking and is also comparable to or better than the performance ranking proposed in (Abeliuk et al. 2015a) that maximizes the probability of a download for each incoming participant.

Our results provide an interesting contrast with the conclusions of (Salganik, Dodds, and Watts 2006). The quality ranking aligns quality and popularity, making the market efficient and predictable. In other words, *it is not social influence per se that makes markets unpredictable: It is the way it is used that may lead to unpredictability and inefficiency.*

Trial-Offer Markets

This section introduces trial-offer markets in which participants can try a product before deciding to buy it. Such models are now pervasive in online cultural markets (e.g., books and songs). The market is composed of n products and each product $i \in \{1, \ldots, n\}$ 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 $\{1, \ldots, n\}$ 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 the product in position i of the list π, and σ_i 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, if we interpret this last action as an inefficiency, maximizing the expected efficiency of the market also minimizes unproductive trials. We also examine a number of questions about the market including (1) What is the best way to allocate the products to positions? (2) Is it beneficial to display a social signal, e.g., the number of past purchases, to customers? (3) Is the market predictable?

Static Market The probability of trying product i given a list σ is

$$p_i(\sigma) = \frac{v_{\sigma_i} A_i}{\sum_{j=1}^{n} v_{\sigma_j} A_j}$$

and the static market optimization problem consists of finding a ranking σ maximizing the expected number of purchases, i.e.,

$$\max_{\sigma \in S_n} \sum_{i=1}^{n} p_i(\sigma)\, q_i \tag{1}$$

where S_n represents the symmetry group over $\{1, \ldots, n\}$. 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.

Dynamic Market Our goal is to study a dynamic market where the appeal evolves over time according to a social influence signal. Given such a signal $d = (d_1, \ldots, d_n)$, the appeal of product i becomes $A_i + d_i$ and the probability of trying product i given ranking σ becomes

$$p_i(\sigma, d) = \frac{v_{\sigma_i}(A_i + d_i)}{\sum_{j=1}^{n} v_{\sigma_j}(A_j + d_j)}.$$

The dynamic market uses the number of purchases $d_{i,t}$ of product i at a time t as the social signal. However, other social signals such as the market share $\phi_i = (A_i + d_i)/(\sum_{i=1}^{n} A_i + d_i)$, used in online site such as `iTunes`, can easily be verified to be equivalent. Hence, in a dynamic market, the expected number of purchases of product i at time $t+1$ given list σ and purchase history d_t is $p_i(\sigma, d_t)\, q_i$. 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 with no social influence the purchase history plays no role and hence the market behaves as a static market.

Popularity as a Social Influence Signal As mentioned in the introduction, the benefits of popularity, i.e., the social influence signal used in this paper, have been validated experimentally. For instance, two recent studies (Engstrom and Forsell 2014; Viglia, Furlan, and Ladrón-de Guevara 2014) were conducted to understand the relative importance of the different social signals on consumer behavior. The first study (Engstrom and Forsell 2014) surveyed the downloads of more than 500,000 apps from the `Android` marketplace Google play, while the second study (Viglia, Furlan, and Ladrón-de Guevara 2014) conducted an online experiment ($n = 168$) where participants were asked to rank hotels based on the number of reviews and the average rating. Both studies arrived at the same conclusion, namely that the popularity signal (i.e., social preference) has a much stronger impact than the rating signal (i.e., a state preference). These two studies provide evidence that quantitative models with popularity as the social signal are of significant importance.

Popularity is Not Quality The experimental evidence from MUSICLAB suggests that the relationship between quality and popularity can be significantly distorted by social influence and position bias. These relationships have also been observed experimentally by Stoddard (Stoddard 2015), who studied the relationship between the intrinsic article quality and its popularity in the social news sites `Reddit and Hacker` news. The author proposed a Poisson regression model to estimate the demand for an article based on its quality, past views, and age among others. The results obtained after an estimation of each intrinsic article from these social news site showed that the most popular articles are typically the articles with the highest quality. Another study of social influence was carried out by (Tucker and Zhang 2011). The authors conducted a field experiment which showed that popularity information may benefit products with narrow appeal significantly more than those with a broad appeal. Along these lines, Sipos et al (2014) analyzed the voting behavior of users from Amazon product reviews when answering the question "Was this review helpful to you?" and how these votes relate to quality. The results showed that votes not only depend on the inherent quality of reviews, but also on the position where the review was presented in the ranked list. The authors also concluded that the ranking process converges and that the relative ordering of reviews stabilizes during the 4 months data was collected.

Our work presents a ranking method that mitigates the disparities between popularity and quality that emerge from social and position bias. A key feature of trial-offer markets is its decomposition into two stages, a sampling stage where participants decide which product to try followed by a buying stage where participants decide whether to buy or not the product sampled at the previous stage. Our results rely on the natural assumption that social influence and position biases have a greater effect on the decisions taken in the sampling stage than on the buying stage. Thus, popularity as proxy of quality is distorted by the noise of the first stage. *Our ranking policy uses a proxy for quality based only on the second stage, which can be interpreted as the posterior probability of buying an item given that it was sampled.*

Relationship to Online Advertising In online advertising, ads shown together compete for user attention and are affected both by position bias and the other ads displayed on the same page, which are called negative externalities in the literature (Gomes, Immorlica, and Markakis 2009; Jeziorski and Segal 2012). Ad auction is a well-studied scenario with negative externalities where the allocation of slots is assigned to ad bidders by an auctioneer (Kempe and Mahdian 2008; Aggarwal et al. 2008; Ghosh and Mahdian 2008; Cavallo and Wilkens 2014; Hummel and McAfee 2014). With the appearance of online social networks, recent work in optimal auction for a single good has also considered (positive) network externalities, where the utility of an individual consumer for the good increases with the number of network neighbors using the same good (Hartline, Mirrokni, and Sundararajan 2008; Haghpanah et al. 2013; Munagala and Xu 2014). Our trial-offer model considers both types of externalities, albeit in a different setting.

Rankings Policies

This section presents the ranking policies studied in this paper. In the following, without loss of generality, we assume that the qualities and visibilities are non-increasing, i.e., $q_1 \geq q_2 \geq \cdots \geq q_n$ and $v_1 \geq v_2 \geq \cdots \geq v_n$. We also assume that the qualities and visibilities are known. In practical situations, the product qualities are obviously unknown but we will show later in the paper that they can be recovered accurately and quickly, either before or during the market execution (Abeliuk et al. 2015a). We use $a_{i,t} = A_i + d_{i,t}$ to denote the appeal of product i under social influence at step t. When the step t is not relevant, we omit it and use a_i instead for simplicity. Finally, also for simplicity, we sometimes omit the range of indices in aggregate operators when they range over the products.

The policy studied in this paper is *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 $\sigma_i = i \quad (1 \leq i \leq n)$. The quality ranking contrasts with the *popularity ranking* which was used in (Salganik, Dodds, and Watts 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.

The *performance ranking* was proposed in (Abeliuk et al. 2015a) to show the benefits of social influence in cultural markets. The performance ranking 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 rank-

ing σ_t^* defined as

$$\sigma_t^* = \underset{\sigma \in S_n}{\text{arg-max}} \sum_{i=1}^{n} p_i(\sigma, d_k) \cdot q_i$$

where $d_t = (d_{1,t}, \ldots, d_{n,t})$ is the social influence signal at step k. The performance ranking uses the probability $p_{i,t}(\sigma)$ of trying products i at iteration t given ranking σ, as well as the quality q_i of product i. The performance ranking can be computed in strongly polynomial time and the resulting policy is scalable to large markets (Abeliuk et al. 2015a).

In the rest of this paper, Q-RANK, D-RANK, and P-RANK denote the policies using the quality, popularity, and performance rankings respectively. The policies are also annotated with SI or IN to denote whether they are used under the social influence or the independent condition. For instance, P-RANK(SI) denotes the policy that uses the performance ranking under the social influence condition, while P-RANK(IN) denotes the policy using the performance ranking under the independent condition. We also use RAND-RANK to denote the policy that simply presents a random order at each period. Under the independent condition, the optimization problem is the same at each iteration as mentioned earlier. Since the performance ranking maximizes the expected purchases at each iteration, it dominates all other policies in this setting (Abeliuk et al. 2015a).

Theoretical Analysis

This section presents a number of theoretical results on the quality ranking. In particular, it shows that the quality ranking always benefits from position bias and social influence, is an optimal and predictable policy asymptotically. All the proofs are in the appendix. For simplicity, the results assume that $q_1 > q_2 > \cdots > q_n$.

The Benefits of Position Bias We first show that position bias always increases the expected number of purchases when quality ranking is used.

Theorem 1. *Position bias increases the expected number of purchases under the quality-ranking policy, i.e., for all visibilities v_i, appeals a_i, and qualities q_i $(1 \leq i \leq n)$,*

$$\frac{\sum_{i\geq 1} v_i a_i q_i}{\sum_{j\geq 1} v_j a_j} \geq \frac{\sum_{i\geq 1} a_i q_i}{\sum_{j\geq 1} a_j}.$$

Note that not all ranking policies benefit from position bias: Theorem 1 exploits the properties of the quality ranking.

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 quality ranking policy, the expected marginal number of purchases in the studied trial-offer model increases when past purchases are revealed. This indicates that both social influence and position bias improve the market efficiency and their benefits are cumulative.

Theorem 2. *The expected rate of purchases is non-decreasing over time for the quality ranking under social influence.*

This result contrasts with the popularity ranking, under which Theorem 2 does not hold. A numerical example where the popularity ranking decreases the marginal expected number of purchases can be found in (Abeliuk et al. 2015b). Theorem 2 does hold under the performance ranking policy (Abeliuk et al. 2015a). However, the quality ranking presents two additional benefits: (1) its implementation does not require any knowledge about the appeal and visibility parameters of the model; (2) its asymptotic behavior is predictable and converges, as we show in the next section.

Asymptotic Behavior of the Quality Ranking We now prove the key result of the paper: The trial-offer market becomes a monopoly for the best product when the quality ranking is used at each step. As a consequence, the quality 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 using the quality ranking. The proof needs the following lemma that characterizes the probability that the next purchase is product i.

Lemma 1. *The probability p_i that the next purchase (after any number of steps) is product i is*

$$p_i = \frac{v_i a_i q_i}{\sum_{j=1}^{n} v_j a_j q_j}.$$

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

$$X_{i,t} \doteq a_{i,t} \hat{q}_i \tag{2}$$

$$Z_{i,t} \doteq \frac{X_{i,t}}{\sum_{i=1}^{n} X_{i,k}} \tag{3}$$

where $\hat{q}_j = v_j q_j$. 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 $X_{i,t+1} = X_{i,t} + \hat{q}_i$ if product i is purchased, the generalized Pólya scheme add $\hat{q}_i$ balls, each time product i is purchased. As a result, the Pólya scheme uses the replacement matrix

$$\mathbf{R} = \begin{pmatrix} \hat{q}_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \hat{q}_n \end{pmatrix}. \tag{4}$$

where $\hat{q_1} \geq \ldots \geq \hat{q_n}$. Higueras 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

$$Z_{i+1} = Z_i + \gamma_{i+1} \left(F(Z_i) + \epsilon_{i+1} + \beta_{i+1} \right), \tag{5}$$

where: (1) the process $\{\gamma_i\}$ is a decreasing sequence of positive 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 $\beta_i \to 0$ almost

surely; (4) $F(x) = xR(I - 1^T x)$, where R is the replacement matrix of the Pólya scheme. The ODE method (Ljung 1977) relates the recurrence Equation 5 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 are obtained by solving $F(x) = 0$ when $\sum_{i=1}^{n} x_i = 1$.

Lemma 2. *Let $F(x) = xR(I - 1^T x)$. Then, the solutions to $F(x) = 0$ when $\sum_{i=1}^{n} x_i = 1$ are the set $\{e_i : 1 \leq i \leq n\}$, where e_i denotes the unit vector whose i^{th} entry is 1.*

We now proceed to check the stability of the equilibrium points. In order to use the ODE Method, we study the asymptotic behavior of the solutions of $\dot{x} = F(x)$.

Theorem 3 (Monopoly of Markets)**.** *Consider a trial-offer market where $\hat{q_1} > \hat{q_2} > \cdots > \hat{q_n}$. Then the market converges almost surely to a monopoly for product 1.*[1]

This result states that, starting from any initial condition where the appeals are non-zero, the market eventually reaches the equilibrium that corresponds to a monopoly for the product of highest quality. This result also implies that the quality ranking is optimal asymptotically, since only the best product is left.

Corollary 1. *The quality ranking is asymptotically optimal in trial-offer markets.*

Static Behavior of the Quality Ranking The quality ranking also has some performance guarantees at each step, which only depends on the visibility coefficients. We define an α-approximation to be an algorithm producing a solution that is guaranteed to be within a factor α of the optimum algorithm.

Theorem 4 (Static Performance Bound)**.** *The quality ranking is an α-approximation of the static market optimization problem, where $\alpha = v_1/v_n$. Moreover, the α-approximation of the quality ranking is tight.*

Under the experimental setting of the MUSICLAB, the (worst-case) approximation factor of the quality ranking is $\alpha = v_1/v_n = 0.8/0.2 = 4$.

Computational Experiments

We now report computational results that illustrate and complement the theoretical analysis presented in the previous section. The computational results use settings that model the MUSICLAB experiments discussed in (Salganik, Dodds, and Watts 2006; Krumme et al. 2012; Abeliuk et al. 2015a). As mentioned in the introduction, the MUSICLAB is a trial-offer market where participants can try a song and then decide to download it. The generative model of the MUSICLAB (Krumme et al. 2012) is the model of consumer choice with social influence described earlier.

[1]When there are products with the same highest quality, the market share of such a product converges almost surely to a random variable following a beta-distribution.

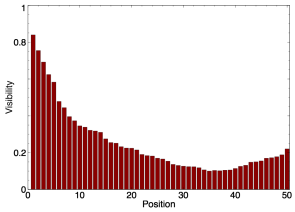


Figure 1: The visibility v_p (y-axis) of position p in the song list (x-axis) where $p = 1$ is the top position and $p = 50$ is the bottom position of a single column display.

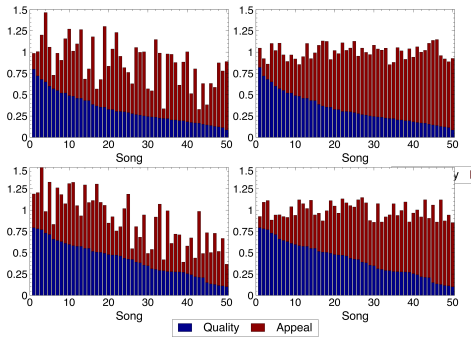


Figure 2: The quality q_i (blue) and appeal A_i (red) of song i in the four settings. In the first setting (top left), the qualities and appeals were chosen independently according to a Gaussian distribution. The second setting (top right) explores an extreme case where the appeal is negatively correlated with the quality used in setting 1. In the third setting (bottom left), the qualities and appeals were chosen independently according to a uniform distribution. The fourth setting (bottom right) explores an extreme case where the appeal is negatively correlated with the quality from setting 3.

The Experimental Setting The experimental setting uses an agent-based simulation to emulate the MUSICLAB. Each simulation consists of N iterations and, at each iteration t,

1. the simulator randomly selects a song i according to the probabilities $p_i(\sigma, d)$, where σ is the ranking policy under evaluation and d is the social influence signal;
2. the simulator randomly determines, with probability q_i, whether selected song i is downloaded; In the case of a download, the simulator increases the social influence signal for song i, i.e., $d_{i,t+1} = d_{i,t} + 1$. Otherwise, $d_{i,t+1} = d_{i,t}$.

Every r iterations, a new list σ is computed using one of the ranking policies described above. For instance, in the social influence condition of the original MUSICLAB experiments, the policy ranks the songs by popularity, i.e., the D-RANK policy which ranks the songs in decreasing order of download counts. The parameter $r \geq 1$ is called the refresh rate. The experimental setting, which aims at being close to the MUSICLAB experiments, considers 50 songs and simulations with 20,000 steps. The songs are displayed in a single column. Figure 1 depicts the visibility parameters used in all

computational experiments. The visibility profile is based on the analysis in (Krumme et al. 2012), indicating that participants are more likely to try songs higher in the list. More precisely, the visibility decreases with the list position, except for a slight increase at the bottom positions.

The paper also uses four settings for the quality and appeal of each product, which are depicted in Figure 2. In the first setting, the quality and the appeal were chosen independently according to a Gaussian distribution normalized to fit between 0 and 1. The second setting explores an extreme case where the appeal is negatively correlated with quality. The quality of each product is the same as in the first setting but the appeal is chosen such that the sum of appeal and quality is 1 plus a normally distributed noise. In the third setting, the quality and the appeal were chosen independently according to a uniform distribution. The fourth setting also explores an extreme case where the appeal is negatively correlated with quality. The quality of each product is the same as in the third setting but the appeal is chosen such that the sum of appeal and quality is exactly 1. The results were obtained by averaging the results of $W = 400$ simulations.

Recovering the Songs Quality We now show how to recover songs quality in the MUSICLAB. The key idea is borrowed from Salganik et al. (2006) who stated that *the popularity of a song in the independent condition* is a natural measure of its quality and captures both its intrinsic "value" and the preferences of the participants. Expanding on their idea, the popularity of a song in the independent condition and with no position bias is a natural measure of its quality. *However, under social influence, popularity may no longer reflect quality and may be strongly influenced by the visibility and early downloads.*

To approximate the quality of a song, it suffices to sample the participants in an independent world. This can be simulated by using a Bernoulli sampling based on the real quality of the songs. The predicted quality $\hat{q}_i$ of song i is obtained by running m independent Bernoulli trials with probability q_i of success, i.e., $\hat{q}_i = \frac{k}{m}$, where k is the number of successes over the m trials. For a large enough sampling size, $\hat{q}_i$ has a mean of q_i and a variance of $q_i(1 - q_i)$. This variance has the desirable property that the quality of a song with a more 'extreme' quality (i.e., a good or a bad song) is recovered faster than those with average quality. In addition, we can merge information about downloads into the prediction as the market with social influence proceeds: At step k, the approximate quality of song i is given by $\hat{q}_{i,k} = \frac{\hat{q}_{i,0} \cdot m + d_{i,k}}{m + s_{i,k}}$, where m is the initial sample size, $d_{i,k}$ and $s_{i,k}$ are the number of downloads and samplings of song i up to step k.

Figure 3 presents experimental results about the accuracy of the quality approximation for two rankings, assuming an initial independent sampling set of size 10 per song. More precisely, the figure reports the average squared difference between the song qualities and their predictions under the social influence and the independent conditions. In all cases, the results indicate that song qualities are recovered quickly and accurately. Note also that the Q-RANK only requires an ordinal ordering of the qualities, not their exact values.

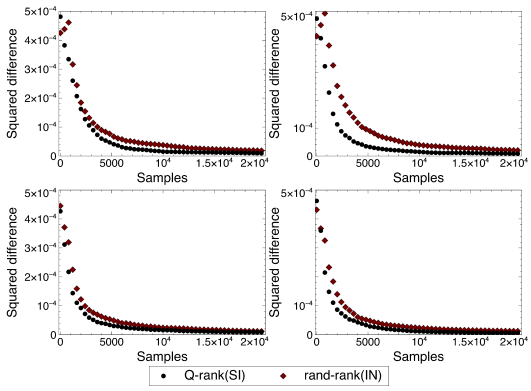


Figure 3: Average Squared Difference of Inferred Quality over Time for Different Rankings for the top 10 quality songs. The figure reports the average squared difference $\sum_{i=1}^{n} \frac{(\hat{q}_{i,k} - q_i)^2}{n}$ between the song quality and their predictions for the quality ranking under social influence and the random ranking in the independent condition. The figure shows the four settings in clockwise direction from the top-left plot. The quality of each song was initially approximated with 10 Bernoulli trials.

Performance of the Market Figure 4 depicts computational results on the expected number of downloads for the various rankings and settings and reveals two findings:

1. The quality ranking exhibits a similar performance to the performance ranking and provides substantial gains in expected downloads compared to the popularity and random rankings. On settings with negative correlations between appeal and quality, the quality ranking performs better than the performance ranking.
2. The benefits of social influence and position bias are complementary and cumulative. Both are significant in terms of the expected performance of the market.

Predictability of the Market Figures 5 and 6 depict computational results on the predictability of the market under various ranking policies. The figures plot the number of downloads of each song for the 400 experiments. In the plots, the songs are ranked by increasing quality from left to right on the x-axis. Each dot in the plot shows the number of downloads of a song in one of the 400 experiments. Figures 5 and 6 present the result for the first and second settings.

The computational results are compelling. Figure 5 shows that the best song always receives the most downloads in the quality ranking (with social influence) and that the variance in its number of downloads across the experiments is very small. The performance ranking (with social influence) also performs well although the variance in its downloads is larger. The popularity ranking is highly unpredictable, while the random ranking is highly predictable as one would expect. It is also interesting to note that these observations continue to hold even when the appeal is negatively correlated with quality, as Figure 6 indicates. The contrast between the popularity ranking used in (Salganik, Dodds, and

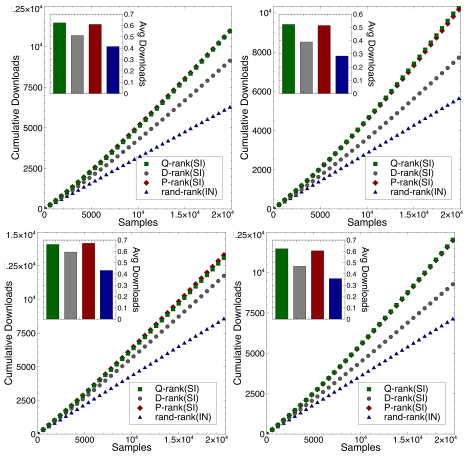


Figure 4: The number of downloads over time for the various rankings. The x-axis represents the number of product trials and the y-axis represents the average number of downloads over all experiments. On the upper left corner of each graph, the bar plot depicts the average number of purchases per try for all rankings. The results for the four settings are shown in clockwise direction starting from the top-left figure.

Watts 2006) and the quality ranking is particularly striking.

Conclusion

This paper studied trial-offer cultural markets, which are ubiquitous in our societies and involve products such as books, songs, videos, clothes, and even newspaper articles. In these markets, participants are presented with products in a certain ranking. They can then try the products before deciding whether to purchase them. Social influence signals are widely used in such settings and help promote popular products to maximize market efficiency. However, it has been argued that social influence makes these markets unpredictable (Salganik, Dodds, and Watts 2006). As a result, social influence is often presented in a negative light.

In this paper, we have reconsidered this conventional wisdom. We have shown that, when products are presented to participants in a way that reflects their true quality, the market is both efficient and predictable. In particular, a quality ranking tends to a monopoly for the product of highest quality, making the market both optimal and predictable asymptotically. Moreover, we have shown that both social influence and position bias improve market efficiency. These results are robust and do not depend on the particular values for the appeal and quality of the products. In addition, computational experiments using the generative model of the MUSICLAB show that there is a fast convergence to our asymptotic theoretical results. These results are an interesting contrast with the popularity ranking studied in (Salganik, Dodds, and Watts 2006), where the market is indeed unpredictable and less efficient.

There are some important lessons to draw from these results. On the one hand, it appears that *unpredictability is not an inherent property of social influence*: Whether a market

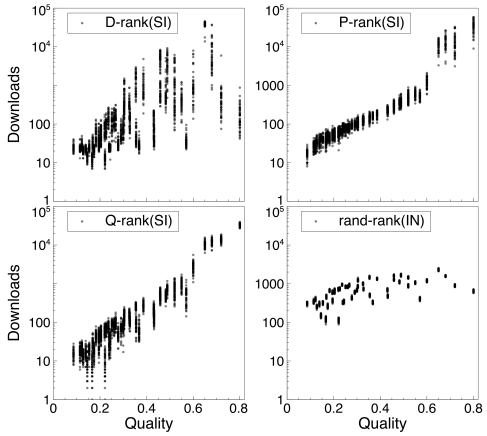


Figure 5: The Distribution of Downloads Versus Song Qualities (First Setting). The songs on the x-axis are ranked by increasing quality from left to right. Each dot is the number of download of a product in one of the 400 experiments.

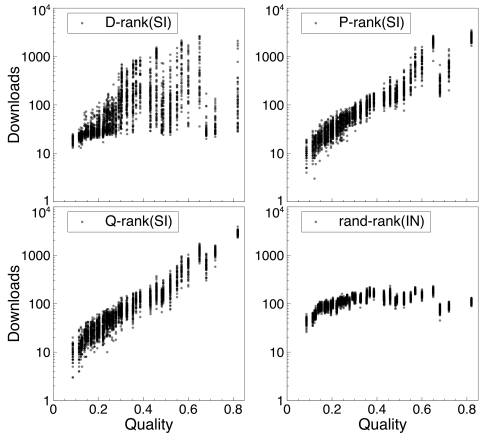


Figure 6: The Distribution of Download Versus Song Qualities (Second Setting). The songs on the x-axis are ranked by increasing quality from left to right. Each dot is the number of downloads of a product in one of the 400 experiments.

is predictable or not really depends on how social influence is used. On the other hand, with the quality ranking, computational experiments also show that "blockbusters" are quickly identified, even when the appeal is negatively correlated with quality. In addition, a high-quality product will overcome a poor appeal but the opposite does not hold.

It is also interesting to contrast our results with the study in (Ceyhan, Mousavi, and Saberi 2011) which uses the MNL model for consumer choice preferences but without position bias and with products of the same appeal. In their model, they incorporate a parameter J that measures the social influence intensity and show that when J is large, eventually a monopoly for some random product (depending on the initial conditions and the early downloads) will occur. This comes from the fact that the social influence signal is much stronger in this model than in ours where the social influence

intensity is set to $J = 1$. Recall also that the choice model used in this paper was shown to reproduce the original experiments of the MUSICLAB (Krumme et al. 2012).

Overall, these results show that the quality ranking makes it possible to align popularity and quality in trial-offer markets by using the quality ranking. This alignment makes the market optimal asymptotically and predictable. In contrast, using popularity as a proxy for quality in the ranking makes the market unpredictable and potentially inefficient as was shown in (Salganik, Dodds, and Watts 2006).

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.

Appendix: Proofs

Proof of Theorem 1. Let $\lambda = \frac{\sum_i v_i a_i q_i}{\sum_j v_j a_j}$ be the expected number of purchases for the quality ranking. We have

$$\sum_i v_i a_i \left(q_i - \lambda\right) = 0.$$

Consider the index k such that $(q_k - \lambda) \geq 0$ and $(q_{k+1} - \lambda) < 0$. Since $v_1 \geq \ldots \geq v_n$, we have

$$\sum_{i=1}^{k} v_k a_i (q_i - \lambda) + \sum_{i=k+1}^{n} v_k a_i (q_i - \lambda) \leq \sum_i v_i a_i (q_i - \lambda) = 0$$

and, since $v_k \geq 0$, $\sum_{i=1}^{n} a_i \left(q_i - \lambda\right) \leq 0$. It follows that $\lambda \geq \frac{\sum_{i=1}^{n} a_i q_i}{\sum_{i=1}^{n} a_i}$. □

Proof of Theorem 2. Let $\mathbb{E}[D_t] = \frac{\sum_i v_i a_i q_i}{\sum_i v_i a_i} = \lambda$ denote the expected number of purchases at time t. The expected number of purchases in time $t+1$ conditional to time t is

$$\mathbb{E}[D_{t+1}] = \sum_j \left(\frac{v_j a_j q_j}{\sum v_i a_i} \frac{\sum_{i \neq j} v_i a_i q_i + v_j (a_j + 1) q_j}{\sum_{i \neq j} v_i a_i + v_j (a_j + 1)} \right) + \left(1 - \frac{\sum_i v_i a_i q_i}{\sum_i v_i a_i}\right) \left(\frac{\sum_i v_i a_i q_i}{\sum_i v_i a_i}\right)$$

$$= \sum_j \left(\frac{v_j a_j q_j}{\sum v_i a_i} \frac{\sum_i v_i a_i q_i + v_j q_j}{\sum_i v_i a_i + v_j} \right) + \left(1 - \frac{\sum_j v_j a_j q_j}{\sum_i v_i a_i}\right) \lambda.$$

We need to prove that

$$\mathbb{E}[D_{t+1}] \geq \mathbb{E}[D_t], \tag{6}$$

which amounts to showing that

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

which reduces to proving

$$\frac{1}{\sum_i v_i a_i} \sum_j \left[\frac{v_j^2 a_j q_j}{\sum_i v_i a_i + v_j} \left(q_j - \lambda\right) \right] \geq 0$$

or, equivalently,

$$\sum_j \left[\frac{v_j^2 a_j q_j}{\sum_i v_i a_i + v_j} \left(q_j - \lambda\right) \right] \geq 0. \tag{7}$$

Let $k = \max\{i \in N | (q_i - \lambda) \geq 0\}$, i.e., the largest index $k \in N$ such that $q_k \geq \lambda$. We have

$$\sum_{j=1}^{n} \left[\frac{v_j^2 a_j q_j \left(q_j - \lambda\right)}{\sum_i v_i a_i + v_j} \right] = \sum_{j=1}^{k} \left[\frac{a_j v_j v_j q_j}{\sum_i v_i a_i + v_j} \left(q_j - \lambda\right) \right] + \sum_{j=k+1}^{n} \left[\frac{a_j v_j v_j q_j}{\sum_i v_i a_i + v_j} \left(q_j - \lambda\right) \right].$$

By definition of k, the terms in the summation on the left are positive and the terms in the summation on the right are negative. Moreover, for any $c > 0$ and $v_i, v_j \geq 0$, we have

$$\frac{v_i}{c + v_i} \geq \frac{v_j}{c + v_j} \Leftrightarrow (c + v_j) v_i \geq (c + v_i) v_j$$
$$\Leftrightarrow c v_i \geq c v_j \Leftrightarrow v_i \geq v_j.$$

Since $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$,

$$\frac{v_1}{\sum_i v_i a_i + v_1} \geq \frac{v_2}{\sum_i v_i a_i + v_2} \geq \ldots \geq \frac{v_n}{\sum_i v_i a_i + v_n}. \tag{8}$$

Moreover, since the quality ranking orders the products by quality and $q_1 \geq q_2 \geq \ldots \geq q_n \geq 0$, Equation (8) and the definition of k implies that

$$\forall i \leq k: \ \frac{v_i q_i}{\sum_j v_j a_j + v_i} \geq \frac{v_k q_k}{\sum_j v_j a_j + v_k},$$
$$\forall i > k: \ \frac{v_i q_i}{\sum_j v_j a_j + v_i} \leq \frac{v_k q_k}{\sum_j v_j a_j + v_k}.$$

This observation, together with the fact that the left-hand (resp. right-hand) terms are positive (resp. negative), produces a lower bound to the right-hand side of Inequality (7):

$$\sum_{j=1}^{k} \frac{a_j v_j v_j q_j}{\sum_i v_i a_i + v_j} \left(q_j - \lambda\right) + \sum_{j=k+1}^{n} \frac{a_j v_j v_j q_j}{\sum_i v_i a_i + v_j} \left(q_j - \lambda\right)$$
$$\geq \frac{v_k q_k}{\sum_i v_i a_i + v_k} \left(\sum_{j=1}^{k} a_j v_j (q_j - \lambda) + \sum_{j=k+1}^{n} a_j v_j (q_j - \lambda) \right).$$

Now, by definition of λ,

$$\lambda = \frac{\sum_{i=1}^{n} v_i a_i q_i}{\sum_{i=1}^{n} v_i a_i} \Leftrightarrow \lambda \sum_{i=1}^{n} v_i a_i = \sum_{i=1}^{n} v_i a_i q_i$$
$$\Leftrightarrow \sum_{i=1}^{n} v_i a_i (q_i - \lambda) = 0.$$

which implies that

$$\frac{v_k q_k}{\sum_i v_i a_i + v_k} \sum_{j=1}^{n} \left[a_j v_j \left(q_j - \lambda\right) \right] = 0$$

concluding the proof. □

Proof of Lemma 1. The probability that product i is purchased in the first step is given by

$$p_i^{1st} = \frac{v_i a_i}{\sum_{j=1}^{n} v_j a_j} q_i.$$

The probability that product i is purchased in the second step and no product was purchased in the first step is given by

$$p_i^{2nd} = \left(\frac{\sum_{j=1}^{n} v_j a_j (1 - q_j)}{\sum_{j=1}^{n} v_j a_j} \right) \frac{v_i a_i}{\sum_{j=1}^{n} v_j a_j} q_i.$$

More generally, the probability that product i is purchased in step m while no product was purchased in earlier steps is

$$p_i^{mth} = \left(\frac{\sum_{j=1}^{n} v_j a_j (1 - q_j)}{\sum_{j=1}^{n} v_j a_j} \right)^{m-1} \frac{v_i a_i}{\sum_{j=1}^{n} v_j a_j} q_i.$$

Defining $a = (\sum_{j=1}^{n} v_j a_j q_j)/(\sum_{j=1}^{n} v_j a_j)$, we have

$$p_i^{mth} = \Big(1 - a\Big)^{m-1} \frac{v_i a_i}{\sum_{j=1}^{n} v_j a_j} q_i.$$

Hence the probability that the next purchased product is product i is given by

$$p_i = \sum_{m=1}^{\infty} \Big(1 - a\Big)^{m-1} \frac{v_i a_i}{\sum_{j=1}^{n} v_j a_j} q_i.$$

The result follows from

$$\sum_{m=1}^{\infty} \Big(1 - a\Big)^{m-1} = \frac{1}{a}.$$

□

Proof of Lemma 2.

$$\begin{aligned} F(x) &= x \left[R - \begin{pmatrix} \hat{q}_1 \\ \vdots \\ \hat{q}_n \end{pmatrix} x \right] \\ &= \left[x_1(\hat{q}_1 - \sum_{i=1}^{n} \hat{q}_i x_i), \ldots, x_n(\hat{q}_n - \sum_{i=1}^{n} \hat{q}_i x_i) \right]. \end{aligned}$$

Hence, an equilibrium point must satisfy for all i,

$$x_i \left(\hat{q}_i - \sum_{j=1}^{n} \hat{q}_j x_j \right) = 0. \tag{9}$$

A point $u \in \{e_i : 1 \leq i \leq n\}$ is a trivial solution to Equation 9. Since $\hat{q}_i \neq \hat{q}_j$ for any $i \neq j$, such points are the only solutions: Indeed, if $x_i > 0$ and $x_j > 0$ for $i \neq j$, then Equation 9 states that $\hat{q}_i = \hat{q}_j$ (since the sum in Equation 9 is the same for all products), which violates our assumption. □

Proof of Theorem 3. We study the asymptotic behavior of the solutions of $\dot{x} = F(x)$, or equivalently

$$\dot{x}_i = F_i(x) = x_i(\hat{q}_i - \sum_{j=1}^{n} \hat{q}_j x_j), \quad \forall i \in \{1, ..., n\}.$$

If $x_i \neq 0, \forall i \in \{1, ..., n\}$, we can rewrite the previous equation as follows:

$$\frac{\dot{x}_i}{x_i} - \hat{q}_i = -\sum_{j=1}^{n} \hat{q}_j x_j,$$

where the right-hand-side of the equation is the same for every product. Hence,

$$\begin{aligned} \frac{\dot{x}_{i,t}}{x_{i,t}} - \hat{q}_i &= \frac{\dot{x}_{k,t}}{x_{k,t}} - \hat{q}_k, \quad \forall i, k \\ \Leftrightarrow \frac{d}{dt}[\log(x_{i,t}) - \hat{q}_i t] &= \frac{d}{dt}[\log(x_{k,t}) - \hat{q}_k t] \\ \Rightarrow \int_0^t [\frac{d}{ds}[\log(x_{i,s}) - \hat{q}_i s]] ds &= \int_0^t [\frac{d}{ds}[\log(x_{k,s}) - \hat{q}_k s]] ds \\ \Leftrightarrow \log(x_{i,t}) - \hat{q}_i t - \log(x_{i,0}) &= \log(x_{k,t}) - \hat{q}_k t - \log(x_{k,0}) \\ \Leftrightarrow \log(\frac{x_{i,t}}{x_{k,t}}) &= t[\hat{q}_i - \hat{q}_k] + \log(\frac{x_{i,0}}{x_{k,0}}) \end{aligned} \tag{10}$$

Now, as the process begins inside of the simplex (i.e., $0 < x_{i,0} < 1$), $\log(\frac{x_{i,0}}{x_{k,0}})$ is bounded. In consequence, the behavior of the solutions is given by the asymptotic behavior of $t[\hat{q}_i - \hat{q}_k]$ which depends of the sign of $\hat{q}_i - \hat{q}_k$. Since $\hat{q_1} > \hat{q_2} > \cdots > \hat{q_n}$, taking $i = 1$, $k \in \{2, ..., n\}$ in Equation (10) yields $t[\hat{q}_1 - \hat{q}_k] \to +\infty$ as $t \to +\infty$. Hence $\log(\frac{x_{1,t}}{x_{k,t}}) \to +\infty$ for all $k > 1$, and consequently $x_k(t) \to 0$. Since $\sum_{i=1}^{n} x_{i,t} = 1$, $x_{1,t} \to 1$, i.e., the market converges to a monopoly for the highest-quality product. □

Proof of Theorem 4. Let σ^* be the optimal sorting and λ^* its expected number of purchases. We have

$$\lambda^* = \frac{\sum_i v_{\sigma_i^*} a_i q_i}{\sum_j v_{\sigma_j^*} a_j} \leq \frac{\sum_i v_1 a_i q_i}{\sum_j v_n a_j} = \alpha \frac{\sum_i a_i q_i}{\sum_j a_j}.$$

Let λ^q be the expected number of purchases for the quality ranking, i.e.,

$$\lambda^q = \frac{\sum_i v_i a_i q_i}{\sum_j v_j a_j}.$$

By Theorem 1,

$$\lambda^q \geq \frac{\sum_i a_i q_i}{\sum_j a_j}.$$

Combining both bounds yields $\lambda^* \leq \alpha \lambda^q$.

We now show that the approximation is tight. Consider 3 products with qualities $q_1 = 1, q_2 = \epsilon, q_3 = 0$ and appeals $a_1 = 1, a_2 = x, a_3 = 0$ and let the visibilities be $v_1 = 1, v_2 = 1, v_3 < 1$. The quality ranking is $\sigma^q = (1, 2, 3)$ and the optimal performance ranking is $\sigma^* = (1, 3, 2)$. The expected number of purchases for the quality ranking is $\frac{1+\epsilon x}{1+x}$ while it is $\frac{1+\epsilon x \alpha}{1+\alpha x}$. for the performance ranking. When ϵ tends

to zero, the ratio between the performance and quality ranking becomes

$$\lim_{\epsilon \to 0} \frac{1 + v_3 \epsilon x}{1 + v_3 x} \frac{1 + x}{1 + \epsilon x} = \frac{1 + x}{1 + v_3 x}.$$

Hence, when x is large enough, the ratio is approximately α:

$$\frac{1 + x}{1 + v_3 x} \approx \frac{1}{v_3} = \frac{v_1}{v_3} = \alpha.$$

□

References

Abeliuk, A.; Berbeglia, G.; Cebrian, M.; and Van Hentenryck, P. 2015a. The benefits of social influence in optimized cultural markets. *PLOS ONE* 10(4).

Abeliuk, A.; Berbeglia, G.; Cebrian, M.; and Van Hentenryck, P. 2015b. Assortment optimization under a multinomial logit model with position bias and social influence. *4OR* 1–19.

Aggarwal, G.; Feldman, J.; Muthukrishnan, S.; and Pál, M. 2008. Sponsored search auctions with markovian users. In *Internet and Network Economics*. Springer. 621–628.

Cavallo, R., and Wilkens, C. A. 2014. Gsp with general independent click-through-rates. In *Web and Internet Economics*. Springer. 400–416.

Ceyhan, S.; Mousavi, M.; and Saberi, A. 2011. Social Influence and Evolution of Market Share. *Internet Mathematics* 7(2):107–134.

Engstrom, P., and Forsell, E. 2014. Demand effects of consumers' stated and revealed preferences. *Available at SSRN 2253859*.

Ghosh, A., and Mahdian, M. 2008. Externalities in online advertising. In *Proceedings of the 17th international conference on World Wide Web*, 161–168. ACM.

Gomes, R.; Immorlica, N.; and Markakis, E. 2009. Externalities in keyword auctions: An empirical and theoretical assessment. In *WINE*, volume 9, 172–183. Springer.

Haghpanah, N.; Immorlica, N.; Mirrokni, V.; and Munagala, K. 2013. Optimal auctions with positive network externalities. *ACM Transactions on Economics and Computation* 1(2):13.

Hartline, J.; Mirrokni, V.; and Sundararajan, M. 2008. Optimal marketing strategies over social networks. In *Proceedings of the 17th international conference on World Wide Web*, 189–198. ACM.

Higueras, I.; Moler, J.; Plo, F.; and San Miguel, M. 2003. Urn models and differential algebraic equations. *Journal of applied probability* 401–412.

Hummel, P., and McAfee, R. P. 2014. Position auctions with externalities. In *Web and Internet Economics*. Springer. 417–422.

Jeziorski, P., and Segal, I. R. 2012. What makes them click: Empirical analysis of consumer demand for search advertising. *Available at SSRN 1417625*.

Kempe, D., and Mahdian, M. 2008. A cascade model for externalities in sponsored search. In *Internet and Network Economics*. Springer. 585–596.

Kleinberg, J. 2008. The convergence of social and technological networks. *Communications of the ACM* 51(11):66–72.

Krumme, C.; Cebrian, M.; Pickard, G.; and Pentland, S. 2012. Quantifying social influence in an online cultural market. *PloS one* 7(5):e33785.

Lerman, K., and Hogg, T. 2014. Leveraging position bias to improve peer recommendation. *PLOS ONE* 9(6):1–8.

Ljung, L. 1977. Analysis of recursive stochastic algorithms. *Automatic Control, IEEE Transactions on* 22(4):551–575.

Luce, D. 1965. *Individual Choice Behavior*. John Wiley and Sons.

Muchnik, L.; Aral, S.; and Taylor, S. J. 2013. Social influence bias: A randomized experiment. *Science* 341(6146):647–651.

Munagala, K., and Xu, X. 2014. Value-based network externalities and optimal auction design. In *Web and Internet Economics*. Springer. 147–160.

Renlund, H. 2010. Generalized Polya Urns Via Stochastic Approximation. *ArXiv e-prints 1002.3716*.

Salganik, M. J., and Watts, D. J. 2008. Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly* 71(4):338–355.

Salganik, M. J., and Watts, D. J. 2009. Web-based experiments for the study of collective social dynamics in cultural markets. *Topics in Cognitive Science* 1(3):439–468.

Salganik, M. J.; Dodds, P. S.; and Watts, D. J. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762):854–856.

Sipos, R.; Ghosh, A.; and Joachims, T. 2014. Was this review helpful to you?: it depends! context and voting patterns in online content. In *Proceedings of the 23rd international conference on World wide web*, 337–348. ACM.

Stoddard, G. 2015. Popularity and quality in social news aggregators: A study of reddit and hacker news. In *Proceedings of the 24th International Conference on World Wide Web Companion*, 815–818.

Tucker, C., and Zhang, J. 2011. How does popularity information affect choices? a field experiment. *Management Science* 57(5):828–842.

van de Rijt, A.; Kang, S. M.; Restivo, M.; and Patil, A. 2014. Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences* 111(19):6934–6939.

Viglia, G.; Furlan, R.; and Ladrón-de Guevara, A. 2014. Please, talk about it! when hotel popularity boosts preferences. *International Journal of Hospitality Management* 42:155–164.